THE PRAGMATICS OF REPETITION, *EMPHASIS*, AND INTENSIFICATION

Rebecca Claire Jackson

Submitted for consideration of the award of PhD

University of Salford

School of Arts and Media

Salford, UK

2016

Table of Contents

Tables and Figures

Acknowledgements

Many thanks to my supervisor Diane Blakemore, without whom this PhD would not have been possible. I thank you for encouragement, inspiration, coffee, and for the most engaging and challenging supervisions. We have come a long way since the first day I sat on the front row of your pragmatics class. In research, I have now found an outlet for all my questions! In teaching, I have found a vocation. I will be forever grateful that you showed me the way.

Thanks to my family for making me laugh and smile, and for giving me something to do other than study during the making of this thesis. My nephew, Seb, has been the light of my life during the last few years. My brother, Chris, and my sister, Laura, have always been able to cheer me up with their abject silliness. Sunday dinners and cups of tea with my Nan have been invaluable. To my mom (and John) and dad (and Michelle), I thank you for your patience, belief, and encouragement. When I could not believe in myself, you believed in me for me. My mom deserves particular credit for lengthy phone calls concerning the intricacies of pragmatic theory. I love you all dearly.

This linguist struggles to 'put into words' the gratitude I feel towards Kelly Dannielle Jones. You gave me the strength, support and breakfast smoothies to knuckle down when things got tough, and finish my thesis in a less stressful fashion than I expected. Without your love and kindness, without your reassurance, and without our early morning research chats, this PhD would be a very different thesis indeed. I can never thank you enough. You are my STAR*.

Thanks go to Vanessa John, who set me on the path to the PhD when she suggested I go back to university to study linguistics. Your advice was excellent, as always. My thanks to Julia Kolkmann, the best wingman a girl could ask for. Our chats on LEXICAL PRAGMATICS* were always very INSTRUCTIVE*, and your remarkable work ethic and diligence have spurred me on in tough times. In so many ways, I would not be where I am now without you. I am also grateful to Ryoko Sasamoto and Kate Scott for their friendship, and for making me think, and dream. I want to do what you do. Adam Gargani, thanks for your advice, and for your friendship. Edit Gyenes must be thanked for her excellent food, which kept my brain ticking over on difficult days. I give love to David Gallardo, my brother from another mother, who introduced me to Inma Reivan and Flos Mariae - the soundtrack to my studies.

If there is one thing I have learned during the PhD process, it is this: whether you think you can, or you think you can't, you will be right either way.

Any and all errors in this thesis remain my own.

Abstract

It may be tempting to think that humans generally say or do things *once* in communication. However, repetition for the communication of a particular stylistic effect is a commonplace and everyday occurance. Think of texts with repeated kisses and emoji, lively conversations with friends who excitedly produce the same utterance again, and adverts and branding campaigns that feature repeated utterances and forms. Within Relevance Theory pragmatics, pragmatic stylistics generally, and in some areas of linguistics proper, stylistic repetition is understudied and under-understood. The term is applied to a rag-tag muddle of phenomena that have little in common from the point of view of form, interpretation or effects. Utterances repeated due to illness, repeated forms mandated by the grammar (reduplication), forms necessarily repeated due to limits on linguistic resources (e.g., re-use of the common conjunction 'and'), and the repetitions in poetry or rhetoric, or the repetitions produced by emphatic speakers are often lumped together, without consideration of speaker intentions, the nature of communication, or the division of labour between linguistic en-/decoding and pragmatic inference. Yet, these are all qualitatively distinct.

This thesis (re)assesses a set of phenomena which have been called repetition, for example, reduplication, epizeuxis, and 'long distance' repetition, as well as repetition phenomena which have not yet been given detailed treatments within cognitive pragmatics, pragmatic stylistics or linguistics, such as repeated gradable adjectives, repeated intensifiers, repeated yes/no particles, and repeated face emoji. Study is restricted to the deliberate and ostensive repetition of such items for communicating vague and non-propositional effects. It is noted that many repetitions, particularly epizeuxis, are often called emphatic, or intensifying, or both. A key aim of this work is to combat the conflation surrounding the effects of stylistic repetitions, and to explain how such repetitions are recognised in the first instance as intended to communicate non-propositional effects.

The work is carried out within Relevance Theory pragmatics (Sperber & Wilson, 1986/1995) and draws on the showing-saying continuum developed by Wharton (2009). From the point of view of how they achieve relevance, the repeated forms examined are all analysed as cases of indeterminate *showing* (Sperber & Wilson, 2015), and stylistic repetition is, as such, a non-verbal behaviour which allows a speaker to communicate a vague range of effects by providing relatively direct evidence for what their communication. Along the way, it is suggested that intensification is a processing phenomenon and not an effect, while emphasis is also judged not an effect, and is defined instead as highly ostensive *showing* on the part of a speaker. With support from repetition data, the author proposes a continuum of

cases from mere display, through what is called highlighting (Wharton & Wilson, 2005), to emphasis, allowing for more fine-grained analyses of pragmatic phenomena in similar contexts.

Chapter One: Repetition and Its Effects: Narrowing the Scope for Enquiry

1.1 Introducing 'repetition': an understudied muddle

One might be forgiven for thinking that we only say or do things once in communication. We don't even think about it. You utter *pass the salt* once in order to get someone to hand it to you. You also point at the salt once to encourage someone to pass it over. If you want your partner to buy some salt, you write one note to that effect. Moreover, our utterances tend to be comprised of phonologically different words internally. It is not common to find two identical words next to each other. Pick up any book and flick through. Whichever way you look at it, a lot of what we do in communication, verbal or nonverbal, written or spoken, happens *once*. This does not mean we only ever produce a word or some other communicative act once. The point is that it is atypical to intentionally produce the same word or phrase (or many other communicative elements) in close proximity to themselves. This is not really surprising; all other things being equal, we ought to expect our utterances and communicative acts to be largely different from one another (and different in terms of internal composition). For one, if every utterance we produced were the same, people would stop listening because they would be bored. For another, complex and intricate communicative acts can - from an informational perspective - communicate more than less complex and less intricate communicative acts. We should therefore expect the form of our utterances, gestures, and writings to be mostly *different*.

However, often, what we say or do in communication seems to be the same as something we have already said or done. In other words, communicative acts sometimes contain what we might term repetitions, in a pre-theoretical sense. We can repeat what we say, tweet, text, gesture, or do with our faces or tone of voice, and we can repeat lines in our poems and novels. Think of any form of intentional human communication, in fact, and you will eventually encounter repetition. The following utterances seem to exhibit repetition:

(1) You make the tuna salad and I'll make the **SALAD-salad**.

(2) He went for a **long long** walk.

(3) He's a **very very** tall guy.

(4) He went for a **long, long** walk.

(5) He's a **very, very** tall guy.

(6) Rosie: Would you like to try some of this nice tea?

Kelly: **Yes yes YES**.

(7) Julia: I'm shattered. **I'm shattered.**

(8) [A text message from a friend cancelling a meeting]

☺ you're going to kill me, really ☺ ☺

(9) Kelly: **Owowow**. that bloody hurt.

(10) [From 'Anthem', by Ayn Rand, 2010] Total pages: 102

Page 5: The sleeping halls are white and clean and bare of all things, save one hundred beds.

Page 6: **The sleeping halls are white and clean and bare of all things, save one hundred beds**.

Page 14: **The sleeping halls are white and clean and bare of all things, save one hundred beds**.

(11) [From 'Miranda', BBC1, 01/2014]

Miranda at 3 minutes 34 seconds: **Plunge** the loo before I go…expunge the bath…any activities ending in –unge should be banned…I am on an –unge protest. Although **PLUNGE** is a lovely word…**plunge**…(to camera) **PLUNGE**!

Miranda at 10 minutes 21 seconds: Ooh! Crack. Cheeky! Crack! **PLUNGE**!...**Plunge** m' crack.

In (1-11) a variety of expression types are repeated, including so-called intensifiers, adjectives, discourse particles, emoticons, and entire chunks of text. All of these, as we will see throughout the course of the thesis, are associated with the recovery of particular effects, and effects which are difficult to pin down descriptively.

In information theory, repeated (aspects of) messages are considered redundant (Danesi, 2012, p 36). Repetitions are seen to 'shore up' a message so as to ensure it gets through in the event of problems in the communication system (Danesi, 2012, p 36). If some repetitions *are* redundant, and if communicative acts without repetitions ought to be informationally

richer, we have to ask why humans repeat (elements of) what they have already said or done in communication. If there were no useful consequence to repetitions like those in (1-11), we would not produce them. Yet, we *do* produce them. As the philosopher of language Grice (1989) noted, communication is a rational endeavour involving rational communicators who are likely following some kind of guiding principle. If communication was not something like Grice's conception of it, it would fail more often than it succeeded. Without principles of some kind, communication would descend into chaos and there would be no expectation that you could hope, by saying and doing certain things, to communicate particular meanings.. If we assume that most communicators are rational, there must be some benefit or effect to producing and interpreting repetitions. We wouldn't produce them otherwise. A non-expert, if asked to reflect on the examples in (1-11) would be able to say that there was some kind of point or effect to the repetitions, even if they could not spell out what it was. Furthermore, as Andreas Jucker (1994, p 47) notes, in our time, efficiency 'is the battle cry', but repetition is pervasive in language nevertheless - this also suggests there has to be a good reason for why people do repeat things.

One reason why people *do* produce repetitions such as (1-11) is to communicate a particular type of effect. However, it is hard to pin down exactly what these effects are. What is communicated is often not a concrete proposition or a determinate range of propositions. Instead, there is a kind of vague interpretation – something that cannot be put into words, but which can be 'lost in translation' if the repetition is removed, e.g.:

(12a) I'm VEry, VEry, VEry, VEry TIREd.

(12b) I'm VEry TIREd.

It is not easy to say exactly what these repeated *very*s communicate. Such elusive effects can be called non-propositional effects in that, whether feelings or impressions, they are not made up of propositional representations, and are difficult to pin down descriptively. In this thesis, I mostly focus on repetitions that are chiefly produced in order to communicate non-propositional effects.

Repetition has been discussed within the fields of information theory, philosophy, rhetoric, literary criticism, poetics, literary linguistics, stylistics, translation studies, and discourse coherence and cohesion. Within linguistics, it comes up the areas of prosody, conversation

analysis, syntax, morphology, and pragmatics. Despite the variety of fields that attempt to analyse repetition in some way, there are few *substantial* treatments of it. Certainly, whole individual papers address the topic, as do book sections or chapters, or definitions in texts such as literary encyclopaedias and stylistic dictionaries. However, to my knowledge, there is only one book-length treatment of repetition *and language*, Swiss Papers in English Language and Linguistics 7 (SPELL) (Fischer, 1994). Though important, it contains a thematically disparate collection of papers. Even within the theoretical framework adopted in this thesis, Relevance Theory, a pragmatic framework with a history of application to stylistic phenomena, little attention is paid to repetition. In the seminal work *Relevance: Communication and Cognition* (1995), merely a few pages are devoted to its treatment. Repetition is also briefly mentioned in an interview between Deirdre Wilson, one of the founders of Relevance Theory, and Isao Higashimori (1996). The discussion runs to less than a single page. Repetition is also mentioned by Wilson and Matsui (1998) to raise problems with coherence approaches to discourse. This discussion is equally brief, and is not designed to explain how repetition works specifically, but is intended to highlight issues with coherence approaches to discourse. In Relevance Theory (Sperber & Wilson, 1995), repetition is used as an example device to explain how stylistic phenomena yield effects, and is not yet fully treated in its own right.

Repetition literature spans multiple disciplines, and is less in depth than research on other phenomena, such as metaphor, for example. In addition to this, researchers don't agree on what repetition is. The term has been applied to a myriad of (unrelated) phenomena. Jean Aitchison (1994, p 15) says that repetition 'skulks under numerous names...depending on who is repeating what where'. She (*ibid*.) says that when children repeat, it is *imitation*, when brain-damaged people repeat, it's *echolalia*, when disfluent individuals repeat, it's *stuttering*, when novelists repeat, it's *cohesion*, when morphemes repeat, it's *reduplication*, and when conversations repeat, it's *reiteration*. Aitchison (1994, p 16) lists terms for repetition and its guises, e.g., *alliteration*, *assonance*, *iteration*, *parallelism*, *epizeuxis*, *rhyme* and *shadowing*. She (*ibid*.) also notes that what could count as repetition is vast. In fact, Aitchison (*ibid*.) suggests that the entire field of linguistics could be considered as the study of repetition 'in that language depends on repeated patterns'. As we shall see in this chapter, this view of repetition is too broad.

Since the topic of repetition *is* potentially so broad, this thesis could never be all things to all people. There is not enough time and space to address everything called repetition in this thesis. However, in Relevance Theory, the question for anyone concerned with the type of phenomena seen in (1-11) above is how repetitions are recognised, processed, achieve effects, and how the answer can be presented in terms of the cognitive mechanisms and principles that underlie communicative behaviour generally. Relevance Theory has been meticulously applied to other phenomena that are also associated with the recovery of various stylistic effects. Close attention is paid to apposition (Blakemore, 2008), parentheticals (Blakemore, 2006; 2009), metaphor (Carston, 2002; Wilson & Carston, 2007; Carston & Wearing, 2014), simile (Walaszewska, 2013; Gargani, 2014), and aspects of poetic and literary form (Pilkington, 2000; MacMahon, 2007). A key aim of this thesis is to pay similar attention to repetition, and expand our understanding of it within Relevance Theory.

This is easier said than done. As I have said, not everyone agrees on what repetition is. There is a vast array of phenomena that have been called repetition, and only a particular subset of these is worthy of closer attention in relevance-theoretic pragmatics, which is concerned with a cognitively-grounded account of the relationship between form and interpretation. The scope of this thesis must be narrowed. Most of this chapter is devoted to (the reasons for) narrowing down the scope for investigation. However, first, I want to think about what it even means to repeat something. We all have a pre-theoretical notion of repetition, I think, but I wish to continue by thinking about whether or not we can actually repeat something in the first place.

1.2 What it means 'to repeat' something

Can we ever truly repeat something in communication? Literally understood, it's not really possible to reproduce anything. If I see a beautiful ring, and want a similar one, I can ask a jeweller to make me a copy. To make this copy, the jeweller examines the original ring, takes measurements, obtains the same materials, and endeavours to produce the 'same' piece of jewellery for me. However, despite the jeweller's best efforts, the ring I receive is not *the same* as the original. At best, it is extremely *similar*. Its arrangement of atoms, mixture of materials, and unique grooves set it apart from the original. The rings are *similar*, but they are not identical. (I should note that I do not say nothing can be reproduced exactly – the point is

that I don't think *humans* can ever produce exactly the same behaviours or artefacts twice. We might argue that a computer could replicate designs or patterns. However, I am not talking about computers – this thesis concerns human communication.)

Just as with rings, neither verbal nor nonverbal acts of communication can ever truly be reproduced. Utterances are time-stamped. They occur in specific contexts, for particular reasons, and are uttered by speakers, whose vocal tracts are in particular configurations, producing utterances with particular frequencies, volumes, and voice qualities. Gestures are also one-time communicative events in a similar way. It is just not possible to point at the salt in exactly the same way as you pointed at the salt before; you would never be able to reconfigure your body in an identical way. It seems as though nothing can ever be repeated exactly. Nonetheless, I think repetition *is* worth investigating. Repetitions, even if not exact, are noticeable, and communicate particular effects, which is why they are studied in the first place. There seems to be, then, a kind of paradox: we feel repetition in communication exists even though it is technically not possible to attain it. There is a way around the paradox - the solution involves drawing on the notion of *resemblance*, as understood within Relevance Theory (Sperber & Wilson, 1995, pp. 226-237; Sperber & Wilson, 2012, pp. 243-244).

The notion of resemblance is used within Relevance Theory to explain how irony and metaphor work, for example (Sperber & Wilson, 1995, p 226). However, it can also be used to explain why we needn't actually produce exact repetitions of form in order to be understood as aiming to produce a repetition for some kind of effect. Sperber and Wilson (1995, p 226) say that if you want someone to entertain the concept of a dog, but there is no dog instantiated in the environment, you can use a representation of a dog (say a drawing) to achieve a similar effect. They (1995, p 227) say that any phenomenon in the world can be used to represent in this way as long as the phenomenon resembles what you want to represent in some respects. Clearly, dogs and pictures of dogs resemble each other in terms of form. Dogs have four legs and a tail; a picture of a dog features four legs and tail. Thus, there are salient properties shared between a dog, and the pictorial representation of a dog. Since the communicator, as I explain in chapter two, §2.4.2, can be understood in these circumstances to have some useful information on offer and to be reasonably economical in his communication by intentionally using one representation to communicate another that it resembles, the audience is encouraged to look for such relevant similarities between the

picture and the actual dog. However, crucially, the picture and a real dog are not *the same*. One is two-dimensional and has no texture, while the other makes a fantastic furry companion. Nevertheless, the two are *similar enough* for one to stand in for the other if required.

Let us think again about the jeweller who made a ring similar, but not identical, to one I saw. I cannot approach trading standards and complain that the ring was not the same as the first one. That would be absurd. Everyone would say that, for all intents and purposes, the ring was *similar enough* and that it is impossible to tell the two rings apart. The two rings are recognisably similar enough, and the jeweller *intended* it that way. What we can say is that in the right context, and with the recognition of the right intentions, just as you can use a picture of a dog to represent an actual dog, an act of communication that sufficiently resembles another act of communication by virtue of sufficient shared properties *of form* would typically be recognised, in the right context, as a repetition, even if the two acts of communication – just like the rings – are not actually identical. For this reason, although I use the term *repetition* throughout, this thesis isn't *really* a study of full-on repetition, but a thesis about repetitions that are *good enough* to be identifiable as intentional repetitions that have some role to play in communication.

In what follows, in order to narrow down the scope of the thesis, there is a necessarily lengthy and detailed treatment of a set of phenomena which have, perhaps mistakenly, been badged with the term 'repetition'. This discussion is also needed to show the problems and conflations that arise in how the effects or functions of repetition are analysed. However, even though some exposition needs to be devoted to excluding certain phenomena from the thesis, there are some initial hunches, hypotheses and questions that the author pursues from the outset. I ask whether or not there are a set of repetitions that have anything in common from the point of view of form, pragmatic interpretation, or effects (and, at this stage, it is tempting to think not). I ask what it is about stylistic repetition that allows us to recognise it, and recognise it as intended to communicate such non-propositional, and often very vague, effects. I ask what kind of effects repetitions of the type treated in this thesis can communicate, and how these effects are communicated. At this early part of the study, it seems that the type of interpretations such repetitions can communicate are fairly indeterminate. I also harbour a hunch that it will not be fruitful to describe the effects of such

repetitions as emphatic or intensifying as others have done, and that we have to investigate what the repeating communicator is doing in terms of his communicative behaviour and the evidence he produces for the interpretation that he intends to communicate. Ultimately, I sense that stylistic repetition might have something to do with providing evidence for a kind of vague interpretation that we could never, ever communicate by linguistic en- and decoding. These are the lines of enquiry I pursue in this thesis but, first, it is necessary to narrow down the scope for enquiry, since not every reproduction of form has a place in this thesis.

1.3 Repetitions that fall outside the scope of this study

1.3.1 Self-repetition and echoic use

I now return to narrowing the set of cases I address in this thesis. First, I must consider who is repeating whom. (1-11) are all cases of *individuals repeating themselves*. Aitchison (1994, p 18) mentions that 'self-repetition' and the repetition of other people ('other-repetition') are variables in repetition studies, implying that there must be a reason to study these separately. In Relevance Theory, the repetition of others is already handled by the notion of echoic use. (It is important to note here that not all echoic uses are repetitions, and not all repetitions are echoic, however.) Sperber and Wilson (1995; 2012) say that an utterance can be used either descriptively or interpretively. An utterance is used descriptively if it describes a state of affairs in the world, such as saying *It is raining* when it is raining. An utterance is used interpretively when it is an interpretation of an existing representation, as in Kelly's utterance in (13) below:

(13) Rosie: What did David say?

Kelly: I'll be late.

Here, Kelly's utterance interpretively represents what David actually uttered. Kelly is not understood as saying *she* will be late. David may not have uttered the exact same words as Kelly did, but he may have uttered words to that effect, e.g., *I won't make it on time*. Kelly's utterance interpretively resembles David's because they share logical properties in terms of the propositions they express. Echoic use, which handles the repetitions of others, is a special

type of this interpretive resemblance. An example of an ironic utterance shows how this works. Consider (14a-14b) below:

(14a) [Uttered at 9am]

Kelly: What's the plan for studying?

Rosie: **I'm going to work extra hard today.**

(Rosie spends the rest of the day watching television.)

[Uttered 5pm, same day]

(14b) Rosie: I've done nothing at all.

Kelly: **I'm going to work extra hard today!** (Uttered in a mocking tone of voice.)

Kelly's highlighted utterance interpretively resembles Rosie's original utterance from the morning in terms of content and form. Kelly's utterance is understood echoically because it does not achieve its effects by representing Kelly's own views, or by reporting what someone else has said, but by expressing Kelly's attitude to thoughts that she tacitly attributes to Rosie. Kelly is being ironic. Her tone is mocking; she expresses a dissociative attitude to what she attributes to Rosie (other attitudes can be communicated depending on tone of voice and other paralinguistic features) (see Sperber & Wilson, 2012, pp. 123-146).

The effects of ironic echo are achieved as follows. There is strong communication (see chapter 2, §2.7.2) in that the thought or utterance attributed to Rosie is easily recoverable, and that the attitude communicated by Kelly is almost impossible to miss. However, there is also weak communication in that the utterance triggers the recovery of a range of assumptions that Rosie has much of the responsibility for recovering, e.g., Kelly thinks I am lazy, I never hit my goals, I'm going to fail my course, and so on. If the array of assumptions communicated is quite weak and vague, non-propositional effects may be recovered. At this point, we can say that the repetition of others in order to communicate a particular effect can be dealt with by the existing notion of echoic use, and, thus, Aitchison's (1994) 'other repetition' can be excluded from the thesis. The only problem is that the notion of echoic use also allows for people to echo themselves.

(15a) Julia (to a friend): I've got to stop eating cake. **I'm going on a diet**.

[Several hours later, Julia is eating a large piece of cake with her friend.]

(15b) Julia: **I'm going on a diet!**

Julia's tone of voice in (15b) communicates a dissociative attitude to the proposition that she attributes to herself from (15a) earlier in the day. (15a) is also ironic.

Based on form alone, it is impossible to distinguish between self-repetition, and self-echo. However, the ways that the repetitions in (1-11) achieve their effects and the way that self-echoic utterances achieve their effects should be different, in many cases. First of all, the relationship between an echoed utterance and its echo is generally one of resemblance of *content*, and the effects arise, in part, from the degree of resemblance in logical implications. Repetitions are recognised by the degree of shared *form* between what is repeated (the original) and the actual repetition, and their effects cannot arise from calculating shared logical implications between an original and its repetition(s) - they would share a full set of logical implications every time. Finally, while self-echoic utterances and 'self-repetitions' can both communicate non-propositional effects, the effects of many repetitions *won't* be derived from weakly communicated implicatures (cf. Sperber & Wilson, 1995), which is how the non-propositional effects of irony above were derived. The notion of echoic use handles cases of other-repetition, and any cases of self-repetition where you communicate an attitude to a thought or utterance you attribute(d) to yourself. Cases like (1-11) are *not* covered by echoic use. The focus of this thesis is self-repetition. Self-repetition involves a deliberate, *ostensive* attempt to reproduce the form of a communicative act that you have produced previously. This thesis is not a study of all self-repetitions, however. I now exclude a further number of repetition scenarios from study.

1.3.2 Unintentional repetitions

Aitchison (1994, pp. 22-23) says that there are two main types of unintentional repetition: perseveration, and echolalia (more severe perseveration). Perseveration is common in aphasias and some dementias (Snowden *et al*., 2002). It occurs when an individual repeats an utterance or gesture in the absence of a stimulus, or fails to recognise that one has already produced an utterance or gesture. In echolalia, individuals just repeat back to the co-communicator something that they have just said due to an inability to respond competently in communication, as in severe cases of autism (*What is autism?*, 2015). In addition to these unintentional repetitions, there exist cases of drunk people who tell you something again and

again, forgetting they have already told you. There are also psychiatric conditions that can present through jumbled and chaotic speech. Some individuals with schizophrenia exhibit the phenomenon called schizophasia, or 'word salad', where speech is largely nonsensical (Verhaeghe, 2008). Such speech can contain partial or complete repetitions. As such, there are many circumstances where people do not, and often *cannot*, intend to repeat themselves. Moreover, forgetful, nervous or distracted individuals repeat themselves all the time due to attentional issues. These cases are also excluded from this thesis because the speaker does not intend to repeat something for the communication of a particular effect, or, in relevance-theoretic terms, these are not cases of *ostensive-inferential communication* (see chapter two, §2.3).

The scope for investigation is therefore narrowed down to self-repetitions which are intended. But, this isn't to say all intentionally produced repetitions belong in this thesis either. There are repetitions which are deliberate in the sense that the speaker or writer intended to form their utterance that way, but which also have to be excluded from this thesis because whatever was repeated was not repeated to communicate a particular stylistic and non-propositional effect.

1.3.3 Non-stylistic intended repetitions

Sometimes people repeat themselves just because an original act of communication has failed, i.e., the speaker was not heard the first time around, or waves at someone again because his first gesture was not seen. Sometimes a speaker's co-communicator is aware that they missed the initial attempt at communication, and sometimes they are not. Depending on the circumstances, you might speak louder or gesture more wildly the second time around, but the wording or the type of gesture you employ is likely similar in form to the one misperceived the first time. In these cases, an individual repeats what they have said or done only because they realise there was a problem with the perception and/or processing of the original communicative stimulus. The intention behind such repetitions is just to communicate the same interpretation as was intended the first time around. In chapter two, §2.4.2, I explain that just because a speaker aims at being optimally relevant, it does not follow that he always succeeds – and this is what happens in these types of cases. Thus, repetitions which are intended to remedy failed or imperfect communication already have an

explanation in Relevance Theory – the speaker simply reproduces his original attempt at optimal relevance. In addition to these cases, there are situations where a speaker does repeat himself on purpose and absolutely intends the form of his utterance(s), but the repetitions of form are incidental, or, in a sense, necessary. As such, these repetitions are also not intended to communicate any particular effects. In her 1987 paper *Repetition in Communication*, Deborah Tannen offers examples of repetition in order to show how pervasive it can be in everyday communication. Her examples are taken from spontaneous conversations recorded by her students. There is something interesting about some of the examples she puts forward. Elements of linguistic form are certainly repeated but the case is easily made that these repetitions are not actually intended to communicate any particular effect. The repetitions are, in a sense, incidental, and a by-product of how speakers can be constrained by the linguistic resources available to them for linguistically encoding the proposition schemas they want to communicate. Consider the following two examples taken from Tannen (1987, pp. 587-588):

(16) Vivian: So I stood on my bed

Marge: She **pounded** on the **ceiling**

Vivian: And I **pounded** on the **ceiling**

Marge: She was **pounding**...

Vivian: I was **pounding** on my **ceiling**

(17) Frank: Well **daydreaming** is something that comes natural! [You don't don't PLAN daydreaming.] ...

Frank: You can't PLAN **daydreaming**…I'm gonna go **daydream** for a couple of hours guys so...

Tannen (1987) does not explicitly state why she chose these excerpts as examples of repetition. It seems reasonable that the repeated lexical items *pounding* and *daydreaming* (and their variants) played a role in her decision-making. However, these low frequency words have few readily accessible synonyms. *Hitting* or *knocking* don't entail the intensity of *pounding*, and I can think of no higher frequency noun that would ordinarily communicate the same concept as the word *daydreaming*. My point is that, compared to a case like the repetition of *plunge* in (11), these words are not repeated to trigger any stylistic effect. In (16)

and (17), there are quite particular concepts that the speakers wish to (re)express, and the expression of these requires the necessary re-use of certain lexical resources. A speaker could say something other than *daydreaming*, for example, but a suitable paraphrase is more effortful to come up with *and* interpret. Speakers in such contexts almost 'have to' repeat such words. Of course, the example in (16) could be analysed as 'stylistic' in a sense, but it is important that this is not picked up on as something to analyse, given the aims of the study at hand.

Necessary re-use is not just restricted to lexis. Recall the point made by Aitchison (1994) that there might be a sense in which all of linguistics can be seen as the study of repetition because there are repeated structures and patterns in many areas of linguistics. For example, a common basic syntactic structure might be something like [DP_{subj} [VP [DP_{obj}]]] as in sentences like *I eat cake* or *Kelly baked bread*. There is a sense in which I repeat this sentence structure many times every day, but I do not do so for any particular effect. When I repeat this syntactic structure, I do so because rules of English require me to when I want to express certain simple and commonplace types of proposition. I could always utter something else, but this would be marked, and I should only utter something else if there is a communicative reason for doing so. Put simply, repetitions of all kinds of linguistic forms occur daily, but the majority of them will not be of interest to this thesis.

Before moving on, it is worth noting, particularly in the case of syntax, that of course we usually *can* repeat any aspect of linguistic form for an effect. Imagine that I pen a few lines of (not particularly sophisticated) lyrics for a children's song:

(18) She was singing,
She was dancing,
She was laughing,
She was smiling,
She was howling,
She was clapping (and so on…)

We could argue there is some kind of effect associated with the morphology and syntax of these lines. Though it is hard to pinpoint it, we know that an effect of some kind is there, since it disappears or changes if you change the morphology and syntax of the lines.

(19) She sang.

She was dancing.

She did laugh.

Laura smiled.

She was howling.

My sister did clap.

A key aim of this thesis is to shed light on how and why an audience is able to work out whether a repetition of linguistic (or non-linguistic) form is intended to be used in the recovery of a particular effect. The point to be made here is that not all repetitions of form are used in the recovery of such effects, but that those forms *could be* in the right context. Essentially, there is nothing inherent in the forms repeated in these examples in terms of whether or not they are permitted to enter into repetitions of one type or another - acceptability depends on the context. Nevertheless, repetitions of form which are present because we must necessarily re-use some elements of our language are not included in the scope of this study.

1.4 The form of ostensive repetitions triggering the recovery of particular effects

This thesis analyses deliberate self-repetitions produced in order to communicate a particular effect. However, I have not yet said much about the forms that can be repeated to achieve these effects, or provided an account of what kind of effects the repetitions of these forms achieve. The second question (what kind of effects) forms the focus of chapters three to six. In the meantime, I briefly say something about the sort of forms which can be repeated to communicate the effects I am concerned with.

In this study, I draw on both linguistic and non-linguistic examples - just as we saw in the preliminary data in (1-11). That is to say, there are examples of written and spoken utterances, but there is also discussion of repetitions of non-linguistic or borderline linguistic material, e.g., expletives, interjections, elements of prosody. I also want to contribute to the

burgeoning debate surrounding what and how emoji communicate. This topic is addressed through the lens of repetition in chapter five. In any case, I draw on linguistic and non-linguistic repetitions - not because I necessarily expect linguistic and non-linguistic repetitions to communicate in the same way, but because any behaviour that can be deliberately produced in communication can also be deliberately *repeated* for effect. Communication is made up of verbal and nonverbal behaviours, and both are repeatable. I do not privilege the former over the latter.

In this thesis, when I use the word *form*, I mean any physical properties, structure, make-up or characteristics of the percept that is produced as a communicative act. Thus, form can be multi-layered, as we will see from the data in chapters three to five. With verbal utterances, aspects of form might be choice of lexis, volume, certain syntactic structures, voice quality, and so on. In written utterances, form could be colour, font size, the use of emoji, underlining or bold, and so on - even the placing of white space. In terms of paralinguistic behaviours such as facial expressions and aspects of tone-of-voice, I would look at configurations of the face, or aspects of vocalisations such as pitch, pause, volume, range, and tempo, for example. These types of forms are the sorts of things can be repeated for the communication of non-propositional effects and a subset are treated in chapters three to five.

Finally, one last point must be raised with regards to the scope of this thesis. There is a particular type of reproduction of linguistic form which has been conflated with repetition, reduplication (see chapter three, §3.2). Reduplication is generally thought of as the repetition of a word or part-word to express a particular grammatical function (Aronoff & Fudemann, 2005). From the point of view of *form*, it might seem like I need to include phenomena termed reduplication within the scope of this thesis. Consider (20a) and (20b) below, which are both considered to be reduplicative:

(20a) You make the tuna salad and I'll make the SALAD-salad.

(20b) [From Indonesian (Aronoff & Fudemann, 2005, pp. 167-168)]

Kuda = horse

Kuda kuda = horses

However, if the broad definition of reduplication set out above is correct, it is not clear that it has a place in this thesis, its function falling chiefly within the domain of grammar, even

though the phenomenon may well have its evolutionary roots in something pragmatic. This issue is the subject of the beginning of chapter three, where I argue that reduplication proper is not to be conflated with stylistic repetition.

1.5 The functions of repetition

I now consider two theoretical constructs claimed to play an explanatory role in the way discourses are constructed, and which address or impact on the study of repetition data: cohesion and coherence. In Relevance Theory, discourse and discourse structure are not the object of study – the focus is on the cognitive processes involved in ostensive-inferential communication. In contrast to Relevance Theory, however, other approaches do take discourse as the object of study, and aim to show how properties of discourse explain how it is understood. Two key properties for those working on discourse are cohesion and coherence, and these notions are related to the connectedness of discourse – our sense that discourse 'hangs together' somehow. An important assumption for those working on coherence or cohesion is that the study of discourse understanding begins with the claim that we can understand discourses *because of* their connectedness

Some researchers on discourse connectedness deal quite explicitly with repetition (e.g., Halliday & Hasan, 1976), while others seek to explain the connectedness of discourse more generally but their work has implications for the wider study of repetition (Mann & Thompson, 1986; Giora, 1985, 1997, 1998).

1.5.1 Cohesion and repetition

As Halliday and Hasan (1976), Hobbs (1979), and Sanders and Maat (2006) note, discourse is not random; discourse exhibits connectedness. Halliday and Hasan (1976, p 1) observe that all speakers of English can recognise whether a set of sentences are connected as a text, or seem completely unrelated. Those researchers of discourse connectedness interested in the cohesiveness of discourse examine the overt linguistic elements present therein which contribute a text's cohesiveness by having the function of helping it to 'hang together' (Halliday & Hasan, 1976). On this view, a text hangs together because the interpretation of certain linguistic elements depends on the interpretation of other elements; pronoun

interpretation in anaphoric reference is the paradigm case (Halliday & Hasan, 1976). Meanwhile, researchers studying the coherence of discourse take a more psychological approach to connectedness, examining it in terms of cognitive representations of texts rather than in terms of linguistic elements of texts themselves (Sanders & Maat, 2006). Coherence theorists have different means of approaching where coherence emerges from. A common approach has coherence as emerging from relations holding between adjacent segments in a text (Mann & Thompson, 1986). Such theorists are generally interested in 'local' coherence relations. Another approach is to consider that coherence does not emerge exclusively from relations between local segments, but, instead, comes from how well a segment behaves with respect to some general principles governing the well-formedness of discourse at a 'global' level. Giora (1985, 1997, 1998) takes such an approach. First, I consider what Halliday and Hasan's (1976) *Cohesion in English* tells us about repetition.

As I said above, humans can judge whether or not a collection of sentences seems to hang together as a text. By *text*, Halliday and Hasan (1976, p 1) mean any collection of sentences, however big or small, spoken or written, which is judged to be a 'unified whole'. Halliday and Hasan (1976, p 1) argue that if we are all able to judge which collections of sentences 'hang together' as a text, there must be some objective features of texts which we use to make judgements about cohesion. The point of *Cohesion in English* is to delimit the features of texts which contribute to cohesion, and explain how they do so. It is important to note here that such features are *linguistic* in nature. Cohesion comes from the presence of particular linguistic elements, e.g., antecedents and anaphors, synonyms, and conjunctions. Thus, cohesion is found in linguistically encoded relationships in texts (cf. Blakemore, 2002, p 159).

Given that the presence of certain linguistic elements is considered to almost 'guarantee' the connectedness of a text, it is unsurprising that Halliday and Hasan (1976) approach the issue of cohesion from a semantic angle. After all, the only means (strictly) of guaranteeing that a particular meaning (or sense of cohesiveness) is recovered is to use a (well-specified) code, and semantics is the study of linguistic en-/de- coding. Texts are considered semantic units on this approach (1976, p 1), and '[t]he concept of cohesion is a semantic one; it refers to relations of meaning that exist within the text, and that define it as a text' (Halliday & Hasan, 1976, p 4). Cohesion emerges from the interpretation of textual elements being dependent

on the interpretation of other elements. There are five types of cohesion distinguished by Halliday and Hasan (1976). The first four are considered types of 'grammatical' cohesion and are set out in (21a-21d):

(21a) **Reference**: Present when two linguistic elements refer to the same entity, e.g., Kelly drives [Rosie's car$_i$]. [It$_i$] is a bit slow, though.

(21b) **Substitution**: Present when an item is replaced by a substitute instead of being repeated, e.g., Seb loves [Eggman ice creams$_j$]. He eats [them$_j$] every time we go to The Blue Ginger.

(21c) **Ellipsis**: Present when an item is left out rather than repeated, e.g., Both cats had meat. Pawl had a sachet of Whiskas, and Tilly did too.

(21d) **Conjunction**: Present when semantic relations are overtly 'marked' in a text, e.g., Laura is tired since she has been job-hunting all day (cause); David was upset but went to work anyway (contrast or contradiction).

This is a very broad sketch of Halliday and Hasan's (1976) framework. The sub-type of cohesion that is most pertinent to this study is not grammatical, though - it is *lexical*, and here repetition enters the picture. Lexical cohesion emerges from the selection of vocabulary items in a text, and lies in sets of items which stand in certain lexico-semantic relations to one another (Sanders & Maat, 2006). There are two key types of lexical cohesion: collocation, and reiteration, and (exact) repetition is considered to be a type of reiteration. Collocation, as understood by Halliday and Hasan (1976, pp. 284-286), pertains to lexical items that often occur in the same lexical environment, or which are semantically related, e.g., *big* and *small*. I now look at reiteration as a superordinate form of lexical cohesion, examining specifically what Halliday and Hasan call repetition.

Halliday and Hasan (1976, p 278) say that '[r]eiteration is a form of lexical cohesion which involves the repetition of a lexical item, at one end of the scale; the use of a general word to refer back to a lexical item, at the other end of the scale; and a number of things in between - the use of a synonym, near-synonym, or superordinate'. Halliday and Hasan (1976, pp. 278-279) propose a 'cline' of cohesiveness for these cases of reiteration, saying that such cases are related to one another because original lexical items and their reiterata always share a

common referent. Reiteration so defined subsumes lexical items that are related to one another both by virtue of their form *and* their meaning on one hand (exact lexical repetitions), and lexical items that are *not at all* related by form, and which are only related by similar or overlapping meanings on the other. (22a) below is an example of lexical repetition proper, (22b) is an example of a 'middle case' involving a synonym, while (22c) exemplifies reiteration 'at the other end of the scale', i.e., reiteration involving the use of a 'general word' to refer back to a previous lexical item:

(22a) I love [that cake$_i$]. [That cake$_i$] is delicious.
(22b) The suspect admitted taking [the vehicle$_j$] – [the car$_j$] was taken from outside the shop.
(22c) There's [a squirrel$_k$] by the barbecue. [The numpty$_k$]'s going to get burnt if he doesn't move!

To reinforce the point, cohesion on this view emerges because of shared reference of co-indexed elements such as those in (22a-22c).

Although Halliday and Hasan discuss repetition, their intention was not to develop a theory of it. They simply wanted to examine factors underpinning discourse connectedness. Due to the necessarily limited scope of their investigation, the type of repetition data it can handle is necessarily restricted. There are many forms of repetition data that they cannot account for but which arguably contribute to some kind of connectedness in discourse because they are acceptable to us.

Regarding the treatment of repetition here, the first point to address is that the type of repetitions that Halliday and Hasan mention are lexical, and only lexical (though sometimes articles or demonstratives are also repeated due to grammatical necessity). This has two consequences: firstly, Halliday and Hasan have nothing to say about repetitions of linguistic material which are longer than a single content word, or longer than, say, a noun phrase. They cannot account for repetition at or above the clause level. Yet, such repetitions are abundant (see chapter four). Secondly, they have nothing to say about repetitions of non-linguistic or borderline linguistic material, e.g., emoji and interjections. Of course, *Cohesion in English* was written before emoji were invented, and it is important to be fair and acknowledge that the

book did not aim to explain such phenomena. Nevertheless, if we look back to the examples raised at the start of this chapter, we will see that while there are deliberate and communicative repetitions that look like Halliday and Hasan's, there are equally deliberate and communicative repetitions which don't. It is in this sense that Halliday and Hasan admit too *narrow* a range of repetitions into their investigation. We should like an account of as many cases of repetition as possible (Wilson & Matsui, 1998; Blakemore, 2002).

There is also a sense, however, in which Halliday and Hasan's understanding of repetition might be too broad. Earlier, I introduced the idea of 'necessary re-use', where a linguistic form is repeated only because a speaker needs to communicate a particular interpretation, and his language has limited resources for doing this without periphrasis. Halliday and Hasan do not make the important distinction between such cases, and cases where a lexical item is repeated for a stylistic effect. Although some of the repetitions I am interested in resemble clear cases of lexical cohesion from the perspective of form and semantics (e.g., *There's a fire, a fire!*), Halliday and Hasan's account has nothing to say about the non-propositional effects that such repetitions can engender, or a relationship between these effects and the connectedness and acceptability of discourse. Yet, such repetitions are perfectly acceptable. Actually, Halliday and Hasan (1976, p 10) actively back away from the investigation of stylistic effects, saying they will avoid the study of devices such as metre, rhyme, and parallelism. By not making a distinction between 'necessary' and stylistic repetitions, Halliday and Hasan cannot actually exclude the repetitions that they say they want to rule out.

Since Halliday and Hasan don't distinguish between stylistic and non-stylistic repetitions, and between intended and 'necessary' repetitions, they don't need to propose any mechanisms by which we would recognise or dismiss certain repetitions. In particular, it is not clear how their framework could provide an account of the stylistic effects of repetition (since it is a semantic, that is, linguistic, framework) because it excludes the role of the context and pragmatic inference in the interpretation process. As I shall show, repetitions may give rise to different effects in different contexts. But, more fundamentally still, the decision by a hearer to expend effort on the recovery of effects from a given repetition follows from his recognising that the production of the repeated form was intended to communicate those very effects, and, as we have seen, Halliday and Hasan's framework does not give us a means

of distinguishing between those repetitions which are a product of the linguistic system and those which are a result of a communicative intention. Any account of the interpretation of repetitions in (1-11) must be able to explain why we identify them and process them in particular ways. It should be noted that there is a sense in which 'necessary' can be applied to stylistic repetitions in that they could be 'necessary' to the communicator's intentions, but the point is that 'necessary' here refers to form being mandated by linguistic constraints of one kind or another.

Before proceeding, it is worth noting that there seems to be a very particular type of repetition present in Halliday and Hasan's (1976) discussion of reiteration. I call this repetition-at-the-level-of-thought. Essentially, as in example (22c) above, it seems that [a squirrel] and [the numpty] trigger or (re)activate the same concept in our mental representations in this context. (22c) clearly contains two different noun phrases which both have the same referent, and it seems at least possible that the hearer's mental representation of (22c) could contain two similar instantiations of the same referent, i.e., the entity that [a squirrel] and [the numpty] pick out in the world. There could be, in a sense, a repetition of conceptual material here. However, I could not scientifically test that this is the case. We can see activation levels on brain scans, but this does not tell us what type of thoughts a person is having. These mental representations are not available to scientific, replicable scrutiny, so I will not study them in this thesis. However, we *can* study repetition of form. For example, we can use software to see just how similar two intonation contours actually are, or use measures of surprise to see whether someone has noticed a repeated form. This thesis is chiefly a theoretical piece. However, the point is that the phenomena addressed in (1-11) could be studied empirically if required. As such, repetition-at-the-level-of-thought is not further addressed in this study.

1.5.2 Local coherence and repetition

Mann and Thompson's (1986) Rhetorical Structure Theory (henceforth RST) is a 'local' coherence theory because it is concerned with identifying meaning relations between adjacent segments of text. I discuss it because the recognition of some of the relations proposed to exist in RST depends on the repetition of linguistic form. Just like Halliday and

Hasan, but in contrast with relevance theorists, Mann and Thompson are interested in accounting for how discourses are interpreted. In contrast with Relevance Theory, Mann and Thompson (1986) see coherence as playing a primary role in discourse understanding in the sense that any pragmatic inference follows from the assumption that a given utterance always stands in a particular coherence relation with the preceding segment. In this section, I examine Blass' (1990), Wilson's (1998), and Blakemore's (2002; 2004) claims that this assumption cannot be maintained, and explain how this debate contributes to the study of repetition data.

RST (Mann & Thompson, 1986) accounts for coherence in texts at a level greater than that of the clause. It was developed for the study of computer-based text generation. Mann and Thompson (1986) account for coherence by positing relations that hold between particular propositions, called coherence relations, which hold because of formal properties of texts such as the presence of particular so-called discourse markers, e.g., 'in other words' or 'namely'. These relations can be hierarchical in nature, and can stretch across varying expanses of text (Mann & Thompson, 1986). Each of the proposed relations is defined in terms of a particular function. It is a main assumption of this framework that a coherent text has a clearly identifiable function for each of its parts; no textual element is out of place (*Intro to RST*, 2015). Coherence is understood as 'the absence of non-sequiturs and gaps' (*Intro to RST*, 2015).

Since the interpretation of a discourse depends on the successful identification of particular coherence relations between adjacent segments of text, it follows that there necessarily must exist a taxonomy of rhetorical relations upon which an audience can draw during interpretation. However, the relations proposed by Mann and Thompson (1986) seem to be a rather arbitrary set, and no one agrees how many such relations there ought to be. Moreover, it is not clear how such taxonomies are acquired during language acquisition, and there is no definable point when communication would break down if your co-communicator drew on a coherence relation you had not acquired. There is nothing wrong with including a taxonomy in a model if there is evidence that the taxonomy is natural – just as we have taxonomies of distinct entities in the natural sciences. It is not clear that Mann and Thompson's taxonomy of relations form a taxonomy in this acceptable sense, however.

Working from Relevance Theory's deflationary standpoint, one might also ask if there is a more elegant explanation to be had.

Within RST, relations are posited between at least two spans of text: a nucleus and (minimally) a satellite (*Intro to RST*, 2015). It is important to note that nuclei and satellites are almost always adjacent. The relationship between a nucleus and a satellite is such that the satellite plays a specific role with respect to the interpretation of the nucleus (*Intro to RST*, 2015). Let us consider how this works for two common Rhetorical Structure Relations (hereafter RSRs): EVIDENCE and PURPOSE.

When two text spans exhibit the RSR EVIDENCE, the nucleus span yields a claim, and the satellite span provides evidence which should increase belief in the claim derived from the nucleus, as in (23a) below.

> (23a) Nucleus: Chris is a hard worker. Satellite: He regularly does fifty hours a week.
> (23b) Nucleus: Laura wants to make a cake. Satellite: She fancies eating something sweet.

In (23b), when two text spans stand in an RSR of PURPOSE, the nucleus presents an intended situation, while the satellite yields information about the intent underpinning the intended state of affairs.

On this view, a reader simply selects the right RSR from a range available to her, and the text will 'make sense', and cohere. This approach is often criticised because it is not clear how the 'correct' relation is selected (Blass, 1990; Blakemore, 2002, 2004). For example, the two-segment utterance in (24) could be interpreted as either ELABORATION (involving additional information), or RESTATEMENT (involving a re-expression):

> (24) Today, we are learning about syntax. We will learn about the formal structural properties of language.

In spite of this, let us consider now how repetitions fit into this picture. Of the twenty-five or so RSRs posited in RST (*Intro to RST*, 2015), there is no specific RSR REPETITION. What is of

interest is how repetitions of segments might lead to the identification of one particular RSR, and what we can learn from this. Here, the RSR ELABORATION is important. An RSR ELABORATION can be identified when a nucleus yields *basic information*, while its satellite yields *additional information*. A typical example of this is (25) below:

> (25) Trove Cafe is near a supermarket. Go down the A6 until you get to Tescos on the right hand side, and Trove is virtually next-door.

The second segment of (25) clearly provides more in-depth information than that expressed by the first segment.

Blakemore (2004, p 102) notes that, for some researchers (e.g., Hobbs, 1979), repetitions of the form x, x (commonly called epizeuxis) are also considered to satisfy the formal criteria for an elaboration relation, as in (26a) and (26b) below:

> (26a) (taken from Blakemore, 2004, p 102) There's a mouse, a mouse.
> (26b) You don't understand. You don't understand.

Blass (1990, p 18) explains that, for Hobbs, repetitions like (26a) and (26b) are seen as elaborations because, although both segments 'superficially' express the same proposition, new information is indeed conveyed through the second segment by virtue of the fact that there has to be a reason that the communicator uttered the first segment again. If this is right, and repetitions of the form x, x are indeed elaborations, then the presence of any x, x repetition would satisfy the formal criteria for ELABORATION, and the text should be both locally coherent and acceptable. However, this is not always the case. Consider (27a) and (27b) below:

> (27a) [I go into the library and say to the assistant] I'm looking for the linguistics shelf. I'm looking for the linguistics shelf.
> (27b) [I help my friend spell a word] 'Psychologist' has a silent 'p'. 'Psychologist' has a silent 'p'. 'Psychologist' has a silent 'p'.

(27a) and (27b) could only be acceptable in some very particular contexts indeed. Thus, it seems possible to construct texts, as Blass (1990) notes, which satisfy the formal requirements for coherence, but which are unacceptable without reference to some kind of context. This suggests that repetitions can only be judged acceptable or not when they are processed with respect to a particular context. What is clear from the discussion above is that satisfying the formal criteria for coherence does not guarantee the acceptability or well-formedness of discourse. Whether or not an utterance is understood as standing in a particular RSR to another will depend on whether the identification of that relation is relevant in the context.

Not only is the satisfaction of the criteria for coherence not sufficient to guarantee coherence, but the satisfaction of the criteria is also not *necessary* for coherence to obtain. Consider (28):

(28) David: What did Laura say?
Kelly: Our bus is late.

The coherent interpretation here is that Kelly's answer is a representation of what Laura said to her, and that it is an answer to David's question. This may well be the interpretation that is recovered. However, it is not difficult at all to imagine a highly accessible context in which this is *not* the interpretation that is recovered. David and Kelly might be talking at the bus stop while Kelly is waiting to go to work, and she is late – she might feel the need to mention that the bus is late if this is more important to her than talking about Laura, for example. Thus, we see that a locally 'incoherent' sequence can actually be perfectly well-formed and acceptable. The suggestion is, then, that local coherence is not actually a prerequisite for utterance understanding. It is the hearer's quest to make sense of things in context that drives the interpretation process. It just happens that, in some cases, it is the locally coherent interpretation that is most relevant, while, in other cases, the 'incoherent' interpretation is most relevant. Local coherence is a by-product of the interpretation process rather than a prerequisite for it (Sperber & Wilson, 1995, p 289). Indeed, we can actually define it in terms of relevance (Blass, 1990) – if an utterance is relevant in a context which includes the interpretation of the preceding utterance, then local coherence obtains. However, if a given utterance is relevant in some other, non-local context, then there is no local coherence. A proper definition of relevance is given in chapter two, §2.4.

Two final criticisms of RST are as follows. Firstly, it can only deal with adjacent repetitions of clauses. As pointed out by Wilson and Matsui (1998), Mann and Thompson are unable to say anything about sub-clausal repetition (e.g., *It was very very hot*), and they can provide no insights about repetition-at-distance (e.g., example (11)). However, these are phenomena that I must address, as both can trigger the recovery of stylistic effects. Secondly, as with Halliday and Hasan's (1976) cohesion account, RST also cannot be applied to communicative but borderline or non-linguistic forms (repeated or not) such as emoji, or certain elements of prosody. If researchers working on local coherence are not able to explain how all the repetitions in (1-11) are considered well-formed, and how they are interpreted, then perhaps it is worth seeking insight from work on coherence which does not rely on local relations between discourse segments – Giora's (1985, 1997, 1998) global account of coherence.

1.5.3 Global coherence and repetition

As we saw above, there are cases where the context for interpretation must involve something other than the information expressed by the preceding segment. Giora (1985, pp. 699-700) states that coherence does not obtain from linear connections between sets of segments, and that sentence-level analysis is not relevant at all to the study of coherence. In other words, coherence must obtain globally, taking into account the whole of a text. She (*ibid*.) goes on to say that the notion of coherence is best understood in terms of orientation to a discourse topic, and that text coherence is, intuitively, best understood in terms of the idea of text well-formedness. Here, I set out what it means for Giora for a text to be well-formed, and set out any requirements that lead to text well-formedness. Then, I show why repetition data are problematic for these requirements.

Giora (1985) explains the coherence of texts by appealing to the notion of orientation of a segment to a discourse topic. In (29) below, Giora (1985, 1997, 1998) would say that the second segment contributes to coherence because it predicates something about the topic of the discourse, i.e., what a segment of text is *about* (Hobbs, 1979, p 24). In the case of (29), we can argue the topic is *David*.

(29) David is from Spain. He speaks Spanish, English, and a little German and Japanese.

The second segment of (29) predicates something about David, who can be argued to be the topic of (29) as a whole. (29) is considered a coherent text because the second segment predicates something about a discourse topic, David. Utterances that do not *comment* on a discourse topic or which do not *predicate* something about a discourse topic are not judged to be coherent (Giora, 1985, p 705). In later work, Giora (1997, 1998) develops this claim into two specific requirements which, when met, should yield a coherent text. These are *the Relevance Requirement*, and the *Graded Informativeness Principle*.

The Relevance Requirement mandates that all propositions expressed by a text be related to a (generally explicitly designated) discourse-topic proposition (Giora, 1998). Giora (1998, p 80) states that discourse topics are generally placed early on in a text, and are *mentioned*. Digressions are permitted if they are explicitly pointed out with a 'digression marker' such as *by the way*. Discourse topics function as a 'reference point relative to which all incoming propositions are assessed and stored' (Giora, 1997, p 220), and the nature of information organisation is akin to how Roschian prototypes are set up (Giora, 1997, p 23; also, see chapter three, §3.4.3 of this thesis). On this view, it would seem that (29) is coherent, because it conforms to the Relevance Requirement in that it has an explicit topic, *David*, which is mentioned early on, and a proposition is predicated of this topic. The Graded Informativeness Principle 'requires that each proposition be more (or at least not less) informative than the one that precedes it in relation to the discourse topic' (Giora, 1997, p 22). Propositions are considered informative only to the degree that they have properties 'unshared by the previous proposition' (*ibid*.). What is important here is that propositional representations share greater or fewer properties with each other. On this view, (29) is well-formed and coherent because the proposition expressed by the second segment is more informative than that expressed by the first segment, having properties that are not exhibited by the first proposition. Let us be very clear what is meant by *informativeness* here. Giora (1991) views informativeness as it is viewed in information theory – a message is more or less informative relative to the number of possibilities it reduces or rules out for a given question. The Graded Informativeness Requirement seems to suggest that each *proposition* expressed by a segment

should be more informative than preceding segments. What is not clear here is whether 'message' is taken to mean purely the proposition expressed, or other aspects of meaning as well. However, as far as I can judge from the discussion, Giora seems to mostly refer to semantic / linguistic meaning.

What predictions can we make about repetitions of the type I am interested in on this account? In a text that we find coherent, well-formed, acceptable and, loosely put, not out of the ordinary, we should only find text segments which express *propositions* that are *about* a topic or comment on a topic, and propositions which are more informative than those that have gone before in that they share relatively fewer logical properties with propositions which have been expressed earlier. Should we encounter ANY acceptable repetitions of form above the clause level in a text that is judged to be generally coherent and acceptable, then they must necessarily conform to these requirements, unless they are explicitly marked by some kind of deviation marker. However, what we actually find is that there are perfectly acceptable repetitions which do not conform to these requirements and are not introduced as deviant by special markers. Consider (30) and (31):

(30) [An extract from 'The Raven' by Edgar Allan Poe, 2002]
Then, methought, the air grew denser, perfumed from an unseen censer
Swung by Seraphim whose foot-falls tinkled on the tufted floor.
"Wretch," I cried, "thy God hath lent thee—by these angels he hath sent thee
Respite—respite and nepenthe from thy memories of Lenore;
Quaff, oh quaff this kind nepenthe and forget this lost Lenore!"
Quoth the Raven "Nevermore."

"Prophet!" said I, "thing of evil!—prophet still, if bird or devil!—
Whether Tempter sent, or whether tempest tossed thee here ashore,
Desolate yet all undaunted, on this desert land enchanted—
On this home by Horror haunted—tell me truly, I implore—
Is there—*is* there balm in Gilead?—tell me—tell me, I implore!"
Quoth the Raven "Nevermore."

(31) [An extract from 'Valentine' by Carol Ann Duffy, 2004]

Not a red rose or a satin heart.

I give you an onion.
It is a moon wrapped in brown paper.
It promises light
like the careful undressing of love.

Here.
It will blind you with tears
like a lover.
It will make your reflection
a wobbling photo of grief.

I am trying to be truthful.

Not a cute card or a kissogram.

I give you an onion.

Both *The Raven* and *Valentine* contain repetitions of whole lines. The repeated line *Quoth the Raven "Nevermore"* certainly predicates something about a discourse topic, a raven, which is mentioned explicitly early on in the poem. Perhaps the Duffy poem predicates something of a person who is loved very much, a 'Valentine'. We might argue that the Relevance Requirement has been met. However, this example of repetition raises problems with regards to the Graded Informativeness Principle. As set out above, this requirement stipulates that a globally coherent text will have segments which express propositions that are more informative than prior segments in that they yield information not expressed by those prior segments. In a strict sense, the repeated segments above could be said to violate the Graded Informativeness Principle in that they do not express any more information from a linguistic-decoding perspective than their original instantiations. If this is correct, then these repetitions would not be acceptable or well-formed. However, we feel they are, and the skilful deployment of such repetitions is, in part, what makes writer such as Poe appeal to us.

It is correct that a repeated segment technically expresses no more from a propositional perspective than its original. However, in a broad informational sense, a repetition can be more informative (or, *relevant*) than its original instantiation because it gives rise to the recovery of extra effects that were not on offer when the segment was processed the first

time. Remove the repetitions, however, and their effects disappear. Since we cannot say what these effects are in conceptual terms, they are *non-propositional*. The problem that we have is that repetitions such as (30) and (31) are technically not more informative as per the Graded Informativeness Requirement but, in cognitive terms, they *are* more informative in that they have new cognitive consequences for us. They make positive modifications to our cognitive environments (see chapter two, §2.2.3). As such, the Graded Informativeness Requirement and Giora's framework would need to be modified to include non-propositional effects and, I think, emotive or expressive effects. It would seem that Giora recognises that some adjustment to her model is required to account for some phenomena that arguably have similar effects to some of the repetitions I examine. Giora (1991) has proposed that jokes should be analysed as violating the Graded Informativeness Principle, for example. Jokes are argued within Relevance Theory (Yus, 2016) to lead to the kind of non-propositional effects that I am interested in. However, Giora seems to treat jokes as I imagine she would repetitions. Giora has no discussion of the effects of jokes outside of the semantic, the propositional, and the purely categorical. Yus (2016) has suggested that mutual parallel adjustment is vital to humour - if pragmatic processes are so involved in the recovery of those effects, it is likely that pragmatics will also shoulder a lot of the burden in the interpretation of repetitions.

Let's move on from the Graded Informativeness Principle as is, and see what the Relevance Requirement might buy us instead. Repetitions are acceptable as long as they predicate something about a clearly defined topic. However, there are repetitions which don't seem to exist to predicate a proposition *about* something. Instead, they are designed to express an emotion, or to yield other non-propositional effects (e.g., humorous), and while emotions can be directed towards something, and while non-propositional effects can be related to or triggered by a particular stimulus, it is not clear that these utterances are *about* something, or exist to be *about* something in the sense that this term is normally applied. Consider (32) and (33).

(32) [An extract from Monty Python's famous 'Spam' sketch, 1970]
Man: Well, what have you got?

Waitress: Well there's egg and bacon; egg, sausage and bacon; egg and **spam**; egg, bacon and **spam**; egg, bacon, sausage and **spam**; **spam**, bacon, sausage and **spam**; **spam**, egg, **spam**, **spam**, bacon and **spam**; **spam** sausage **spam spam** bacon **spam** tomato and **spam**...

(33) [Extracts from 'Scheherazade' by Carol Ann Duffy, 2012, p 7]
Dumb was as good as dead;
better to utter.
Inside a bottle, a genie.
Abracadabra.
...
Fact was in black and white;
fiction was a colour.
Inside a dragon, a jewel.
Abracadabra.
...
Imagination was the world.
clever to chatter.
Inside a she-mule, a princess.
Abracadabra.

It is possible to argue that (32) is 'about' Spam in a sense, but I struggle to say what is predicated about the topic other than there seems to be an awful lot of spam on offer. Clearly, the point of the repeated *spam*s here is to communicate non-propositional humorous effects. In (33), the title is the name of a Persian princess. You could argue that she is a kind of topic of which something is predicated. However, *abracadabra* is an incantation and could be analysed as an incantatory speech act. It would be seen in Relevance Theory as expressive (see chapter two, §2.8.4). The point is that these repetitions don't express *propositions* that are *about* anything in the sense intended by Giora. As they are currently formulated, it seems that the Relevance Requirement and the Graded Informativeness Principle predict that commonplace and pleasing repetitions are not well-formed, and are not coherent or acceptable. This is not the case, as attested by our love for poetry and literature featuring

repetitions, and the prevalence of various repetitions in everyday speech. Bearing all this in mind, this raises the questions of why these repetitions are acceptable (when others aren't e.g., (27a) and (27b) above), and how their effects actually contribute to any sense of acceptability or coherence for a text or discourse.

Theorists working on the coherence of discourse consider that it is understandable because it is well-formed. Relevance-theorists, however, consider that a discourse is acceptable only when it is understood – acceptability is a derivative notion here. Repetitions are judged acceptable in a context and are judged acceptable because of how they are recognised and processed within that context. The recognition and processing of different examples of repetitions is addressed throughout chapters three, four, and five of this thesis, and, essentially, it becomes clear that stylistic repetitions are acceptable and contribute to well-formedness because their ostensiveness creates expectations of particular effects which are then satisfied.

Thus far, I have set out the kinds of repetition that fall outside the scope of this thesis, and I have begun to raise questions about how stylistic repetitions are processed. However, I have not yet touched much on what the *effects* of repetition might be. A central aim of this thesis is to contribute to the literature in pragmatic stylistics by explaining the effects of such repetitions. To that end, it seems particularly fruitful to gather up statements about the putative effects of repetitions in order to to ascertain what issues to address.

1.6 The 'emphatic' and 'intensifying' effects of repetitions

1.6.1 Motivating a cognitive account of repetition

In this thesis, I do not focus on socially-grounded accounts of repetition. I will not discuss repetitions that signal or complete a turn in a conversation (e.g., Curl *et al*., 2003), or which are used to position oneself socially in communication, or to meet particular interactive goals (see Norrick, 1987; Kjellmer, 2008). This decision concerns the nature of the relationship between social accounts of communication, and cognitive accounts. Socially-grounded accounts of repetition sit within social accounts of communication, or what has been termed in linguistics *E-language*, or, 'External Language'. This is the set of habits and knowledge that

a language community shares. E-language is *not* the internal mentally-represented linguistic knowledge that an individual has of their language - this corresponds roughly to the idea of *performance* in Chomsky's competence/performance distinction (Chomsky, 1986), which is roughly analogous to *I-language* or 'Internal Language'. This internal linguistic knowledge is the proper object of linguistic enquiry. Performance is understood as how language is deployed in actual situations (Chomsky, 1965).

The relationship between E-language and I-language is that the latter is primary, and the former would not exist without the latter. We can only understand how people use language in society if we first explain what goes on linguistically inside people's minds. Sub-personal explanations are first necessary to understand what individuals do socially. It is possible to make the same case for communication. To understand the social, we need to look at cognition, at the sub-personal level. I do not offer a social account of repetition in communication because, first, we need a cognitive, sub-personal account of how repetitions are interpreted. 1.6.2 The conflation of emphasis and intensification

As mentioned earlier, there is, to my knowledge, only one book-length treatment of repetition in linguistic stylistics: SPELL (Fischer, 1994). Many contributors in this volume draw attention to an association with intensity or intensification. Aitchison (1994, p 19) says that '[r]epetition is used primarily for intensification'. Intensification, here, is understood as some kind of augmentation in quantity or quality, and may have an expressive function (1994, p 20). Similarly, Jucker (1994, pp. 52-53) notes that there are repetitions that serve an *intensifying* function, and that the repetition of intensifiers such as *very* can have an additional *intensifying* effect over and above that of a lone intensifier. Attridge (1994, pp. 61-81) focuses mostly on repetition in poetry, and, in particular, on the effects of repetition of metrical form. He (*ibid*.) notes the 'prolonging' effects of certain repetitions, an association with expectation and resolution, and with a sense of progression. He also notes that repetitions can build to a kind of climax. Finally, Attridge (*ibid*.) says that poems that include adjacent repetitions in their final line have a flavour of expressiveness or *intensity*.

In SPELL (Fischer, 1994), a link is also drawn between repetition and emphasis. Brian Vickers writes about repetition and *emphasis* in rhetoric, wherein rhetoric is taken to be a codified set of tropes and practices used in particular speech acts (Vickers, 1994, p 86). Vickers (1994, p 93) says that rhetorical figures such as repetition have been thought to share similarities with looks, gestures, and attitudes in terms of the embodiment of *emotion*, and proposes that repetition can be used to produce *force* and *emphasis*. He also reports Cicero's comment that repetitions are *vigorous* and *emphatic*. Vickers (1994, p 97) claims that repetition has long been heavily associated with *emphasis* and *intensity* of emotion. In this volume, multiple mentions are made of the emphatic nature of repetition – both in terms of what writers *do* (e.g., how Charles Dickens *emphasises* a point), and in terms of the so-called *emphatic effects* of repetition. I must highlight that none of these accounts explain in detail what is meant by intensity/intensification or emphasis, and, with the exception of Jucker's account which draws on Relevance Theory, no cognitive models of these reproductions of form are explored.

The authors in SPELL are by no means the only researchers to mention *intensification* and/or *emphasis* when describing the effects of repetition, or to mention prolongation, or strengthening, or expressivity.

The fact that Bolinger (1972) deals with repetition in a chapter called *Intensification by Stretching* suggests he considers the effects of the repetitions he addresses to be intensifying. 'Stretching' refers to a type of lengthening – in other examples, it concerns prosodic or phonetic lengthening, i.e., a manipulation of sound form (Bolinger, 1972, pp. 281-292). In Bolinger's (*ibid*.) cases of repetition, 'stretching' means stretching out semantic *features* of repeated words, rather than their sound. He (*ibid*.) considers cases like *they walked and walked*, and says '[s]aying something twice not only doubles it semantically but also doubles the noise with which we say it, and noisiness is certainly one form of *intensification*' (Bolinger, 1972, p 288). We must assume, from this characterisation, that the effects of such repetitions affect the processing of what is decoded from the semantics of an utterance. However, it is not made clear what the term 'semantic doubling' means. Such an account seems to chiefly apply to cases where some kind of lexically encoded concept is somehow intensified, since the discussion is restricted to semantics. Bolinger (1972, p. 290) also notes that when 'intensifiers' or certain adjectives are repeated, as in *that's* ***very, very*** *interesting* or *it was a*

big, big *bear*, there is an additional intensificational effect. All these cases are identified by Bolinger as cases of 'semantic repetition', but it is not really clear what is meant by this term. Repeated adjectives and repeated intensifiers of a particular kind are addressed in chapter five.

In Leech and Short's (1981) *Style in Fiction*, repetition of words is considered a cohesive device, just as in Halliday and Hasan's (1976) framework. They (1981, p 247) note, however, that '[r]epetition is *expressive* in that it gives *emphasis* or *emotive heightening* to the repeated meaning'. Preminger *et al*.'s (1993, pp. 1035-1037) comments on repetition are descriptive; they list 'types' of repetition. However, the authors note that repetitions can be 'a force for *continuation*', and that repetition can create patterns of expectation that are *strengthened* and fulfilled with even further repetition tokens (Preminger *et al*., 1993, p 1036). They (*ibid*.) say '[w]hen a line, phrase, or even a sound is repeated, the experience of the first occurrence is continuously *maintained* in the present'. Finally, Preminger *et al*. (*ibid*.) say that repetition can be *expressive*, and incremental, suggesting that repetitions can give rise to effects or meanings which are prolonged or maintained. Again, however, there are no cognitive explanations of how this is achieved.

Brody (1986) notes that single-speaker repetitions can 'express emphasis', indicating that emphasis is some kind of meaning or effect. Gerleman (1951) considers that repetitions can also yield *emphasis*, which also suggests emphasis is part of an interpretation. Koguchi (2009) discusses lexical repetitions in Dickens and claims that repeated words *achieve* a kind of special emphasis. Nadarajan (2006) mentions that there are repetitions which *yield a sense of emphasis*, Ulatowska *et al*. (2000) discuss repetitions *for emphasis*, while Tannen (1983) notes that repetition of longer utterances can be exploited to *emphasize a point*. Bazzanella (2011) draws a link between repetition and *intensity*, but also states that repetition can be associated with rhetorical *emphasis*. Some authors, then, seem to see emphasis as an *effect*, while some see it as something communicators *do*.

1.6.3 Repetition within Relevance Theory

In *Relevance: Communication and Cognition* (1995), the intention is not to provide a theory of repetition. Repetition is only mentioned to demonstrate that the link between a device and its interpretation is pragmatically inferred in context, and that there is no one-to-one link between devices and effects. Nevertheless, there are a number of observations made by Sperber and Wilson about repetition. In terms of the effects of repetition, or, rather, the case of epizeuxis, they (1995, p 219) say that the *emphatic* effects of repetition are achieved in different ways. Thus, here, Sperber and Wilson imply that repetitions have effects which are *emphatic*. Emphasis seems to be framed as an aspect of an interpretation. Diane Blakemore (2004, 2008, 2009, 2011) has also mentioned the effects of repetition in discussions of coherence, free indirect thought, appositions, and expressive meaning, and often mentions that the effects of repetitions have something in common with the effects of other devices.

Blakemore (2008) follows Sperber and Wilson (1995) in saying that it is difficult to 'pin down' the effects of repetitions in propositional terms. In her paper on apposition (2008), there are a number of points made which are of interest concerning the effects of stylistic repetitions. Blakemore (2008, p 37) notes that apposition can communicate an *impression of emphasis or intensification*, implying that both of these things are some kind of outcome of interpretation, or effect. Blakemore (2008, p 41) does, however, note that it is not clear that all appositions should be understood in terms of emphatic effects. She suggests that examples such as *he was depressed, flattened*, are better understood in terms of intensification, since *flattened* is a more intense kind of *depressed*. However, along with many of the authors mentioned in this section, the terms *intensification* and *emphasis* seem to be used in a general and intuitive sense – these notions are not defined in cognitive terms.

Blakemore (2008, p 42) does seem to say that certain appositions give rise to impressions of intensification or impressions of heightened vividness that are 'different from the effects yielded by repetitions'. This implies that emphasis and intensification might be distinct entities. Blakemore (2008, p 43) does, however, make the point that intensity and emphasis are *not* constituents of mental representations themselves, i.e., intensity and emphasis are not part of propositional representations. I build on this claim and argue that intensity and

emphasis, as understood with respect to stylistic repetitions, are not meanings, or effects in themselves at all.

The intention here was not to provide a full account of everything that these authors have to say about analysing repetitions of linguistic form. I just wanted to show that, within and outside Relevance Theory, the words *emphasis*, and *intensification* or *intensity* crop up regularly in discussions of repetition. I also wanted to show that it is not clear what intensification or emphasis are taken to mean, but that many theorists consider that the effects of repetitions are *emphatic* or *intensified*. Thus, the terms *emphasis* and *intensification* mostly seem to be applied to the outputs of utterance interpretation – not the inputs or processes involved. There is clear conflation of the terms in the literature. They are almost used interchangeably, or seen as types of one another, as when Blakemore (2008) implies that intensity is a form of emphatic effect, or when Vickers (1994) says that repetition is associated with emphasis *and* intensity of emotion.

People can emphasise the importance of doing something, emphasise a point, and emphasise cheekbones. In art, the term emphasis applies when a pictorial element is made prominent by an artist. In linguistics, emphasis is generally understood as speakers placing prosodic prominence on a word by manipulating pitch or volume, for example (Cruttenden, 1997). You can emphasise the topic of a sentence by pre-posing it, by bringing it to prominence. All of these are things that people *do*. Emphasis, then, should be a type of behaviour. It cannot be an effect. (Intentional) behaviours in communication produce effects and, as such, cannot *be* effects. Forms and effects are distinct. One is the input to the interpretation process, while the other is the output.

1.7 The road ahead

Above, I showed that repetition is associated with both emphasis and intensification, but that these terms are not defined, and are certainly conflated at times. The discussion of repetition in this thesis makes it possible to provide distinct and cognitively-driven explanations of both phenomena. I investigate the hypothesis that emphasis and intensification are distinct entities. The former is a particular type of human behaviour. The latter is a processing

phenomenon, and also not an effect. Loosely put, intensification emerges when we expend more effort on adjustment or calibration processes in interpretation. Emphatic behaviours can encourage hearers to expend effort on particular processes during the interpretation process, and so intensification is characterised in terms of processing. I argue for unitary accounts of both emphasis and intensification and I will demonstrate how the study of repetition contributes to these accounts. However, at the same time, I do not assume that every repetition of form is emphatic, or that every repetition achieves relevance in the same way. In particular, I draw a distinction between the way repetitions across distinct intonation groups, and repetitions of form within a single intonation group are interpreted. The ideas of expressive or emotive meanings or functions of repetition were raised by some of the authors above. I hope to explore, if briefly, any connection between expressivity and repetition, and consider what that reveals about the nature of emphasis.

Chapter four addresses repetitions across intonation groups, and includes both adjacent cases (epizeuxis) and (very) non-adjacent cases. I show there is intensification in the identification of implicit content, and show how the intonation boundaries involved in the examples interact with the context and the Principle of Relevance (see chapter two, §2.4.2) to yield effects which must be accounted for at the level of implicit communication. This chapter is important in explaining why we notice stylistic repetitions, and why we decide that they are intended to communicate particular effects.

Chapter five addresses repetitions of linguistic form within intonation groups, e.g.:

(34a) He went for a **long long** walk.

(34b) Would I like to go to the cinema? **Yes yes YES**!

Many of these repetitions can be analysed in terms of their contribution to the explicit content of the utterance. However, this discussion leads to the discovery that there are items which are repeated in communication whose meanings cannot be analysed in conceptual terms, e.g., expletive *fuck* or *shit*, and non-linguistic examples such as repeated emoji, giving us a more general account of repetition in human behaviour.

Before exploring different kinds of repetition data, in chapter two, I outline the pragmatic framework which underpins this thesis, Relevance Theory, including some key ideas concerning emotions and expressivity. This information plays a supporting role in the analysis of the data in chapters three and five. In chapter two, I set out the distinction between semantics and pragmatics, which is particularly important for the analysis of reduplication in chapter three. I also explain what is meant by ostensive communication and ostensiveness, as these notions are vital for my account of emphasis, as well as the discussion of stylistic repetitions in general. Chapter two also contains a discussion of the relevance-theoretic distinction between explicit and implicit content which is essential for the distinctions drawn in chapters four and five. This chapter also sets out key assumptions of Carston's (2002) lexical pragmatics, which underpins the discussions of so-called contrastive focus reduplication in chapter three, and repeated modifiers in chapter five. A discussion of procedural meaning allows me to suggest how we interpret repeated non-conceptual items repeated within the same (intonation) group in chapter five. There is also discussion of the *showing-saying* continuum, which sets out the relationship between our communicative behaviours and what we want to communicate, or, rather, between the kinds of evidence we present for what we want to communicate and what we communicate. The notion of *showing* is crucial to this thesis. It helps to explain how many stylistic repetitions achieve relevance, and is the cornerstone of my explanation of emphasis. The account of style and stylistic effects provided in this framework chapter underpins the entire discussion of the way the stylistic repetitions achieve their effects.

Chapter Two: Relevance Theory: a Framework for Analysing Repetition, Emphasis and Intensification

2.1 Introducing Relevance Theory

This thesis is carried out within the framework of Relevance Theory (Sperber & Wilson, 1995). The study addresses both linguistic issues, and problems from different fields of linguistics. However, Relevance Theory (Sperber & Wilson, 1995) is not a theory about language, but communication, and about cognition. It is a theory about how we might cognise when communicating. Both communication and cognition are linked with language, but neither requires language. We can communicate without language, and we can think without language. Relevance Theory is a cognitive theory of pragmatics. It is worth beginning this chapter by stating briefly how Relevance Theorists view the mind, and what is meant by a cognitive theory of pragmatics. Relevance Theorists (Sperber & Wilson, 1995) hold that the mind is both representational and computational. It is representational because we are able to mentally represent things – objects, states of affairs in the world, and, as we shall see, other people's beliefs, desires and intentions. The mind is computational because we can perform computations over these representations. Some of these processes are deductive inferences. For example, if you tell me *if it rains, you must take an umbrella*, I can deduce from looking outside whether or not I should take an umbrella. Deductive inference plays an important role in how Relevance Theorists view communication and cognition (Sperber & Wilson, 1995). In the course of this chapter, we see that inference plays a significant role in how we understand each other during communication. More specifically, we see that inference plays an equally significant role in how we interpret utterances containing linguistic material. Relevance Theory is a cognitive pragmatic theory because its account of communication is based on general cognitive principles (Sperber & Wilson, 1995). General tendencies, mechanisms and heuristics are responsible for how we go about communicating and, in cases of verbal communication, processing utterances of linguistic material (Sperber & Wilson, 1995).

At the heart of this thesis is an exploration of the relationship between linguistic form and meaning. As such, this chapter must set out what Relevance Theory has to say about aspects of linguistics that are concerned with meaning. In particular, it must explain how Relevance Theorists view semantics and pragmatics. However, I will not start with a discussion of the

semantics-pragmatics interface. The relevance-theoretic understanding of the semantics-pragmatics interface is couched in its notion of communication: ostensive-inferential communication (Sperber & Wilson, 1995). This chapter begins by explaining what communication cannot be like, and what it *must* be like in order to accommodate our experiences of it. I show that communication involves inference as well as coding. Then, I outline how semantics has traditionally been viewed. In doing so, I discuss what is meant by *context*. The problems that emerge from these conceptions of semantics, pragmatics and context are stepping-stones to a picture of utterance interpretation where, generally, *semantics* comes to mean *decoding*, and *pragmatics* comes to mean *inference*. The view of utterance interpretation that emerges is one in which pragmatics is fully inferential, and bears much more responsibility for the recovery of interpretations than previously thought.

Following this, I introduce the general cognitive principles underpinning Relevance Theory, and the notion of *relevance*. I then outline a particular procedure by which humans process communicative stimuli - the relevance-theoretic comprehension procedure. Once these issues have been addressed, I discuss what Relevance Theorists consider to be the relationship between form and meaning. In examining how relevance theorists see the relationship between utterances and thoughts, we see there are stronger and weaker degrees of communication, and that the types of thoughts communicated are not all the same. On the one hand, we communicate propositions but, on the other, we can also communicate vaguer interpretations, and impressions – the sort associated with many repetitions.

Throughout the chapter, I present data which a code model or propositional model of communication cannot explain but which are clearly communicative. It becomes clear there is much more to communication than just *the sharing of propositional representations*. It can lead to humans sharing cognitive environments and feelings or emotions. Thus, expressivity and emotion are also discussed. The chapter ends with an explanation of some basic ideas in intonation (which are added to in chapter four), and a brief introduction to the software employed for analysing intonation in this thesis.

2.2 Communication and context

2.2.1 The code model of communication

On the first page of *Relevance: Communication and Cognition*, Dan Sperber and Deirdre Wilson (1995, p 1) say:

> 'Communication is a process involving two information-processing devices. One device modifies the physical environment of the other. As a result, the second device constructs representations similar to representations already stored in the first device. Oral communication, for instance, is a modification by the speaker of the hearer's acoustic environment, as a result of which the hearer entertains thoughts similar to the speaker's own. The study of communication raises two major questions: first, what is communicated, and second, how is communication achieved?'

In this section, we see how a traditional model attempts to answer the latter question. Sperber and Wilson (1995) call this the 'code model' of communication. For now, I assume that it is thoughts that are communicated.

According to the code model, thoughts are converted into utterance signals by means of a code. A code is a system for pairing messages to signals (Sperber & Wilson, 1995, p3). A Morse-code operator converts his message into a signal for transmission by way of the Morse system, the code. The message is transmitted to another operator who decodes the message from the signal using knowledge of the code. It might be thought that human verbal communication works in the same way. This is the basis of semiotic approaches to communication (Pike 1967; Saussure 1974; Leach 1976). On this view, if I want you to think a particular thought, I can take my thought and convert it into a signal utterance using a code, language. You receive my signal and convert it back into the thought I want you to identify using the rules of the system. This idea, as Sperber and Wilson (1995, p5) note, is old, but is alive and well in the minds of everyday language users. We talk of 'getting our ideas across'. We struggle to 'put our thoughts into words'. However, such an attractive idea is only an idea. While the code model seems neat and explanatorily adequate, it is descriptively inadequate; the *how* of the model does not match the facts. As Sperber and Wilson (1995, p 6) note, 'comprehension involves more than the decoding of a linguistic signal'. They (1995, p9) go on to say '[i]t is true that a language is a code which pairs phonetic and semantic representations

of sentences. However, there is a gap between the semantic representations of sentences and the thoughts actually communicated by utterances'. Examples (1a-1b) illustrate this:

(1a) He's too tired. (Too tired for what? To watch TV late at night? To carry on living?)

(1b) It's cold in here. (The hearer should infer to close the window.)

The utterances in (1a) and (1b) do not encode everything the speakers want to communicate. How should the hearer of (1b) know she should close the window? We have to say that she works it out. However, she can only work it out drawing on the intentions behind the speaker's utterance. Commonplace utterances such as (1a) and (1b) do not linguistically encode everything that we want to communicate. That gap is filled by inference, and is guided by the recognition of speaker intentions.

2.2.2 Towards an inferential approach to communication

In the late 1950s, the philosopher H.P. Grice (1989) was attempting to sketch out a theory of 'meaning'. In order to find out what speakers mean by utterances, Grice had to first make observations about communication. In *Meaning*, Grice (1989) shows that understanding what someone means by a communicative behaviour depends on intention recognition (even though his model of intention recognition in communication is not quite correct). His characterization of a communicator's intentions is as follows:

In order to mean something by the utterance *x*, an individual *S* must intend the following:

(a) *S*'s utterance of *x* to produce a certain response *r* in a certain audience *A*;

(b) *A* to recognise *S*'s intention (a);

(c) *A*'s recognition of *S*'s intention (a) to function at least part of *A*'s reason for *A*'s response *r*.

(This characterization, due to Strawson (1964), is cited by Sperber & Wilson, 1995, p 28.)

Imagine that *S* wants *A* to come to think that it is raining outside. If *S* utters *It is raining outside*, that utterance can certainly result in this response. Condition (a) is fulfilled. *A* will likely realise that *S* intended *A* to think that it is raining outside. Condition (b) is also fulfilled. *S*'s intention

(a) might interact with how and why *A* comes to think that it is raining. It seems that condition (c) is also fulfilled. Clearly, intentions are involved in *A* understanding what *S* wants to communicate by his utterance. The question is whether or not all three intentions are necessary conditions for communication?

As Sperber and Wilson (1995, p 28) point out, only condition (b) needs to be fulfilled for communication to obtain. Conditions (a) and (c) might well be fulfilled, but they need not be. Consider (2):

(2) I have five A grades at A-level.

I intend to make you believe that I indeed have five A grades at A-level - my (a) intention. You might recognise that I intend you to hold this belief, but you decide (rightly) not to believe me. This means that you understood me without holding the belief that I want you to hold. You have to understand what proposition I wanted to communicate in order to decide whether or not to believe it. Since you did recognise my intention, (b) is fulfilled. (a) is not fulfilled, however, because you do not believe me - the intended response. Intention (c) does not enter into the picture because you did not produce the intended response - believing I have five A grades at A-level.

Communication can succeed without intention (a) or (c), and the Grice-Strawson view of communication is thus not accurate. Nevertheless, it demonstrates that speaker intentions are crucial in inferring what speakers intend to communicate. Sperber and Wilson (1995, p 29) develop this idea, showing that intention (b) is the only truly *communicative intention* - 'the intention to have one's informative intention recognised'. It is this intention which causes intention (a) – *the informative intention* – to be fulfilled. If you recognised the communicative intention, you necessarily recognised what that intention is about, the informative intention. The communicative intention is a second order intention and the informative intention is a first order intention. The communicative intention is an intention about what should be communicated.

There is a gap between what our utterances linguistically encode, and what they actually communicate. This gap is filled using inference, and intention recognition is crucial in determining how the gap is closed inferentially. Grice (1989) wanted to use observations

about communication to help him work out what utterances mean. He (*ibid*.) wanted to use his conception of what speakers mean by utterances of linguistic material to speculate on semantics. Sperber and Wilson (1995, p 21) note that this is a possible avenue of enquiry but say '...Grice's model can also be used as the point of departure for an inferential model of communication, and this is how we propose to take it'.

2.2.3 Context in communication: the problem of context selection

Sperber and Wilson (1995, p 15-16) say:

> '[t]he set of premises used in interpreting an utterance (apart from the premise that the utterance in question has been produced) constitutes what is generally known as the *context*. A context is a psychological construct, a subset of the hearer's assumptions about the world. These assumptions, rather than the actual state of the world, affect the interpretation of an utterance. A context in this sense is not limited to information about the immediate physical environment or the immediately preceding utterances. Expectations about the future, scientific hypotheses or religious beliefs, anecdotal memories, general cultural assumptions, beliefs about the mental state of the speaker may all play a role in interpretation.'

This view of context as a construct raises the question of how we choose the right contextual assumptions for use in utterance interpretation. Consider (3) below:

> (3) Laura: Would you like another pint?
> Chris: Beer makes me want to dance.

In a context which includes the assumption that Chris does not like dancing, Chris's response is interpreted as a refusal. In a context which contains the assumption that Chris likes dancing, Chris's answer is as an acceptance. How does Laura know which context is intended? Potentially, this is an impossible problem. Hearers have access to an enormous range of contextual assumptions, any of which could be brought to bear upon the interpretation process. How do hearers choose the right contextual assumptions on the fly? The Mutual Knowledge Hypothesis (MKH) (Clark & Marshall, 1981) restricts contexts to the set of assumptions that are *mutually known* to both speaker and hearer.

(4) [Laura and Chris having a late dinner at a restaurant. Both are regulars, and know that it closes at 11pm daily. It is now 10.55pm.]

Laura: This has been a lovely meal.

Chris: Yes, but it's almost time to go home, isn't it?

To interpret Chris' utterance in (4), Laura must access the assumption that the restaurant closes at 11pm. According to the MKH, this assumption is mutually known to Laura and Chris because they both visit the restaurant regularly. When producing his utterance, Chris knows that Laura knows when the restaurant closes. When interpreting Chris's utterance, Laura knows that Chris knows when the restaurant closes, which means that he knows that she knows this, and she could use that assumption in utterance interpretation. Chris can be sure that Laura will use that assumption in utterance interpretation because he knows that she knows that he knows she knows that assumption, and so on, *ad infinitum*. There is an obvious problem here. Establishing mutual knowledge takes an indefinite amount of time. Moreover, we soon lose track of who-knows-who-knows-what because our minds cannot cope with more than a few such layers. Perhaps communicators don't need to establish full mutual knowledge for communication to obtain. Maybe people only establish mutual knowledge to a degree. What degree is this?

Clark and Marshall (1981) offer another solution: there is a finite inductive procedure for identifying mutual knowledge. A speaker and a hearer can assume mutual knowledge of a given proposition if they are jointly present in a situation which supplies evidence for the truth of that proposition. Physical co-presence can supply direct evidence for mutual knowledge (Clark & Marshall, 1981). In the case of (3), this is the co-presence of Laura and Chris in a situation where Chris actively avoids dancing. Linguistic co-presence is considered less direct evidence for the truth of a proposition. For (3), Chris might have told Laura that he hates dancing on a previous occasion. Membership of a particular group or community is also thought to provide evidence (Clark & Marshall, 1981). Perhaps Laura and Chris are part of a family in which Chris's hatred of dancing is long-standing family joke.

This solution improves upon the stronger MKH. However, this view of mutual knowledge assumes that all knowledge is based on evidence of observable data. This is not the case. To acquire knowledge is to acquire a mental representation, and the construction of that

representation requires an individual to draw on assumptions which are not part of that observable data – say, assumptions from memory. Different people who are co-present in the same physical environment may construct totally different representations of that environment. Reconsider (3). Suppose that Laura has a poor memory and cannot remember whether or not Chris likes dancing. Although she was co-present with Chris in many situations where his attitude towards dancing was raised, she cannot access the assumptions needed to work out whether or not Chris's answer is a refusal or an acceptance. As such, co-presence does not guarantee that individuals access or represent the same information in the same way. The only way two individuals could actually have mutual knowledge of any event is to have access to each other's memories. This is impossible.

The MKH approach also assumes speakers and hearers do not proceed in communication until both have a guarantee of communication succeeding. In (4), Laura and Chris would establish that they both knew the assumption about closing time before proceeding in their interaction. Proponents of this view seek a heuristic that *guarantees* successful communication. But communication does not succeed all the time. Speakers take risks. Hearers are often ill equipped to process what speakers say. A speaker need only utter a word that the hearer has never heard before and communication could break down. There is no fail-safe way to ensure successful communication. But it still works. Moreover, speakers seem to proceed with communication without establishing mutual knowledge beforehand.

(5) Julia: I think I'll have the broccoli bake.

Roz: I won't! I never eat brassicae.

Julia is not familiar with the term *brassicae*. She can't know in advance that broccoli is a member of this vegetable family. Yet, Julia can still understand Roz, and communication succeeds. Julia can recover the assumption that broccoli is a member of the brassicae family (and, thus, Roz won't be eating it) as a result of understanding the utterance, *not as a prerequisite for being able to interpret it*. In other words, the identification of particular contextual assumptions is not required for successful communication, but emerges as a result of it.

The main problem for the MKH is that it aims to posit rules for guaranteeing successful communication – just as the code model of communication does through the use of shared

codes. However, the MKH approach doesn't actually reflect how communication takes place. Communication always proceeds at a risk. Moreover, the MKH approach is mistaken in thinking that successful communication requires the duplication of thoughts. As we saw above, two people may communicate about the same event even if their experiences or memories of that event are not identical. Finally, it is perfectly possible that two people have similar experiences of an event, but these individuals need not share mutual knowledge about it. However, the fact remains that people do *share* information through communication, and so, what is needed is an explanation of how this obtains. Sperber and Wilson (1995) proposed the notion of *manifestness* to that end. They (1995, p 39) say that a fact is manifest to an individual if and only if that individual can mentally represent that fact and hold it to be true or probably true. For a fact to be manifest, it need not actually be mentally represented. It must simply be *represent-able*. This notion of manifestness supports the relevance-theoretic view of the context outlined above. Pre-determined contexts are unnecessary. Contextual assumptions must simply be manifest to an individual should she need to represent them. A hearer can then choose just the assumptions from the set manifest to her that she should bring to bear on the interpretation process. The set that is chosen is selected based on the notion of *relevance*, which I introduce shortly.

The set of assumptions that are manifest to an individual is called a *cognitive environment* (Sperber & Wilson, 1995). If I want to communicate something to you, I just have to make it manifest to you. By making something manifest, I alter your cognitive environment. If I'm in a rose garden and I wish to communicate that the roses smell beautiful, I need only catch your eye and breathe in deliberately and clearly, which makes more manifest the assumption that I think the roses smell nice, among other things. This is all well and good, but it remains to be shown that the notions of manifestness and cognitive environment allow us to share information without suffering the same problems as the MKH approach. Communication does not require mutual knowledge, nor does it require duplication of thoughts. Communication is about the mutual adjustment of cognitive environments. In the rose garden, if I catch your eye and breathe in, it becomes mutually manifest to both of us that this has happened, and it becomes mutually manifest to us that we entertain similar assumptions about the rose garden. Do I need to know exactly what assumptions you entertain? Must we entertain exactly the same assumptions? No. In this case, it just needs to be mutually manifest to us

that we are likely entertaining similar thoughts about the rose garden. All that matters is that it is mutually manifest that our cognitive environments overlap to a degree that allows communication to proceed. The set of mutually manifest assumptions between two communicators is their *mutual cognitive environment* (Sperber & Wilson, 1995).

If we say that communication concerns the manipulation of mutual cognitive environments, we can explain how humans share information without having to claim that the context for utterance interpretation is mutually known. This avoids infinitely long chains of thought where speakers and hearers have to work out who-knows-who-knows-what. By what mechanism can an individual cause an assumption to become more manifest to another individual? In the rose garden, I sniffed in such a way as to attract your attention in order to communicate with you. It is to such attention-attracting behaviours that I now turn.

2.3 Ostensive-inferential communication

Suppose we are waiting for an old friend to arrive at a coach station. I am standing by the road and have a better view of approaching traffic. I am between you and your view of the road. Suppose I see our friend's coach approaching, and wish to communicate this. I could say something. But, I can also step deliberately and clearly to the side to give you a clear view of the approaching coach. In stepping to the side, I alter your cognitive environment. I give you the opportunity to pay attention to the approaching coach, which, of course, you might do or you might not. However, it is highly likely in this case that two things happen. First, it is highly likely that you *do* notice the coach. Second, it is highly likely that you realise that I intended you to do so because of my *ostensive* behaviour. By deliberately stepping aside, one of the facts that becomes more manifest to you is the fact that I deliberately stepped to the side. It was no accident, so you decide it is worth paying attention to it. You conclude that such a deliberate and substantial movement has been carried out for a reason. You begin to look around for reasons as to why I moved. Was there someone in the way? Or, is our friend's coach nearing the coach station where we are eagerly awaiting him? By stepping to the side, I make certain facts more manifest to you, and one of those facts is that the coach is approaching. Sperber and Wilson (1995, p 49) say '[w]e will call such behaviour - behaviour which makes manifest an intention to make something manifest - *ostensive* behaviour or simply *ostension*'. My stepping to the side made manifest my intention to make the

approaching coach manifest, and my behaviour was ostensive in this sense. Sperber and Wilson (1995, p 49) explain how individuals work out what has been made manifest to them by way of an ostensive behaviour. Ostensive acts attract attention. By behaving ostensively, I suggest that there is a stimulus that you should process. However, as there is a tendency to optimise cognitive efficiency, you won't process a stimulus unless you think there is some reason for you to process it or, in Sperber and Wilson's (1995) terms, there is some cognitive reward on offer. There is no rational reason for me to draw your attention to an approaching bus unless that is beneficial to you. There is no reason for me to draw your attention to a stimulus unless it is *relevant*. All ostensive acts come with a tacit guarantee of *relevance* (Sperber & Wilson, 1995, p 49). Relevance is defined in §2.4 below. In a pre-theoretical sense, the approaching coach was relevant to you because it allowed you to infer the information that our friend was arriving.

Above, it was explained that communication depends, in part, on the ability of the audience to make inferences. Amongst the assumptions inferred by the audience will be assumptions about the speaker intentions that underpin his utterance(s). Speakers help hearers in the recognition of these intentions by behaving ostensively. Sperber and Wilson (1995, p 50) say '[o]stension provides two layers of information to be picked up: first, there is the information which has been, so to speak, pointed out; second, there is the information that the first layer of information has been intentionally pointed out'. Naturally, it is possible that the first layer of information could be recovered without the need for recovery of the second. In the coach example above, you may notice that our friend's bus is approaching whether or not you notice that I was intentionally 'pointing out' this information to you. However, recognising the intentions behind ostensive acts is usually required for the successful recovery of the first layer of information. After all, it is perfectly reasonable to say that if you had not asked yourself why I stepped aside, you might not have noticed the coach. Thus, communication in Relevance Theory is termed *ostensive-inferential*.

I have shown so far that communication is not about the duplication of thoughts and it involves much more than just a code. Communication also involves providing evidence from which hearers can derive representations of speakers' thoughts using inference. Contexts for utterance interpretation are constructed on the fly, and relevance-theorists consider that the purpose of communication is to modify the cognitive environments of co-communicators.

Cognitive environments are modified by behaving ostensively to influence the set of assumptions which are manifest to an individual at a given time. Ostensive behaviours involve making manifest an intention to make something manifest. Ostensive acts can be acts of *showing*, acts of *saying*, or some blend of these. For example, I may make manifest the fact that we have nothing to eat tonight by *saying* that we have no food, or by opening the fridge theatrically and *showing* you the empty fridge. I come back to the *showing*/*saying* distinction in §2.8, as neither stylistic repetitions nor emphasis can be explained without it.

Why does ostensive behaviour succeed in modifying cognitive environments? What is so important about the fact that communication is ostensive in the explanation of how people communicate? According to Relevance Theory, the answer lies in the fact that ostensive behaviour comes with a guarantee of *relevance*.

2.4 Relevance and communication

2.4.1 The cognitive principle of relevance

The cornerstone assumption of Relevance Theory is that the way in which hearers understand utterances (or any case of ostensive communication) is constrained by a specific principle which is grounded in a more general principle governing human cognition. Sperber and Wilson (1995) ask us to think about how we process information in general before considering communication. If I'm out hunting on the savannah and hear a loud roar, I don't waste time worrying about whether to expend valuable seconds and cognitive effort in paying attention to this sound. Clearly, any effort will pay off in the form of assumptions which have an obvious advantage, e.g, there is a nasty predator nearby which is about to kill me and I should probably run away. The same sort of point also applies to less dramatic cases. Suppose Chris is at a party where someone decides to put on some music, and clears away the furniture. Why would Chris pay attention to this?

According to Sperber and Wilson's (1995) Cognitive Principle of Relevance, the fact that Chris pays attention to this information is explained as soon as we assume that his cognitive system is 'geared towards the maximization of relevance' (Sperber & Wilson, 1995, p 260). That is, Chris' cognitive system is geared towards getting as much cognitive reward for as little cognitive effort as possible. Cognitive reward is a useful improvement to our representation

of the world. In the above case, Chris might derive rewards in the form of assumptions such as [There is going to be dancing] and [It is time for me to go]. Sperber and Wilson (1995) call these rewards cognitive effects, identifying three kinds of effect (see Blakemore, 2002, pp. 94-96):

> **Deletion (of an assumption)**: My housemate has gone to the shop to buy milk. Out of my window, I see him coming down the path with bread in one hand, and eggs in the other. I delete my assumption that my housemate bought milk.
>
> **Strengthening (of an assumption)**: My housemate has gone to the shop to buy milk. I see my housemate in the distance, carrying something white. As he approaches, I see he is carrying milk. My assumption that my housemate has bought milk is strengthened.
>
> **Contextual implication**: I am at a train station to take the 10 am train to Manchester. I look at my watch and see it is 10.15. I draw the contextual assumption that I have missed the train. (A contextual assumption is drawn when an inference is drawn in which 'new' information supplied from the context and 'old' assumptions make up the premises.)

Sperber and Wilson's (1995, pp. 142-151) claim is that the more cognitive effects derived the greater the relevance of the information to an individual. However, they also point out that relevance is a function of the effort expended in obtaining those cognitive effects, and that the greater the effort expended, the less relevant the information will seem. Someone who sees a headline *Tourist Escapes Wild Lion* but has never experienced lions outside a zoo can probably derive some cognitive effects; however, this will probably take considerable effort of imagination. In contrast, for someone who *has* heard a lion prowling around their tent while on a safari holiday, this will take considerably less effort of imagination and the headline will be that much more relevant.

The Cognitive Principle of Relevance governs the way we approach all information and so it also applies to information that has not been communicated. However, this principle provides an empirical basis for a more specific principle which, according to Sperber and Wilson (1995),

explains the way in which we approach information which *has* been ostensively communicated.

2.4.2 The communicative principle of relevance, optimal relevance, and the relevance-theoretic comprehension procedure

If Chris decides to pay attention to the dance music and the removal of the furniture at the party, he does so because he has hopes of relevance. Of course, there's no guarantee that his attentions will be rewarded. However, when it comes to *ostensive* communicative behaviour, people don't pay attention because they have *hopes* of relevance. Ostensive behaviour creates *an expectation of relevance*. Individuals who act in an ostensive and attention-attracting manner are presumed to have something to communicate that is worth paying attention to. It is just not rational for a person to overtly demand someone's attention if they have no useful information to share. An individual's act of ostensive communication carries the presumption that the audience can derive at least some cognitive effects from it. Of course, this presumption is not necessarily valid. I may gesture to an empty seat in a lecture theatre as you enter, not realising that you have already seen the seat. However, the point is that you can assume that I would not have engaged in this behaviour unless I believed the information would have cognitive effects for you. I should note that different things are more or less relevant to different individuals (Sperber & Wilson, 1995, pp. 142-151). I find utterances about pragmatics more relevant, all other things being equal, than utterances about sociolinguistics - the reverse would generally be true for a sociolinguist. Moreover, stimuli are not relevant in and of themselves. Relevance is evaluated, and it is a property that emerges from offsetting the effort expended on processing a stimulus against the cognitive rewards that processing yields (Sperber & Wilson, 1995, pp. 142-151).

General human cognition is concerned with the maximisation of relevance. However, a hearer is not entitled to expect *maximal* relevance. It may not be within a speaker's capabilities to produce the most hearer-friendly utterance – he may be nervous, tired, or overwhelmed. Moreover, a speaker may not be able to produce the utterance that has the most cognitive effects. When my students ask what is on their pragmatics exam, I could physically tell them. If I did so, however, I might get fired for academic misconduct. Nevertheless, according to Sperber and Wilson (1995, pp. 157-160), the expectation of relevance communicated by an

act of ostensive communication goes further than the expectation of *some* cognitive effects. What Sperber and Wilson end up with is a guarantee of maximal relevance within the parameters of speakers' abilities or interests, or, in other words, optimal relevance. An utterance is optimally relevant if and only if:

> (a) The set of assumptions which the communicator intends to make manifest to the address is relevant enough to make it worth the addressee's while to process the ostensive stimulus.
>
> (b) The ostensive stimulus is the most relevant one the communicator could have used to communicate (Sperber & Wilson, 1995, p 270)

The Communicative Principle of Relevance is simply the claim that 'every act of ostensive communication communicates a guarantee of its own optimal relevance' (Sperber & Wilson, 1995, p 260).

Once again, it does not follow from this principle that every utterance is in fact optimally relevant. Not every student manages to write reader-friendly text, and there are always people who forget to tell you the most important information. Moreover, we can all be mistaken about the contextual and processing resources of our audiences. The point is simply that if a hearer recognises that a speaker has communicated with them ostensively, they can assume that the speaker believed he was being optimally relevant. Wilson and Sperber (2002, pp. 258-259) say:

> '[t]he communicative principle of relevance and the definition of optimal relevance suggest a practical procedure for performing [certain] subtasks and constructing a hypothesis about the speaker's meaning. The hearer should take the linguistically encoded sentence meaning; following a path of least effort [she] should enrich it at the explicit level and compliment it at the implicit level until the resulting interpretation meets [her] expectation of relevance.'

This strategy is formalised in the relevance-theoretic comprehension procedure (*ibid*.):

(a) Follow a path of least effort in computing cognitive effects. Test interpretive hypotheses (disambiguations, reference resolutions, implicatures, etc.) in order of accessibility.

(b) Stop when your expectations of relevance are satisfied.

To see how this works, consider the lexical ambiguity in Laura's utterance in (6):

(6) Dave: I can't take you to Starbucks, I haven't got any cash on me.

Laura: Well, you'd better find a bank!

To identify the proposition expressed by Laura's utterance in (6), Dave must disambiguate *bank*. He should test possible interpretive hypotheses in order of their accessibility. As Dave has already spoken about money, assumptions about where money is obtained are already slightly activated. Following a path of least effort, Dave will entertain hypotheses about financial institutions. To entertain hypotheses about riverbanks would put him to processing effort that would not be offset by appropriate cognitive effects. Having disambiguated *bank*, Dave ceases to search for further interpretations of *bank* because his expectation of relevance has been satisfied. To keep processing would cause Dave to expend processing effort that would not be offset by cognitive effects in the context.

There is no linear order to this process. Hearers do not identify the propositional content of an utterance and then go on to derive contextual implications from this content. Interpretive hypotheses are entertained and mutually adjusted in parallel (Carston, 2002, p 143). In the case of (6), while Dave is disambiguating *bank*, he will also be entertaining hypotheses about implicit content, namely that Laura implicates that he should find a financial bank to obtain money to buy her coffee. As Carston (2002, p 143) points out, the process of interpretation may involve several 'backwards and forwards adjustments of content before an equilibrium is achieved which meets the system's current expectation of relevance'.

2.4.3 Some Gricean insights and their relationship to Relevance Theory

I would like to now return to Grice, which is where Sperber and Wilson, and, indeed, this chapter began. As Sperber and Wilson (1995, p 25) have said, Grice's contribution to pragmatics was not to suggest that human communication involves the recognition of intentions; it was to suggest that the recognition of communicative intentions is *enough*. The

very act of communicating creates expectations which it then exploits, and this allows for the possibility that people can communicate with each other even when there is no code – as, for example, in the case where I theatrically open the fridge to *show* there is no food. Grice developed this idea by proposing his Cooperative Principle and its attendant maxims of conversation – maxims which Grice claimed speakers are assumed to follow or, in some cases, flout for specific sorts of effect (see Levinson 1983 for a fuller treatment of this). For example, in the exchange in (7), Dave assumes that Chris has followed the maxim of relation and this assumption will lead him to use contextual information in the derivation of certain assumptions which have not been communicated explicitly:

(7) Dave: Are you coming to the pub?

Chris: I haven't got a babysitter

In (8), however, Grice would say that Dave has deliberately flouted the maxim of quantityin order to communicate that his coffee is not particularly good.

(8) Chris: How is your coffee?

Dave: It's hot and wet

While Sperber and Wilson recognised the importance of Grice's contribution, they have shown us that Grice's approach cannot be maintained as it stands (Wilson & Sperber, 1981). First, they have shown that while Grice recognises the importance of contextual assumptions in his 'working out procedures', there is no account of how these assumptions are chosen or, indeed, used. Second, they have shown that it is not clear that all the maxims are needed, and that, in particular, assumptions can generally be recovered on the basis of just the maxim of relation ('Be relevant') (see Wilson & Sperber, 1981 for a discussion of Grice's theory of conversation).

Grice (1989) also recognised the potentially centrality of the maxim of relation. Sperber and Wilson might be said to have followed Grice's instincts. However, what they ended up with is a principle which is very different in nature from the maxims which Grice proposed. Grice's maxims are known by hearers and are either followed or not followed. In fact, as we have just seen, rules even have to be flouted deliberately. In contrast, the Communicative Principle of Relevance is a generalisation about ostensive communicative behaviour which itself is

grounded in a generalisation about human cognition. As Sperber and Wilson (1995, p 162) say, '[c]ommunicators and audiences need no more know the principle of relevance to communicate than they need to know the principles of genetics to reproduce'. The Communicative Principle of Relevance is not a principle which is 'followed' and there is no sense in which hearers could choose not to follow it, either. What hearers use in interpreting an utterance or any act of ostensively communicative behaviour is not the Communication Principle of Relevance, but the fact that a presumption of optimal relevance has been communicated.

2.5 The semantics-pragmatics distinction

Grice (1989) did not actually use the term *pragmatics*. However, he did draw a distinction – the distinction between *what is said* and *what is implicated* - which has played a very influential role in the way people have defined the role of pragmatics and its relationship with semantics. Here, I outline Grice's distinction and show that it cannot be the basis for a distinction between semantics and pragmatics within the framework of the inferential approach to communication I have sketched. In particular, I show that Grice did not afford pragmatics enough of a role in the identification of the propositional content of utterances. The role of pragmatic inference is not restricted to the identification of *what is implicated*, but also plays a role in the recovery of *what is said*, or the proposition expressed. The upshot of this is that we cannot maintain the correlation, implicit in Grice's distinction, between propositional content and linguistically-encoded meaning, and we cannot adopt the traditional truth-conditional view of the semantics-pragmatics distinction, namely that semantics is the study of truth conditions and pragmatics is the study of everything else (see Gazdar, 1979). Finally, I outline the relevance-theoretic view of semantics as providing an input to pragmatic inference, showing that this allows us to develop a non-unitary account of linguistic semantics which explains many of the cases of linguistic meaning which are problematic for the truth-conditional approach to semantics.

2.5.1 Grice: what is said and what is implicated

The question of how the distinction between explicit and implicit content should be drawn has been the subject of considerable controversy (see Recanati, 1989; Bach, 1994; Carston, 1998, 2002). My aim here is limited to showing that the distinction drawn by Grice between

what is said by a speaker in making an utterance and what the speaker of that utterance *implicates* does not correspond to the relevance-theoretic distinction between explicit and implicit content. By *what is said*, Grice (1989) means the conventional meanings of words that deliver the proposition expressed by an utterance. By *implicated*, Grice (1989) means the aspects of utterance meaning that do not seem to be concerned with the recovery of truth-conditional content. Example (9) illustrates Grice's distinction:

(9) It's freezing in this room.

What is said in (9) is the proposition [it is freezing in this room]. According to Grice (1989), the speaker of (9) *says* that it is freezing in this room. However, it is not difficult to imagine that in some contexts he also communicates something like the proposition [I want you to shut the window]. This is what Grice would call the conversational implicature derived from the utterance. The proposition expressed by (9) is decoded directly from the conventional meaning of the words making up the utterance. In contrast, Grice would say (1989), the implicated proposition [I want you to shut the window] is not decoded from the meanings of the words making up the utterance, but is derived inferentially on the basis of the assumption that the communicator speaks in accordance with the maxims of conversation (in this case, perhaps, the maxim of information). For example, the hearer might be expected to reason along the following lines: the speaker is cooperative, and following the maxims of conversation, but both the speaker and I know that it is freezing in this room and the speaker knows that I know this, so the speaker can only be construed as being informative provided he is communicating something in addition to the proposition that it is freezing. In this context, the most likely implicature which would allow me to preserve the assumption that the speaker is being informative is the proposition that he wants me to close the window. This sort of account is not without problems. However, my concern here is with Grice's assumption that the proposition expressed (*what is said*) is derived almost entirely through linguistic decoding and does not sufficiently involve pragmatic inference.

Sperber & Wilson (1995) and (particularly) Carston (2002) have shown that Grice seriously underestimated the role that pragmatics plays in recovering the proposition expressed. Pragmatically constrained inference *is* involved in the recovery of many aspects of propositional content. Consider the following:

(10a) He ate the cake.

(10b) George sat by the bank.

(10c) Tired.

(10d) Julia likes a pint.

(10e) My bedroom is rectangular.

All the utterances above require substantial pragmatic work for the recovery of the proposition which the hearer takes the speaker to have intended. For example, in (10a), we need to infer whom *he* refers to. In (10b), we need to disambiguate *bank* in order to understand if George sat by a riverbank or a financial institution. Grice (1989) admits that speakers would have to have a means of assigning reference and a means of dealing with ambiguity in order to determine the proposition expressed. However, this was as far as Grice was prepared to go in allowing inference to determine the *what is said* of an utterance.

As early as 1981, Wilson and Sperber demonstrated that pragmatically constrained inference must play a role in the recovery of what Grice had called *what is said*. Their examples include the following answer to the questions in (11b) and (11c):

(11a) I refuse to admit them.

(11b) What do you do with gate crashers?

(11c) What do you do when you make mistakes?

(Adapted from Wilson & Sperber, 1981, pp. 156-158.)

The only way we can explain that *admit* is interpreted as *let in* (rather than *confess to*) when (11a) is intended as an answer to (11b) is to assume that the speaker was giving a relevant answer to the question – or, in Grice's terms, that the speaker was following the maxim of relation.

As Sperber and Wilson (1995, p 193) say, explicit content is much more inferential and much more worthy of pragmatic investigation than Grice envisaged. I now outline the inferential approach to explicit content developed in Relevance Theory by Sperber and Wilson (1995), and by Carston (2002), and show how they draw the implicit/explicit distinction.

2.5.2 Explicit content and implicit content

There are two types of assumption that a speaker can communicate: an *explicature* or an *implicature*. Consider (12):

(12) You've baked a cake!

Sperber and Wilson (1995, p 182) say that an assumption communicated by an utterance is explicit if and only if that assumption is a development of a logical form decoded from the linguistic semantics of that utterance. 'An explicature is a combination of linguistically encoded and contextually inferred conceptual features' (Sperber & Wilson, 1995, p 182). On this view, the linguistic semantics of utterances encodes a propositional schema or blueprint, which is inferentially developed to something fully propositional. One of the explicatures communicated by (12) is [X has baked a cake] (X is a representation of the referent of *you*). However, this explicature might be embedded under propositional attitude descriptions to yield what Sperber and Wilson call *higher-level explicatures*, for example [the speaker is happy that X has baked a cake].

The recovery of an explicature clearly involves inference. However, crucially, it also involves decoding. The hearer is intended to decode the linguistic meaning of (12) and develop it by pragmatic inference into the proposition which she takes the hearer to intend. In contrast, the recovery of an implicature involves *only* pragmatic inference. Sperber and Wilson (1995, p 182) say that:

> 'Any assumption communicated, but not explicitly so, is implicitly communicated: it is an *implicature*. By this definition, ostensive stimuli which do not encode logical forms will, of course, only have implicatures'.

We might imagine that (12), in the right context and uttered with a teasing tone of voice, yields the implicature [I would like you to give me some cake]. This is an assumption that can be communicated by (12), but is not a development of the logical form decoded from the linguistic semantics.

Explicitness is a matter of degree. Speakers can be more or less explicit. Sperber and Wilson (1995, p 182) say that an utterance is relatively more explicit if the contribution of contextual

information to the explicature is relatively small. Each of the below utterances is increasingly more explicit, as each utterance requires less contextual supplementation:

(13a) Cake!

(13b) A cake!

(13c) You've baked a cake!

(13d) I see you've baked a cake!

(13e) I am happy to see that you have baked a cake.

Now it is clear why the explicit-implicit distinction is not the same as Grice's: 'On a more traditional view, the explicit content of an utterance is a set of decoded assumptions, and the implicit content a set of inferred assumptions. Since we are claiming that no assumption is simply decoded, and that the recovery of any assumption requires an element of inference, we deny that the distinction between the explicit and the implicit can be drawn in this way' (Sperber & Wilson, 1995, p 182). Relevance Theory adds in an extra layer of pragmatic 'work' between semantic decoding and full propositions. This layer of work is needed because of the gap between what words decode to, and the propositions that they express. Carston (2002) has shown that this gap is substantial, and it is pragmatics that closes it.

2.5.3 Lexical pragmatics and linguistic underdeterminacy

(Monomorphemic) conceptual expressions have been treated as encoding atomic concepts (Carston, 2002, p 321). Relevance Theory follows the Fodorian view on concepts, generally (Sperber & Wilson, 1995; Carston, 2002). Carston (2002, p 321) says '...an atomic concept consists of an address or node in memory which may make available three kinds of information: logical content, encyclopaedic or general knowledge, and lexical properties'. To see what this looks like, I borrow Carston's (2002, p 321) outline of the atomic concept CAT (words in capitals are concepts):

> 'Consider the concept CAT: its logical entry contains an inference rule whose output is ANIMAL OF A CERTAIN KIND; its encyclopaedic entry contains general knowledge about the appearance and behaviour of cats, including, perhaps, visual images of cats and, for some people, scientific knowledge about cats, such as their anatomy, their genetic make-up, or their relation to other feline species, etc., and, for most people,

> personal experiences of, and attitudes towards, particular cats; its lexical entry, for an English-speaker, includes the phonetic structure and grammatical properties of the word 'cat'' (Carston, 2002, pp 321-322).

In light of the above, consider *kind* in (14) and, then, in (15).

> (14) I found the Japanese to be particularly **kind**.
>
> (15) I think you were being too **kind**. If they had said that to me, I'd have felt insulted.

Intuitively, the same concept of KIND is not communicated in both cases. In (15), what the speaker has in mind is a concept which might also be communicated by the *tolerant* or *indulgent*. However, the concept I had in mind in (14) would not have been communicated by these words. It is a concept which captures the feeling I had towards the acts of unsolicited and unexpected gift-giving I experienced in Japan – acts which fall outside any experience of kindness in Britain. I had in mind a concept which might be closer to the concept communicated by the word *generous*. The point is that in each case, the speaker communicates a different concept – a concept which is derived via pragmatically constrained inference from whatever is linguistically encoded by that word. Carston (2002) calls such pragmatically derived concepts *ad hoc concepts* and uses the following notation to distinguish different ad hoc concepts communicated on different occasions: KIND*, KIND**, KIND***, etc.

It might then be best to think of the concept encoded by *kind* as a general concept schema that is interpreted in specific ways to meet the expectations of relevance raised by particular utterances in particular contexts. Carston (2002) has suggested that words such as *kind* or *happy* might not encode concepts at all but rather 'point' to conceptual regions. 'This pointing or mapping', she says, 'provides access to certain bundles of information from which the relevance-constrained processes of pragmatic inference extract or construct the conceptual unit which features in the speaker's thought' (Carston, 2002, p 361). As Sperber and Wilson (1998) have argued, this suggests that a concept cannot be seen an internalisation of the word used to communicate it. They argue that it is possible to use a word to communicate a concept

which one has not previously encountered and which would never be derived as a constituent of a thought outside of that particular situation. Consider (16):

(16) Julia: Do you want to go to the cinema?
Roz: I'm tired.

As Sperber and Wilson (1998) show, assuming that Roz has aimed at optimal relevance, Julia pragmatically develops the encoded meaning of *tired* to the point that it allows her to infer that Roz does not want to go to the cinema. In this way, the concept that Roz intends to communicate must be understood as an ad hoc concept of tiredness that is linked to the particular circumstances of the utterance – a narrowed concept of tiredness which warrants the derivation of the conclusion that Roz does not want to go. Concepts can also be broader than one might consider a lexicalised atomic concept to be – 'bachelor' can be legitimately applied to men who are married in order to say something about their behaving as if they are still single. The fact that there is no specific lexical item for this context-specific concept of tiredness, or for bachelor, does not matter: it can still be communicated. What this discussion shows, however, is that even individual conceptual expressions underdetermine the concepts they express, and that modulation processes must be at work in the computation of word meaning.

2.5.4 Modulation processes

There are a number of pragmatic enrichment processes that speakers undertake to modulate, or adjust, conceptual representations. These processes are 'free' in that they are not linguistically mandated, but step in if it is optimally relevant to have them do so (Carston, 2009). For example, approximation occurs when a word is used to communicate an approximate concept, i.e., *square* communicates something like square-ish, metaphorical extension is present when we apply a term to something it wouldn't normally apply to, i.e., *battleaxe* for a frightening person (Wilson, 2004, p 343). An important lexical-pragmatic process for this thesis is narrowing, a type of adjustment whereby an ad hoc concept is computed that is specific and finely-tuned (narrower), i.e., using *drink* to communicate ALCOHOLIC DRINK* (Wilson, 2004, p 343). 'The effect of narrowing is to highlight a proper subpart of the linguistically-specified denotation' (Wilson, 2004, p 344). Narrowing is driven by considerations of optimal relevance when following the relevance-theoretic

comprehension procedure. Hearers narrow in order to satisfy, and do so *until* they satisfy, their expectation of relevance. A link between narrowing and intensification is explored in chapter five of this thesis. Broadening occurs when a word is applied in a more general sense, and the remit of its denotation is extended (Wilson, 2004, p 344).

It is clear that the hearer of virtually any utterance can shoulder a great deal of the responsibility for the recovery of the conceptual content derived. This has implications for the relationship between language and thought, which, in turn, has implications for our understanding of repetition. For now, I would like to note that, in Relevance Theory, it is held that the words we use massively underdetermine the propositions we intend to express by them. This is known as the radical underdeterminacy thesis (Carston, 2002). The range of repetition data in chapter five in particular are particularly indicative of how much pragmatic work we have to do to get to the (intended) proposition expressed.

I now ask: how should we define semantics in a fully inferential approach to communication?

2.6 Non-unified semantics: enter procedural encoding

The picture of semantics which emerges from the considerations of linguistic underdeterminacy above suggests an approach to semantics is needed which is distinguished from approaches that assume that semantics delivers truth-evaluable propositional representations (Burge 1974; Higginbotham 1985, 1988). Instead, the linguistic semantics delivers schematic logical forms which provide an input to the pragmatically constrained inferential computations involved in utterance interpretation (Blakemore, 2002; Carston, 2002). Given this approach, the question for linguistic semantics is not what contribution an expression makes to truth conditions, but rather what kind of contribution it makes to pragmatic inference, or, in other words, what kind of cognitive information it encodes. Moreover, as Blakemore (2002) has shown, from a relevance-theoretic perspective, we can ask whether there is just one answer to this question. Is semantics necessarily unitary?

The answer suggested by Relevance Theory must be *no*. Relevance is determined not just by the rewards derived from processing information, but also by considerations of cognitive efficiency. It would be cognitively efficient if languages developed expressions which encode not just conceptual information, but also instructions (procedures) for deriving those

representations. This would optimise the chances of the 'right' conceptual representations being recovered, and reduce the chance that processing effort is wasted. Our utterances contain within them expressions which constrain us in how we process them. For me, it is a bit like how a road network has signposts to help people to find destinations more successfully. You might get to where you want to go without the signs, but surely your chances improve if the 'system' has some guidance built into it. An analogy between road signs and procedural expressions is not perfect, but it helps us to visualise why we ought to expect procedural expressions in the linguistic repertoire.

More generally, Wilson (2011) has argued within the modular approach to the human mind which is adopted in Relevance Theory that we might expect languages to develop expressions which activate procedures in specific cognitive domains, e.g., domains of inferential communication, or domains concerning emotions. Considered from this point of view, the expressions analysed by Blakemore (1987, 2002) (e.g., *but*, *well*, *so*) can be treated as means for activating procedures in the domain of inferential communication. In contrast, the expressions and behaviours analysed by Wharton (2003, 2009, 2015) (e.g, *wow*, *ow*, or genuine smiles) can be viewed as means for activating procedures within the domain of emotion reading. Here, I outline the notion of procedural meaning as it was first proposed by Blakemore (1987, 2002), and show how it has been developed to account for a wider range of phenomena.

Semantic Constraints on Relevance (Blakemore, 1987) was an attempt to re-analyse the phenomena treated by Grice (1989) as yielding conventional implicatures. Words such as *but* or *therefore* presented particular problems for Grice because they suggested that he could not maintain his definition of *what is said* in terms of a correlation between linguistic content and truth-conditional content. Such expressions are (notorious) counterexamples to truth-conditional semantics: they don't affect the truth conditions of the sentences containing them. At the same time, however, their meanings were clearly conventional or linguistically encoded, as they seemed to have systematic effects on utterance interpretation. Because such expressions did not contribute to truth conditions, Grice decided to treat them as implicating, or, more particularly, as conventionally implicating. On these grounds, other authors (e.g., Gazdar 1979; Karttunen & Peters, 1979) treated them as pragmatic phenomena rather than semantic. Blakemore's (1987) proposal was that these expressions instead

semantically encode constraints on implicit content. At the time, Blakemore envisaged that these expressions might correspond to non-truth conditional aspects of utterance meaning, for example *so* and *after all* as they are used in the following sequences:

(17a) Seb can share out the cakes. So he can count.

(17b) Seb can share out the cakes. After all, he can count.

The second segments in both (17a) and (17b) express the same proposition [Seb can count]. While this proposition is interpreted as a conclusion derived from the first segment in (17a), in (17b), it is understood as a premise in an argument which has the first segment as a conclusion. Blakemore's (1987, 2002) idea was that *so* and *after all* activate different inferential processes – processes which result in different cognitive effects in each case. These expressions ensure that the hearer recovers the intended interpretation for a minimum cost in processing. However, subsequent research at the semantics-pragmatics interface has established that procedural encoding is also involved in the computation of explicit content. It has been argued that pronouns encode procedures for computing referents by ensuring certain contextual information is recovered and developed (Wilson, 2011; Scott, 2015). According to this view, the semantics of *she* can be analysed as a procedure that encourages the hearer to identify a singular, (usually) female entity in the development of a conceptual representation of a referent. Similarly, it has been argued that some mood indicators encode procedures for the construction of higher-level explicatures (see Wilson & Sperber 1988; Clark 1991). If this is right, the two types of coding do not correspond to truth-conditional and non truth-conditional meaning. More recently, the notion of procedural meaning has been extended further still to include devices which are used to activate procedures that are used in the identification of emotions or emotional states. Some of these devices are only borderline linguistic, e.g., *ow*, or *wow*, or are completely non-linguistic (see Wharton, 2009 on interjections and *shown* natural behaviours). The point is, however, as Wilson (2011) has argued, any procedure from any cognitive domain may be activated by such devices, whether they involve linguistic encoding or not. For example, as I explain in §2.8.4 of this chapter, many phenomena which are involved in the communication of expressive meaning are now afforded a procedural interpretation in Relevance Theory.

One of the chief characteristics for identifying procedural meaning is semantic ineffability (Blakemore, 2011). While we might attempt to model the semantics of a conceptual expression using lists of assumptions, or descriptions of conceptual information, we cannot do this for expressions like *but* or *thus*. Concepts are the sort of thing we can visualise in our mind's eye. This is not the case for procedural expressions. If I ask you to tell me what *cat* or *orange* means, you can provide definitions for these expressions, and even point to exemplars. This is not possible for procedural expressions. It is difficult to pin down the exact meaning of so-called discourse markers like *well* and *nevertheless*, or expressives such as expletive *bugger*, or *ouch*. By the same token, it is difficult to pin down the 'meaning' of certain prosodic features associated with particular tones-of-voice, interjections such as *ah*, or diminutives such as *the wee softy* (see Potts, 2007; Blakemore, 2011). Even native speakers seem unable to say what such expressions mean in conceptual terms. When asked for the meaning of a word like *bastard* or *but*, a speaker is more likely to give an example of its use (Blakemore, 2002).

Unsurprisingly, such expressions are notoriously difficult to translate (see Gallai, 2013). While language 1 might have one expression that encodes one particular procedure, language 2 might have two separate procedural expressions that encode the same procedure, or closely related procedures, and vice-versa. This translation issue does not really arise for conceptual expressions. Consider the German utterances in (18a), (18b), and (18c).

(18a) **Die Katzen** sind sehr **freundlich**.

'The cats are very friendly.'

(18b) Ich bin muede **aber** ich will nicht ins Bett gehen.

'I'm tired **but** I don't want to go to bed.'

(18c) Ich bin muede **doch** ich will nicht ins Bett gehen.

'I'm tired **but** I don't want to go to bed.'

Katzen, barring any overtly figurative cases, will usually translate to *cats*. Similarly, *freundlich* will usually translate as *friendly*. This is not the case for procedural expressions. Students of German learn that *but* translates as *aber*. However, depending on the context, this may not be appropriate. Sometimes, *but* translates as *aber*, and sometimes it translates as *doch*, as in (18b) and (18c). Which procedural expression a translator chooses depends on him taking

into account the speaker's intentions, and the context for utterance interpretation. German has a range of expressions (*doch*, *aber*, *jedoch*) that converge on something like the semantics of *but*. Translators must consider exactly which German procedural expression will lead to the recovery of the intended implicature(s).

Non-native speakers of languages find it very hard to acquire the meanings of procedural expressions. The definite article and the indefinite article in English are prime candidates for having a procedural semantics. The utterances in (19a-c) are, produced with permission, from a Polish man to whom I taught English.

(19a) *I am mechanical engineer.

(19b) ? I am outgoing because I go to the theatre group. (This was an answer given in a mock job interview; there was no mutually manifest theatre group so I found his utterance unacceptable.)

(19c) *? Can we please rearrange lesson? I have to go to a doctor's.

Clearly, the above examples relate to difficulty in acquiring the procedural expressions concerned with definiteness. *A/an* and *the* are notoriously hard to acquire. Successful acquisition usually requires exposure to the expressions in very obvious contexts.

Why is it that the meanings of procedural expressions are so ineffable? If you ask me how I can see, I might be able to tell you something about the eye, and perhaps some scientific knowledge about vision. I will not, however, be able to describe the mental processes involved in vision. The experience of seeing is a person-level process. However, the cognitive systems underpinning vision are sub-personal. I have no conscious access to them. The same can be said of procedural meaning. I can provide contextualised examples of procedural expressions but I cannot describe the procedures themselves. For this reason, if we encounter a linguistic form whose semantics is resistant to introspection, and seems to be hard to translate, or to acquire, it is likely procedural. These 'tests' are revisited again in chapter five, §5.4.4.

2.7 Communication and Style

2.7.1 The relationship between utterances and thoughts

Utterances and thoughts are very different sorts of things. Thoughts are private and even when we talk about 'getting thoughts down on paper' or 'getting thoughts across', in reality, these thoughts remain where they are – in our heads. What we put down on paper are squiggles and lines which are somehow used as evidence for these private thoughts. Utterances are events which take place in particular places and at particular times. And with the exception of those produced when talking to oneself, they are produced typically for the benefit of others. Utterances can be thought of as public events which provide evidence for private thoughts. However, as we have seen in this chapter, the path from utterances to representations of thoughts is necessarily indirect. It involves inference. Given the role which pragmatic inference plays in the interpretation of utterances, how can we be sure that the representation a hearer recovers is a representation of the thought the speaker intended to communicate? This is actually the wrong question to ask – if by *representation of the thought communicated* we mean *thought which is identical to the one communicated by the speaker*. Sperber and Wilson (1995, pp. 226-237) have argued that the relationship between the thought recovered by the hearer and the one the speaker intended to communicate is not one of *identity*, but of *resemblance*, where resemblance is determined by the extent to which two mental representations give rise to the same logical and contextual implications (see also chapter one, §1.2). In other words, every utterance is an interpretation of a private thought of the speaker.

How do we *know* whether two mental representations give rise to the same logical and contextual implications? This would entail looking inside the contents of each other's minds – and this is impossible. Indeed, if it were possible, we would not need communication. Sperber and Wilson (1995) would say that the only thing that a hearer should do is go ahead and recover the interpretation that satisfies her expectations of relevance. If a speaker has produced an optimally relevant utterance, then the interpretation the hearer recovers will be sufficiently like the thoughts which the speaker wanted to communicate. In other words, the picture of communication in Relevance Theory is not one in which communicative success depends on the duplication of thoughts, but is one in which communication results in what

Sperber and Wilson (1995, p 193) describe as the enlargement of mutual cognitive environments. On this view, an utterance is simply (public) evidence for a (private) thought, and the interpretation recovered by a hearer can only be an interpretation of the thought communicated. Communication will succeed to the extent that the optimally relevant interpretation achieves a sort of 'loose' coordination which, as Sperber and Wilson (1998, p 199) say, is best compared to the coordination between people taking a stroll rather than that between people marching in exact step. For example, in communicating the thought that Chris is worried, the speaker in (20) below can only assume that the audience's search for relevance will yield a concept WORRIED* which resembles his sufficiently for it to play a role in the (loose) co-ordination of their behaviour.

(20) Chris is worried.

2.7.2 Strength of communication

Utterances contain constituents because they are real-world representations made up of linguistic constituents. Thoughts also contain constituents. Thoughts are mental representations about things in the world and, as such, contain constituents that correspond to the things they are about. In other words, thoughts about theses contain constituents corresponding to theses, and thoughts about your supervisor contain a mental constituent corresponding to your supervisor. We might be tempted to think, then, that if a constituent is found in an utterance, there should be a corresponding constituent found in any representations recovered by interpreting that utterance. This is not necessarily the case. Consider the cases of epizeuxis in (21a-21c):

(21a) I am never, never going to do another PhD.

(21b) There are toys, toys, toys all over the floor.

(21c) My career is over, over.

As we shall see in chapter four of this thesis, the effects of epizeuxis can be recovered in a number of ways. In (21a), the repetition achieves relevance by scaffolding the recovery of the communicator's degree of commitment to the proposition expressed. In (21b), it achieves relevance by suggesting there are more toys on the floor than one might expect. In (21c), the communicator is not suggesting that his career is more over than one might have thought or

that he is more strongly committed to the proposition that his career is over than one might have thought. I will not set out here what the speaker in (21c) communicates. The point is that while in each case the repetition *suggests a direction* for interpretation, the audience does not have any particular information about the interpretation she should end up with. In this way, the responsibility for the interpretation is given to the hearer. And while the interpretation the hearer ends up with is scaffolded in some way by the repetition, we would not say that the hearer ends up with a thought which itself contains a repetition or, rather, repeated constituents; the repetition is simply a way of setting the hearer off in a particular interpretive direction, as we shall see throughout this study. Sperber and Wilson (1995) have argued that the resulting interpretation for a case like (21c) consists of a very wide range of weakly communicated implicatures which the audience will assume provides a faithful interpretation of the communicator's feelings; in other words, what the hearer recovers in this sort of case is 'a sense of apparently affective rather than cognitive mutuality' (Sperber & Wilson, 1995, p 224). Communication is strong, or weak to varying degrees and often a mixture of the two:

'When the communicator makes strongly manifest her informative intention to make some particular assumption strongly manifest, then that assumption is strongly communicated', say Sperber and Wilson (1995, p 59). They (1995, p 60) continue '[i]n the case of strong communication, the communicator can have fairly precise expectations about some of the thoughts that the audience will actually entertain'. In cases of strong communication, the speaker provides very clear ostensive evidence for an intention to make a particular assumption extremely manifest, and, in doing so, the speaker can be relatively confident of steering a hearer's thoughts in a particular direction. With some utterances, bearing in mind the context, there is often only one way that they can be taken. Consider the example of the indirect answer in (22):

(22) Dave: Are you coming to the pub?

Chris: I haven't got a babysitter.

In this example, Dave derives implicated premises and conclusions which are consistent with the assumption that Chris has aimed at optimal relevance. In this way, Dave is directed

towards the argument in (23a) which will enable him to derive the implicated conclusion in (23b):

(23a) If Chris hasn't got a babysitter, he won't be going to the pub.

(23b) Chris won't be going to the pub.

There is very little else that Dave can do here. In this sense, the assumption in (23b) is strongly communicated. At the same time, however, one might ask why a speaker who is aiming at optimal relevance should produce an utterance which requires processing effort which is not required by the very direct answer in (24).

(24) Chris: No.

Chris' answer in (22) could be interpreted in any way at all by Dave. However, any assumptions he derives would not be strongly implicated by Chris, because they could not be said to have been strongly intended by him. The extra processing entailed by Chris' utterance in (22) could be said to be offset by cognitive effects which might have been intended by Chris to a lesser degree – for example, the assumption that Chris is feeling sorry for himself, that Chris wishes he could go to the pub, that Chris would like Dave to come up with a solution to his baby-sitting problems, and so on. Since there is no specific cognitive effect which can be said to be part of Chris's communicative intention, we must say that the cognitive effects which justify the extra processing effort entailed by the indirect answer are *weakly communicated*. The hearer must take responsibility for deriving them in line with their own expectation of relevance in the context.

Here, I explained that assumptions can be strongly or weakly communicated. That is not to say, though, that everything that we communicate is propositional. Consider (25):

(25) [David's friend has borrowed his computer and has deleted some very important files.]

David (to ex-friend): You're a stupid, stupid, STUpid wanker and I hate you!

Although the reproduced *stupid*s in (25) modify *wanke*r at the level of explicit content, it is likely that they also play a role in causing the hearer to experience a negative emotion. In (25), David wants the hearer to feel bad, and produces his utterance in such a way as to bring this

about. Many emotional states are not propositional in nature. The one in (25) might be said to be *about* something, but not everything emotional is. Emotions are addressed in this chapter, §2.8.3. However, many of the repetitions I examine in this thesis result in the communication of weakly communicated non-propositional effects, so I discuss these in the next section.

2.7.3 Propositional and non-propositional thoughts

Propositions are representations that are 'about' things in the world, are semantically complete, and have truth conditions.

(26) [It is raining in Salford]

(26) is a proposition which we might derive either as an explicature or an implicature. It represents a state of affairs in the world. This proposition can be a premise in an inference for deriving other propositions, e.g., [it is raining in Salford] entails [it is raining somewhere]. However, there are representations which we are capable of communicating which do not seem to be anything like (26).

(27a) [Twitter, 02/2012]

User: Fuck. **Fuck fuck fuck**. Petty problem of the day? Accidentally turned stove on and destroyed coffee pot.

(27b) [Text from a friend in relation to the invitation to go out drinking, 05/2012]

Friend: :) you are going to kill me **:) :) :)**

(27c) [Old woman talking about her life, 04/2012]

Woman: So '71 wasn't really that great. **That's life...that's life**.

(27d) [Sylvia Plath, 'Daddy', in 'Ariel', 1965]

I never could talk to you,

My tongue stuck in my jaw.

It stuck in a barb-wire snare.

Ich, ich, ich, ich,

I could hardly speak.

We are interested in the bold elements of (27a-27d). What proposition is communicated by *fuck fuck fuck*? Perhaps the hearer of *fuck fuck fuck* recovers some propositional representations on the basis of it, e.g., [the writer is angry] or [the writer swears lots]. However, it seems that it is more likely that what the writer is actually doing is providing evidence for his emotional state, and the hearer uses this evidence to calibrate a representation of it. What this representation is like, however, is not clear. But it is probably not like the type of representation we have in (26). The facial emoticons in (27b) could also be analysed as leading to the recovery of emotional states (this is the subject of chapter 5, §5.6). Note that in (26), the representation was complete. That is to say, the representation was *determinate*. This is not the case for (27b), and probably for the other examples, too. We cannot delimit the range of effects that the hearer will recover. In (27c), it is true that the words *that's life* can express a proposition. However, it seems that the old woman wants to communicate a sentiment of some sort. Sentiments are not propositions. Finally, in (27d), the reproduced *ichs*, do not communicate a proposition, either. Their effects likely arise through resemblance between the palatal fricative of German, and actual barbed wire. These effects are all impossible to paraphrase conceptually.

The fact that non-propositional effects are hard to paraphrase is a diagnostic that we can use to identify them. If we take an utterance that seems to have non-propositional effects associated with it, and compare it to an altered version that does not yield these effects, we get a sense of something 'lost in translation', and can try to identify any linguistic form that is responsible for the communication of the non-propositional effects (without encoding them).

(28a) I shall never, never, never drink Zubrowka again.

(28b) I shall never drink Zubrowka again.

The hearer of (28a) recovers a range of non-propositional effects relating to the speaker's emotional state. If we remove the repeated linguistic material to get (28b), these non-propositional effects are lost.

Sperber and Wilson (1995) deal with the communication of non-propositional effects through the notion of weak implicature. They (1995, p 222-224) say '[l]et us give the name poetic effects to the peculiar effect of an utterance which achieves most of its relevance through a wide array of weak implicatures', and that '[poetic effects] do not add entirely new

assumptions which are strongly manifest in [the cognitive] environment. Instead, they marginally increase the manifestness of a great many weakly manifest assumptions. In other words, poetic effects create common impressions rather than common knowledge'. On this view, a speaker makes weakly manifest an intention to make ever so slightly more manifest a wide range of implicatures. This leads to the recovery of interpretations of the kind associated with the examples in (27a-27d) and (28a). For example, in the case of (28a), the speaker makes weakly manifest an intention to make slightly more manifest a wide range of assumptions such as the ones below in (29) – although the hearer need not recover all of these specific assumptions for communication to work.

(29) [She feels sick]

[She is hungover]

[Zubrowka is very strong]

[She can't face the thought of drinking ever again]

[She doesn't want an alcoholic drink right now]

[She wants to go to bed]

[She is in no fit state to work on her thesis]

This account seems to work for examples like (28a). A similar account might be applied to sentimental cases, as Pilkington (2000) has suggested. He treats sentiments as achieving their effects by the hearer accessing stereotyped bundles of weakly communicated assumptions that lead a hearer to feel sentimental (Pilkington, 2000). It should be clear, however, that any information recovered during utterance interpretation can have an impact on a hearer's emotional state, and, also, their own physiology. As I will show, there are repetitions that have an extremely emotional flavour, and while some of them may also make slightly more manifest a wide array of weak implicatures, an important intention behind them is not to do this, but to scaffold the calibration of emotional states (e.g., some repeated face emoji). The effects of these repetitions will be non-propositional, but in a different sense to how Sperber and Wilson (1995) explained non-propositional effects originally, i.e., they are non-propositional but not poetic effects. I address repetition cases that cannot be explained in terms of weak implicature in chapter five. How I view emotions is set out in §2.8.3 of this chapter.

2.7.4 Relevance theory and style

A simple Internet search for 'style' brings up texts that deal with topics such as literary composition, rhetoric, poetics, irony, satire, comedy, register, prose, and literary language, among others. These (admittedly unscientific) search results lead me to two conclusions. Firstly, whatever 'style' is thought to be, a theory of it should be able to explain the type of phenomena that people agree should fall under its umbrella. Secondly, modern stylistics appears to still be heavily concerned with the literary in some respects, focusing on poetic notions like metaphor or alliteration, and other aesthetically pleasing aspects of language. I suggest that if phenomena such as metaphor or repetition are to fall under the umbrella of style, you have to take a view of style that is embedded within an appropriate view of communication - a view of communication that reflects that we do much more than communicate propositions. Moreover, since metaphor, repetition, and a host of other 'devices' are actually commonplace in everyday ostensive-inferential communication, and engender the type of effects that literary theorists are concerned with in 'normal' contexts, the study of (the effects of) genuine repetition data brings the study of style down from the clouds and into the realm of the everyday, as I shortly discuss.

2.7.5 Against a propositional view of communication and style

Another reason why communication using language is not concerned with the duplication of thoughts is that, if it were, we would not be able to explain stylistic phenomena that don't encode what they communicate, and/or don't encode propositional representations. In order to explain such phenomena, a view of communication is needed that accepts that we do communicate non-propositional effects, and this can be done without certain expressions or stylistic 'devices' *linguistically encoding* these effects as pre-packaged, guaranteed interpretations. To make this clear, I use examples of metaphor, syntactic parallelism, and repetition of linguistic material.

> (30a) [Romeo and Juliet, Act 2, Scene 1, when Juliet appears at a window (Shakespeare, 2000)]
>
> Romeo: But soft, what light through yonder window breaks.
>
> It is the East, and **Juliet is the sun**.

(30b) Diane lives in a gorgeous house, Akis in a lovely apartment, and **Vanessa lives in Salford**.

(30c) [From 'Text', in The Rapture, by Carol Ann Duffy (2006, p 2)]

I tend the mobile now

Like an injured bird.

We **text, text, text**

Our significant words.

What does Romeo intend to communicate when he says Juliet is the sun? That she is a hot ball of gas orbited by planets? It seems that what Romeo wants to communicate is a sense of radiance and wondrous beauty that befits Juliet within the context. Although (30a) has an explicature, this case of metaphor is also associated with a range of non-propositional effects without encoding them. SUN* comes to be a constituent in the explicature of (30a), but the weak effects are not encoded by SUN* - they rather arise from the mutual adjustment of explicit and implicit content. SUN* plays a role in the recovery of weak effects, but does not linguistically encode these. (30b) is a case of syntactic parallelism, defined as the deliberate reproduction of syntactic structure designed to result in stylistic effects. The first two reproduced syntactic forms are associated with parallel semantic forms. This raises the expectation that any further reproduced syntactic forms in the immediate context will likely be associated with propositions about buildings where people live. However, the third section of (30b) is about a place where people live, rather than a building. The hearer strives to seek a context in which all three segments of the utterance can be accommodated, in spite of the semantic difference of the third segment. As Sperber and Wilson (1995, p 223) put it '[t]he problem is then one of finding a context in which all three clauses have parallel contextual effects.' The reproduced syntactic structures in (30b) play a role in the recovery of the non-propositional effects, but without linguistically encoding them. We can see this by looking at (31a) and (31b) below:

(31a) Seb: But I've tidied my room!

Chris: Not properly. Here's a piece of lego and here's a piece of lego.

(31b) Seb: But I've tidied my room!

Chris: No you haven't. Here's a piece of lego, and here's a piece of lego and here's a piece of lego and here's a piece of lego.

In (31a), the reproductions of syntactic form are incidental; Chris needs them to communicate that there are two separate pieces of lego on the floor. No stylistic effects are intended. In (31b), however, the reproduced syntactic structures *are* to be taken as ostensive evidence of Chris's communicative intention to communicate some kind of non-propositional effect. Certainly, Chris wishes Seb to understand that there is a lot of lego on the floor. However, the fact that the same syntactic structure is used over and over is highly noticeable and makes manifest assumptions about Chris's emotional state. Chris uses the reproduced syntactic structures to communicate his emotions. In this case, the reproduced structures are associated with the recovery of non-propositional effects, but these are not encoded by the structure. If they were, the same weak effects would have been recovered in (31a).

In (30c), the reader entertains a wide range of images and memories as a result of reading *we text, text, text our significant words*. These images and memories are not linguistically encoded by the three tokens of *text* as a unit. An analysis of similar cases is presented in chapter four, §4.3; I do not wish to pre-empt the discussion too much. However, I can say each *text* encourages the hearer to re-activate, or activate more strongly her conceptual address for *text* to search for ever more weakly communicated assumptions on her own responsibility. Thus, the initial token does not encode these effects, and neither do the words as a unit - the effect is incremental and cumulative.

2.7.6 A deflationary account of style (and repetition)

Researchers with an interest in stylistics and related fields are interested in particular tropes, devices, or writing and speaking 'styles'. Many of these phenomena are lumped together because they result in some kind of pleasant poetic or stylistic effect. These devices or tropes are communicative, and are deliberately produced to induce these effects in audiences. It seems to me that much work in stylistics focuses on aesthetically pleasing aspects of language, and gives phenomena such as metaphor, parallelism and repetition an elevated status. By focusing on more 'literary' or 'artistic' phenomena, researchers deprive themselves of a wealth of data, and end up needing two accounts of style: one that deals with the pleasant and the appealing, and one that accounts for ordinary aspects of language use.

There are repetitions that are associated with non-propositional effects, but are not really considered to fall under the umbrella of style. The segments in the following utterance pair differ in their linguistic form, and so could differ in terms of their interpretation. (32b) results in the recovery of a wide range of weak effects that are not recovered in (32a). To my knowledge, this phenomenon (the subject of part of the next chapter) has only been treated as a *linguistic* phenomenon - and a case of reduplication at that. No one has addressed it as a matter for style.

(32a) Make me a PROPER pizza.

(32b) Make me a PIZZA-pizza.

However, if we do not treat the repeated linguistic material in (32b) as a matter for (pragmatic) stylistics, and treat it only as a 'linguistic' phenomenon, we need two separate theories of style. We would also need a mechanism for deciding when a case of language use, or a given repetition, was stylistic or 'ordinary'. It is much more economical from a theoretical point of view to have one theory of style which accounts for as many cases as possible. The term I adopt to describe this approach is suggested by Sperber & Wilson's (2006) paper *A Deflationary Account of Metaphor*, which argues that metaphor is nothing special or fancy; it is found in everyday conversation. In the same way, I adopt a deflationary account of style; style is concerned with the ordinary, and the everyday. Every decision *concerning form* that we take when we construct our utterances falls under the remit of style. In this way, this thesis consequently provides a deflationary account of repetition in that it demonstrates its pervasiveness in everyday ostensive-inferential communication.

2.7.7 Style and affective mutuality

Communication leads to cognitive mutuality - the extent to which the cognitive environments of individuals overlap in terms of all the assumptions that are manifest to them. This kind of cognitive overlap is not the only kind that obtains as a result of communication - as a result of two individuals engaging in communication both entertaining similar ranges of non-propositional effects, we can say that affective mutuality obtains (Sperber & Wilson, 1995, p 224). The non-propositional effects which Sperber and Wilson called poetic effects, and which comprise of a wide range of weak implicatures lead to common impressions rather than common knowledge (Sperber & Wilson, 1995, p 224). Individuals to whom it is manifest that

they are experiencing similar emotions or other internal states can also be said to be experiencing affective mutuality. Sperber and Wilson (1995, p 217) say that style is a relationship. By this, they (1995, pp. 217-218) mean that, through considering a co-communicator's choices of style, individuals can recognise the amount of potential mutuality between them, which has consequences for judgments about our social relationships. All of the repetitions addressed in this thesis lead to affective mutuality of a kind.

2.8 *Showing*, emotions, and intonation

The repetitions I address are involved in the recovery of non-propositional effects. Some of these effects are explainable in terms of Sperber and Wilson's (1995) poetic effects, while others seem to have a much more emotional element to them. Regardless of the type of non-propositional effects, it is quite clear that what repetitions communicate is NOT something that could be lexically encoded by a single word or device. The effects of many repetitions are quite indeterminate; the hearer is to recover a set of effects with no clear cut off point on her own responsibility. If an individual needs to communicate an interpretation that is indeterminate, non-propositional, and which cannot be linguistically encoded (because there would be no way to isolate and encode a guaranteed range of effects for acquisition), one way that he can do this is to *show* evidence for what he wants to communicate. Repetitions can communicate such vague interpretations and, so, it seems reasonable to expect that, at least in some cases, repetitions will be a case of *showing* evidence for what the speaker intends to communicate. I also sense, at this early stage, that the notion of *showing* will be essential to providing an explanation of emphasis. Moreover, I have said, without much explanation, that some repetitions appear to be involved in the communication of emotions. In order to speculate about how this might be achieved, it is necessary to think about what emotions are, and how they might be triggered by ostensive-inferential behaviours in communication. In §2.8.1 and §2.8.3, I discuss *showing* and emotions. To end the chapter, I consider briefly whether the kind of phenomena called expressives can shed light on how I could characterise emphasis, and, finally, I outline points concerning prosody and the use of the software Praat, which is used in this thesis to show the bounds of intonation groups, and the pitch heights of nuclei in adjacent intonation boundaries.

2.8.1 *Showing* and *Saying*

While Grice (1989) was embarking on his project on the nature of meaning, he developed the distinction between *showing* and meaning$_{NN}$ (non-natural meaning). Meaning$_{NN}$, for Grice (1989), involved linguistic coding, and, crucially, the recognition of intentions, which is one of his greatest contributions to pragmatic theory. *Showing*, for Grice (1989, p 218), is contrasted against 'telling' or 'saying' something, and is framed as 'deliberately and openly letting someone know' something (1989, p 218). Wharton (2009, p 11) notes that Grice, and those pragmaticians following him, tended to focus on meaning$_{NN}$, i.e., on the study of linguistic semantics, and, for example, the role that inference played in the recovery of certain implicatures within his framework (but it should be noted that Grice did consider non-natural meaning to concern more than linguistic semantics). Wharton's (2009) suggestion is that cases of *showing* are also very worthy of attention because cases of intentional *showing* are communicative, and we need to be able to handle them within a pragmatic theory of communication.

A significant contribution of Wharton (2009) was to show that the role of intentions was much more crucial to the interpretation of cases of *showing* than had previously been thought. Wharton's (2009) work is an extension and refinement of Grice's (1989). It is not necessary for me to set out the full history of the *showing-saying* distinction for this reason, as my intention is to build on Wharton's (2009) refinement that emerges from his discussion of the debate. His characterisation of *showing* and *saying* is now comfortably embedded in some of the most recent work in relevance-theoretic pragmatics (e.g, Scott, 2013, 2015; Sasamoto & Jackson, 2015; Wharton, 2015; Carston, 2016), and the reader is referred to his 2009 book for in-depth background to *showing* and *saying*. My concern here is to suggest that Wharton's (2009) approach is essentially correct, and outline important theoretical apparatus needed for the analysis of repetition and emphasis in this thesis.

Wharton (2009, pp. 40-41) says:

> '[I]n any act carried out with the intention of revealing an informative intention, there are two layers of information to be retrieved. The first, basic layer, is the information being pointed out, and the second is the information that the first layer is being pointed out intentionally. What makes a certain ostensive act a case of either

> 'showing' or ['saying'] is the precise nature of the evidence provided for the first layer of information. In cases of showing, the evidence provided is relatively direct...In cases of [saying], the evidence provided is relatively indirect - a linguistic utterance, for example.'

To use Sperber and Wilson's (1995, p 53) characterisation, *showing* involves 'strong direct evidence' for the first layer of information.

What is meant here is that, for acts of *showing* in communication, there is no great inferential leap from the evidence that is *shown* to the information the audience is meant to recover by reasoning from that which is *shown*. On this view, (33a-33c) are cases of *showing*:

> (33a) [You and I are sitting in my kitchen. You ask me if I have done the shopping yet.]
>
> I open the cupboard dramatically and gesture towards its empty interior.
>
> (33b) [You and I are sitting in my living room, chatting about our teenage years. You ask me if I had to wear a uniform at school.]
>
> I point to a picture of a school class in uniform that is hanging on my living room wall.

In (33a), there's no great inferential leap that you must undertake to infer from my *showing* of the empty cupboard to the conclusion that I have not been shopping. In (33b), it's not particularly effortful to reason from the picture I *show* to the conclusion that I wore a uniform at school - you can infer that I must be in the picture as it would not be relevant for me to *show* it to you otherwise, and, therefore, I must have worn a uniform. This conclusion preserves your expectation of relevance. Of course, the difference between the two cases here is that the evidence in (33b) is slightly more indirect in that there are additional inferences you need to draw to reach your conclusion that I did wear a uniform at school. This is suggestive of a continuum of cases between *showing* and *saying*, and, indeed, this is what Sperber and Wilson (1995, p 53) and Wharton (2009) propose. We move along this continuum depending on the indirectness of the evidence put forward for the first layer of information, linguistically coded evidence being the least direct, as suggested by the extensive involvement of inference in the recovery of explicit content seen earlier in this chapter.

What Wharton (2009) demonstrates is that it is important to infer what you are being *shown* and why in communication in order to determine what the speaker intends to communicate. For example, if someone thrusts a photograph in front of you, you need to infer what (element of the) photo is *shown*, and why it is *shown*, in order to understand what the speaker intends to communicate (Wharton, 2009, pp. 42-43). You won't know whether the 'mad photograph shower' is touting for work, telling you he has snapped your partner being unfaithful, or has recognised you in a recently recovered family photograph and thinks you're a long lost relative unless you can work out what he is *showing* you and why.

A continuum of cases between *showing* and *saying* is useful because it allows for hybrid cases of communication where it seems both *showing* and *saying* are present. For example, Wharton (2009), and Sasamoto and Jackson (2015) highlight that some cases of onomatopoeia seem to *show* and *say* at the same time, and Wharton (2003, 2009, 2015) suggests that certain interjections represent a mixture of *showing* and *saying*. Thus, the continuum is a useful tool that captures the complex, multi-layered and multi-modal nature of our communicative acts.

What sort of 'things' can we *show* as evidence for what we want to communicate? There are two ways to approach this question. The first might be to categorise some of the things we can *show* into natural signs, and natural signals. Natural *signs* are things like fingerprints on a glass window, or birds' nests in trees. They carry information which is evidence for a particular conclusion, but their function is not to do this (Wharton, 2009, p 13). For example, the function of a bird's nest is to house birds; however, nests are a natural *sign* that birds are present in the immediate environment. Natural *signals* have evolved to have the function of communicating specific information, e.g., the calls of monkeys or the dances of honey bees (Wharton, 2009, p 13). In human behaviour, the *sign/signal* distinction is also present. For example, shivering is a *sign* that you are cold (Wharton, 2009, p 13). Smiling, however, has the function of communicating an emotional state, and so is a natural *signal* (Wharton, 2009, p 13). A significant contribution of Wharton (2009) was to demonstrate that both natural *signs* and *signals* can also be *shown* in communication. You can deliberately and openly *show* that you are genuinely shivering in order to communicate, say, that someone should bring you a hot chocolate. You can deliberately and openly *show* that you are smiling in order to draw

further attention to your emotional state. The second way to approach the question of what we *show* is to consider the modes of the information presented as evidence.

If we think about the cases presented so far (*showing* an empty cupboard, a photograph, a smile, a shiver), paradigm cases of *showing* all seem to be within the visual mode. The evidence presented for the first layer of information can be *seen*. I suppose this is because these are the most intuitive cases when trying to explain what *showing* is. Entities in the visual mode are easily *shown*. We can just point to them or hold them up for people to see. However, it is not clear that *showing* ought to be restricted to the visual mode. Relevance-theorists don't actually restrict cases of *showing* to the visual mode, but I do suggest that there is a lot of focus on the visual mode, and *showing* concerns much more than evidence that we can physically see. Thus, it appears sometimes that *showing* appears restricted to the visual because there is not much *explicit* discussion of other cases. Wharton (2009) has suggested that onomatopoeia is a case of part-*showing* that retains a link to the domain of sound. Sasamoto and Jackson (2015) have suggested that novel onomatopoeia can be analysed as *showing* across various modalities. Sounds are not visible. Moreover, highlighting is *showing*, and Wilson and Wharton (2005) have suggested that prosodic behaviours can highlight constituents to draw attention to them. Prosodically highlighted words are also not visible. However, and this is important for this thesis, it appears they can be *shown*. We can *show* the linguistic forms that we utter. If there is some kind of connection between some stylistic repetitions and *showing*, this raises the question of what mode any evidence presented in *showing* via repetition would fall into. It is very possible that, if any cases of stylistic repetition are cases of *showing*, what we *show* is simply a form that has already been produced by displaying it again more or less ostensively. In verbal communication, we cannot see these transient utterance forms. Thus, some repetitions may represent another case of *showing* that is not visual evidence for the first layer of information. I would now like to turn to a possible diagnostic tool that might be instructive in analysing certain stylistic repetitions as cases of *showing*.

2.8.2 Sperber and Wilson (2015) on *meaning-showing* and *determinate-indeterminate* interpretations

In 2015, Sperber and Wilson published a paper, *Beyond Speaker's Meaning*. Their (1995, p. 117) claim is that an adequate theory of communication has to go far beyond how Grice envisaged speaker's meaning, and should be able to reflect the fact that human communication:

(a) Seems to comprise of a continuum of cases between *saying* and *showing* (where, at its extreme, *showing* is producing direct evidence for a conclusion); and
(b) What we communicate can be either extremely determinate, or very indeterminate. Thus, there is a continuum of paraphrasability/determinacy in terms of what we communicate (Sperber & Wilson, 2015, pp. 117-122).

Sperber and Wilson (2015, pp. 122-123) propose that these two continua (which they frame as *meaning-showing*, and *determinate-indeterminate*) correspond to two dimensions of utterance import, where import is understood as 'the overtly intended cognitive effect of a communicative act'. These continua are considered to be arranged orthogonally to one another, producing a two-dimensional grid space upon which the imports of various communicative acts can be plotted (Sperber & Wilson, 2015, p 123). My reason for setting out this grid is that this tool could be useful in diagnosing particular repetitions as certain cases of *showing*. Below is Sperber and Wilson's (2015, p 123) two-dimensional space in figure 1.

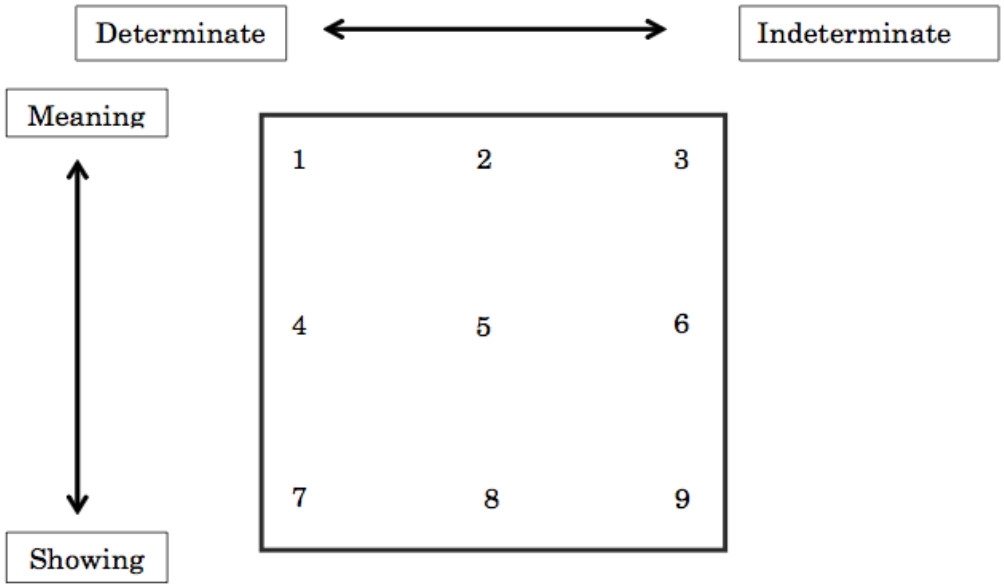

Fig. 1 Sperber and Wilson's (2015) two-dimensional space for plotting import of communicative acts

Using the numbers 1-9 on the grid, I set out below how the continua can interact to yield different imports (the values represent approximate locations, as we are dealing with continua). Some of Sperber and Wilson's (2015) justifications for why certain imports cluster around certain positions in the two-dimensional space hinge on speaker intentions - if all the evidence for the intended import of a communicative act comes from the speaker's intentions, the import will be in the vicinity of 1, for example.

> Position 1: **Determinate meaning**. Cases of 'pure meaning' such as this have all the evidence for the intended import coming from speaker intentions. Moreover, the meaning is determinate because it is readily paraphrasable as a proposition. An example of this is a railway employee neutrally uttering *12:48* when asked the time of the next train (Sperber & Wilson, 2015, p 123).
>
> Position 2: **Semi-determinate meaning**. This kind of case is diagnosed when the import is vaguer than for position 1. Hyperbole is given as an example, e.g., *I could kill for a glass of water*. The speaker's intentions are fairly clear, but the evidence for what is communicated is still quite indirect, but the meaning is less determinate than for position 1 (Sperber & Wilson, 2015, p 123).
>
> Position 3: **Indeterminate meaning**. Poetic metaphor is given as the paradigm case. The intended import is vaguer than for positions (1) and (2), and cannot be

paraphrased at all propositionally. It is still a case of 'meaning' as the evidence presented for the import is still indirect, however (Sperber & Wilson, 2015, p 123).

Position 4: **Determinate meaning/showing**. Here, the import is paraphrasable as a proposition, and so is determinate, but the speaker provides indirect and direct evidence for this at the same time. The example given is saying *He is* while pointing to someone in response to the question *Who is the tallest in the class?* (Sperber & Wilson, 2015, p 124).

Position 5: **Semi-determinate meaning/showing**. On a tourist trip, an individual utters *What a view!* The linguistic evidence he presents along with his tone of voice and facial expressions indicate the sort of conclusions the audience is to draw, but these cannot be easily captured propositionally, and, moreover, the evidence of conclusions to be drawn is both indirect and direct at the same time (Sperber & Wilson, 2015, p 124).

Position 6: **Indeterminate meaning/showing**. This time, the individual on the tourist trip utters *Wow!*. If there is any linguistic meaning, it only gives a rough clue as to the conclusions the audience could draw, and what the speaker communicates cannot be paraphrased propositionally at all. What is communicated is an impression (Sperber & Wilson, 2015, p 124).

Position 7: **Determinate showing**. Pointing to a clock when asked for the time is the example Sperber and Wilson provide here. All the import is determinate. It can be paraphrased fully propositionally, and the evidence is direct as it comes directly from the clock (Sperber & Wilson, 2015, p 124).

Position 8: **Semi-determinate showing**. Here, an example involving rain clouds helps us to understand *semi-determinate showing*. If we go out for a walk and I point to ominous-looking black clouds in the distance, it is not hard to reason to what I want to communicate - that it will soon rain. However, the import is less than fully determinate because there are a number of propositions that I can communicate by doing this. Also, though Sperber and Wilson don't spell this out here, I think there is some role for inference here in working out what is being *shown* (Sperber & Wilson,

2015, p 124). The evidence seems less direct than pulling information straight from a clock face.

Position 9: **Indeterminate showing**. The example provided here for imports clustering around 9 is simply showing someone a photograph of, say, your children. Certainly, no proposition would capture the import of the communicative act here. And, again, though Sperber and Wilson do not make this explicit here, I think there is inference involved concerning what is actually *shown*, and more so than for position 8. Are you shown the whole photograph? A portion of it? One particular child? (Sperber & Wilson, 2015, p 124). The evidence here seems much less direct than that pulled directly from a clock face.

In this thesis, as part of the analysis of each ostensive, stylistic repetition addressed, I attempt to locate its import on the grid. This should be important for determining whether or not stylistic repetitions have anything in common in terms of form, interpretation, or effects.

2.8.3 Emotions (and mirror neurons)

In this thesis, I adopt Wharton's (2009, 2015) characterisation of emotions, which is based on work by Rey (1980), and Cosmides and Tooby (2000). Wharton's (2015) reason for addressing emotion is to explore what expressives communicate, and I return to these in the next section. Wharton (2015) explores a link between procedural meaning and emotions to consider *how* expressives communicate. As I show in chapter five, the optimal way to process some repetitions is to say that they feed into and contribute to the calibration of a representation that is an output of a procedure that has already been activated. That is not to say that repetitions encode procedures. That seems highly unlikely given the discussion so far, and, in fact, if any repetitions of form can be shown to linguistically encode something, they are likely reduplicative - as I address in the next chapter. For this reason, I do not focus on procedural meaning, and consider mainly what Wharton (2015) has to say about emotions themselves.

Wharton's (2015) discussion of emotions focuses on what they are like, and what they do, so I deal with each of these in turn. Following Rey (1980), Wharton (2015, p 20) explains that emotions can be seen as involving several aspects: cognitive, qualitative, and physiological. On this view, the emotion of happiness would include cognition of something good that has

happened, a kind of qualitative feeling of happiness, and the physiological element - brain states and neurochemistry. Each of these is a sub-component of an overall emotion. I suppose that, if relevant, one of these sub-elements could become more salient in the interpretation of communicative behaviours. For example, an extremely high pitch could communicate excitement, which could cause the neurochemical aspect of an audience's emotional response to an interpretation to be foregrounded. This would be compatible with what emotions are considered to *do*. For Cosmides and Tooby (2000), the mind is made up of domain-specific programmes, each of which has evolved to solve a particular problem we have to face such as recognising faces or finding a partner. They (2000) say that emotions are overarching cognitive programmes with subcomponents (or, subroutines) that activate particular sub-programmes in response to particular situations. For example, if we're scared, the fear emotion will help us to focus our attention, and will release certain hormones in the body, and so on (*ibid*.). Wharton (2015) notes that there are emotion-reading procedures in existence, and they interact with these emotion programmes. Wharton (2015, p. 23) also says that emotional states can be contagious. What is clear from some of the repetitions I address in this thesis is that the express intention is for an audience to catch 'the feels' of their co-communicator. How might this work?

More than once in this thesis, I suggest that communicators may almost share emotional states or, in some pretty emphatic cases, particular aspects of emotions are foregrounded and passed on, as it were, such as shared physiological states - e.g., a knot in the stomach. I am neither an evolutionary psychologist nor a neuroscientist, and so I leave these issues to the experts in those fields. However, there is a link made in cognitive science between emotion recognition and mirror neurons, and between mirror neurons and empathy (Lamm & Majdandzic, 2015). Moreover, it has been suggested that there are specific mirror neuron systems for facial configuration recognition and for emotion recognition that work together (van der Gaag *et al*., 2007). It is likely, then, that such cognitive systems play a role in the recognition and experiencing of the emotional aspects of the interpretation of some repetitions.

2.8.4 Clues about emphasis from expressives

There is a category of expressions or behaviours that are considered in Relevance Theory to *show* the emotional states of speakers (Wharton, 2015), and this involves the activation of emotion-reading procedures (Blakemore, 2011; Wharton, 2015). These are called *expressives*. Expressives are a varied range of phenomena that include linguistic items such as epithets, some expletives, and diminutives, and non-linguistic elements such as gesture and tone-of-voice (Blakemore, 2011). Expressives are considered to indicate or express emotional states rather than describe them by conceptual means (Blakemore, 2011; Wharton, 2015.)

The following emboldened elements are analysed as cases of expressive use. In (34a) and (34b), the emboldened expressions are treated as expressives that also feature expressive prosody. In (34c), we see that words which are not necessarily expressive receive expressive prosodic features.

> (34a) [Adapted from Wharton, 2015, p 2. Imagine the emboldened element is produced in the lower portion of the speaker's range, at a low tempo, and with a growling voice quality.]
>
> **That bloody so-and-so** has retired.
>
> (34b) [Uttered with a high falling tone with a nucleus that features very high pitch, and loud volume.]
>
> **Ouch**! That hurt!
>
> (34c) [The emboldened nuclei are uttered with extremely high pitch (in the top portion of the speaker's range) and very loud volume.]
>
> | Oh my **DAYS**, | Justin **BIE**ber! |

The expressives in (34a-34c) would be treated as leading to the recovery of emotional states of the speaker via *showing* (Wharton, 2015). For example, the expressive prosody in (34c) *shows*, rather than describes in conceptual terms, the emotional state of the speaker: excited. I also want to draw the reader's attention to the fact that the examples here are *highly ostensive*. I particularly want to note that prosodic features seem to be particularly exploitable for 'cranking up' the ostensiveness of expressive communicative acts.

My intention here is not to provide a full account of expressives. The reader is referred to Potts (2007), Blakemore (2011), Scott (2015), and Wharton (2015) for an overview of the semantics and pragmatics of particular expressive phenomena. Diane Blakemore (personal communication) has often said that the effects of many repetitions have to be worked out in ways different from those of expressives, and she is right. Here, I want to consider whether there is anything that some repetitions and expressives have in common from the point of view of form that might be useful in this thesis. We have already seen that repetitions are sometimes called *emphatic*. However, I suggested in chapter one that no one had really been able to propose an account of what it actually meant to call a repetition emphatic. Potts (2007) described certain modifiers as *expressive* and *emphatic* at the same time. Romero-Trillo (2014) refers to emphatic expressives in greeting situations. The type of prosodic behaviours featured in the expressive cases above are sometimes called *emphatic* (Cruttenden, 1997). *Emphatic stress*, for example, is considered to be something like a combination of high volume and a high falling tone (Wales, 2001). When we consider all these things, what stands out to me is that anything called emphatic involves *highly ostensive* communicative behaviour on the part of the speaker. Expressives *show*. Repetitions might *show*. Perhaps emphasis concerns the level of ostensiveness found in the speaker's showing behaviour. I have this possibility in mind throughout the thesis.

The analysis of most of the repetitions addressed in this thesis requires some support from prosodic analysis. There are several reasons for this: (1) I wish to separate out which effects in an interpretation are triggered by the repetition, and which might be triggered by its prosodic features, (2) I wish to show that elements of prosodic behaviour can be repeated, and, (3) I need to show that particular repetitions are within or outside of an intonation group.

2.8.5 Prosodic analysis software and representing aspects of prosody

In this thesis, I use Paul Boersma's free software Praat to enable me to comment on:

- Obtrusion of pitch contours
- Pitch height of prosodic nuclei
- Intonation boundary placement

Praat can be downloaded at http://www.fon.hum.uva.nl/praat/, and a basic instruction manual can be found at http://savethevowels.org/praat/.

Pitch contours represent the perceived pitch or tune of an intonation group (Cruttenden, 1997). It is difficult to set out criteria for the delimitation of intonation groups (Cruttenden, 1997). However, there are some characteristics at the edges of these, and internal to groups, which can help us to perceive intonation boundaries. Praat can help to show these, and the identification of intonation boundaries is addressed in much more detail in chapter four, §4.3. Intonation groups have one nucleus, the last syllable in an intonation group that receives pitch accent (Cruttenden, 1997, p 44). Pitch accent is defined as an 'obtrusion' of pitch at the moment of accent away from the surrounding syllables (Cruttenden, 1997, p 40). (However, 'accent' in general is realised through pitch, length, *and* loudness (Cruttenden, 1997, p 40)).

In this thesis, I use | symbols to show some intonation boundaries. I use capitals to show which syllable in the intonation group is the nucleus, where this is relevant. Intonation boundaries are transcribed, and nuclei are marked in capitals only when it is pertinent to the analysis to do this. I may mention (but will not transcribe) the direction of pitch for a particular syllable or intonation contour. Generally, it will suffice in support of my points to show a Praat-generated representation of nucleus and / or particular intonation groupings.

Chapter Three: Repetition and reduplication: 'contrastive reduplication' as stylistic repetition

3.1 Introducing 'contrastive focus reduplication'

In chapter one, I raised the point that not everything that has been called repetition forms a unitary group. This chapter further narrows the scope of study by excluding morphological reduplication from a thesis on ostensive repetitions that communicate non-propositional effects. The chapter also draws attention to a phenomenon which superficially resembles morphological reduplication but is actually nothing like it, and is instead intended to establish greater cognitive mutuality during identification of explicit content: the so-called SALAD-salad

construction, or, contrastive focus reduplication. In this thesis, I am interested in reproductions of linguistic form which provide ostensive evidence for working out a speaker's communicative intentions in the communication of non-propositional effects. Identifying the appropriate data for this endeavour is not necessarily straightforward - as shown in chapter one. This is because there are also linguistic forms which include repetitions of form and can give rise to stylistic effects, but which we may have reason to distinguish from other stylistic repetitions of form. These are the cases of so-called contrastive focus reduplication mentioned above:

(Examples adapted from Ghomeshi *et al*., 2004)

(1a) I'll make the tuna salad and you make the **SALAD-salad**.

(1b) I'm up, I'm just not **UP-up**.

(1c) We're not **LIVING-TOGETHER-living-together**.

(1d) John: Have you seen the leak in the bathroom?

Julie: What leak?

John: The **LEAK-leak**!

The repetitions in (1a-1d) above are associated with a particular prosodic pattern, as shown in figure 2 below - accent on the first item, and a move from a higher pitch to a lower pitch.

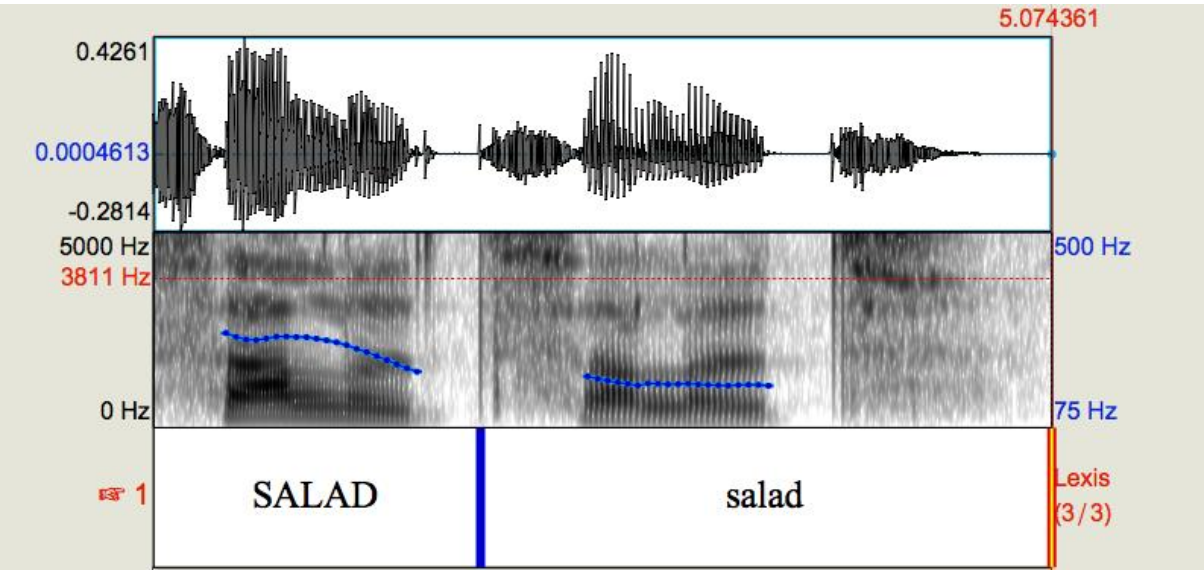

Figure 2 Pitch contour for 'SALAD-salad'

The above cases achieve relevance in a different way to those in (2a-2c) below, cases similar to those I investigate in chapter four, §4.3:

((2a) and (2b) adapted from Sperber & Wilson, 1995, p219.)

(2a) My childhood days are gone, gone.

(2b) There's a fox, a fox in the garden.

(2c) There were tourists, tourists everywhere.

I argue that the repetition of linguistic material in (1a-1d) is better analysed as modification which contributes to the identification of explicit content. At the same time, however, these reproductions of form seem to be sometimes further exploited or *played on* to give rise to weakly communicated stylistic effects which play a role in establishing affective mutuality between speaker and hearer. However, the matter of 'contrastive focus reduplication' is made more complicated by the fact that it has been incorrectly analysed as a case of reduplication in the first place. In order to show that contrastive focus reduplication is, in fact, not reduplication, I first establish diagnostic criteria for reduplication.

When conducting analysis of repeated linguistic material, we must ascertain whether an element is uttered because it is required by the grammar to express a feature or function, or because the speaker intends it to give rise to an effect which cannot be derived from linguistic or grammatical properties alone. In other words, we must consider what the semantics and pragmatics of repeated expressions contribute to an interpretation. The distinctions drawn in chapter two are important in discussing the phenomena addressed in this chapter. I believe that the failure to acknowledge the sorts of distinctions set out in chapter two has lead to the misdiagnosis of so-called contrastive focus reduplication (hereafter X-x for brevity and neutrality) as reduplication.

This thesis examines the relationship between repeated linguistic forms and their interpretations. As such, there is a sense in which this is a thesis about style. This chapter builds on and develops the relevance-theoretic notion of style introduced in chapter two. Style is much more than aesthetics and decoration. Style, as Sperber and Wilson (1995) have claimed, also concerns how and why we manipulate linguistic form to help and direct our co-communicators to recover every aspect of meaning - propositional and non-propositional -

that we intend to communicate. Decisions about style can even apply down at word level. A speaker may choose to utter a word because it unlocks the conceptual address that he wishes the hearer to access, but he may also choose it for other reasons. Our decisions about style are multifaceted. For example, you may say to a friend that he is a nincompoop not just to communicate a concept NINCOMPOOP* that befits his idiocy, but because the word *nincompoop*, loosely put, sounds humorous.

There are many reasons why a speaker utters what he does. However, just as I wouldn't want to say that the word *nincompoop* linguistically encodes any humorous effects, I don't want to say that X-x encodes any stylistic effects associated with it. When analysing the meaning of an expression or construction, it's important not to attribute an aspect of the interpretation to the linguistic semantics of that expression/construction unless there are good grounds to do so. I believe that some of the stylistic effects which are pragmatically recovered in context on the basis of an X-x have previously been incorrectly analysed as being linguistically encoded, if picked up on at all.

True reduplication is a grammatical phenomenon (Aronoff & Fudemann, 2005), and not stylistic. To show this, I set out morphological and semantic criteria for identifying reduplication. Applying these criteria to X-xs, X-x fails the diagnostic tests for reduplication. If not reduplication, what is X-x? Why do speakers choose X-xs over other forms that could communicate a similar concept? As suggested above, X-x are a special type of modification which involves a repetition of form that exploits processes carried out in the construction of ad hoc concepts. However, these X-xs are also often associated with the recovery of very weakly communicated stylistic effects. In the final section of this chapter, I explain how these effects might be recovered and suggest how their recovery might (or might not!) set X-x apart from other cases of stylistic repetition, and discuss whether there is any element of *showing* present in cases of X-x modification.

3.2 Reduplication

Reduplication is a morphophonological operation in which all or part of a base morpheme is reproduced, or copied. Where the entire base is reproduced, we have full reduplication ((3a) below). Partial reduplication occurs when only part of the base is copied ((3b) below).

(Examples from Aronoff & Fudemann, 2005, pp. 167-168)

(3a) From Indonesian

Kuda = Horse

Kuda Kuda = Horses

(Full reduplication expressing plurality.)

(3b) From Latin

Mordeo = I bite

Momordi = I bit

(Partial reduplication expressing perfective aspect.)

Reduplication is rule-driven and sits at the interface between phonology and morphology; it concerns the construction of phonological form during word-formation. The calling card of reduplication and most of the phenomena discussed in this thesis is repetition of *form* so we require some robust formal diagnostic criteria to identify cases of true grammatical reduplication. Otherwise, it would be easy to conflate grammatical reduplication on one hand with stylistic repetition on the other. This can easily happen. For example, Gil (2005, p 33) sets out six diagnostic criteria for splitting apart reduplication and repetition:

	criterion	repetition	reduplication
1	*unit of output*	greater than word	equal to or smaller than word
2	*communicative reinforcement*	present or absent	absent
3	*interpretation*	iconic or absent	arbitrary or iconic
4	*intonational domain of output*	within one or more intonation groups	within one intonation group
5	*contiguity of copies*	contiguous or disjoint	contiguous
6	*number of copies*	two or more	usually two

Table 1 (reproduced from Gil, 2005, p 33) Gil's Criteria for telling apart reduplication and repetition

Gil (2005, p 31) *does* note that telling apart reduplication and repetition can be difficult. However, the table of criteria that he proposes for doing this is obviously problematic. The

fact Gil speaks about 'units of output' suggests that repetition at greater than the word level is also governed by the grammar. This cannot be correct given that we cannot acquire the aspects of meaning that stylistic repetitions can communicate. More worryingly, however, the diagnostics do not meaningfully split reduplication and repetition apart. You could have a reproduction of linguistic form where communicative reinforcement is absent, where the interpretation is iconic, where the intonational domain of output is within an intonation group, where the 'copies' are contiguous, and where you have only two copies. How would you know if it was reduplication or repetition? The only useful diagnostic might be the first, which says that reduplication occurs at or below word-level. However, as I show in chapter five, §5.4.7, there are repetitions at the word level which superficially look like cases of reduplication, but are not. What is needed are robust morphological criteria explaining what inputs and outputs a reduplicative rule is associated with.

3.2.1 Formal properties of reduplication

Speakers of languages with reduplication don't reduplicate just anything. They are constrained by grammars, which only license reduplication for bases of a certain size and category. These constraints apply during word-formation. According to Gil (2005, p 23), full reduplication takes an entire word as an input and delivers an entire word as an output, and partial reduplication takes a string of sounds smaller than a phonological word and yields another string of sounds smaller than a phonological word as output. Thus, base 'size' is a potential diagnostic for reduplication. If there is repetition at or below word-level, we *may* be dealing with reduplication. It is not only base 'size' that constrains speakers of languages that have reduplication. Speakers cannot just select any base equal to or smaller than a phonological word; they have to choose an input form of the right category. Recall example (3a) above from Indonesian. The rule that derives *Kuda Kuda* from *Kuda* specifies that it can operate on nouns. It might also specify what it cannot operate on - certain types of functional morphemes, for example. Now, we have three criteria for identifying true reduplication: reduplication is a type of word-formation involving the copying of base morphemes of a specifiable size and producing outputs of a specifiable morphological category.

Earlier, I suggested that the failure to acknowledge some of the distinctions made in chapter two is one reason why researchers have conflated reduplication with other repetition

phenomena. One key problem is not considering what degree of the interpretation can be put down to grammar/linguistic semantics, and what degree or aspect is pragmatically determined. To properly describe reduplication, we have to identify the point at which it is the grammar that has the final say in specifying the linguistic form of a constituent of an utterance. Based on the formal criteria discussed above, we are now in a position to do this.

The output of reduplication must linguistically encode the feature or function that is its meaning. Gil (2005, p 34) suggests that this semantic requirement allows us to distinguish between true reduplication and repetitions which are associated with but do not linguistically encode the type of effects that I am interested in. Gil (2005, p 34) says:

> '[r]epetition is often devoid of any meaning whatsoever. (By meaning, I mean the term in the narrow sense associated with linguistic semantics, not in any of the broader senses associated with poetics, semiotics and other disciplines.)'

What sort of things can reduplication linguistically encode? Rubino (2005) links reduplication with number, distribution, tense, aspect, intensity and conditionality, among other functions. How might a case of reduplication linguistically encode its contribution to an interpretation? Perhaps cases of true reduplication might be analysed as encoding a procedure which constrains how the function of reduplication interacts with the construction of explicit or implicit content. However, this is not a thesis about reduplication or procedural meaning, and I leave this analysis for another time.

I can now explain why true reduplication has no place in this thesis. Reduplication involves a morphological rule which operates on base inputs equal to or smaller than a phonological word. Base inputs to a reduplicative rule are accepted on morphological grounds. Since decisions about linguistic form, as far as reduplication is concerned, depend on identifying inputs of particular categories and producing outputs of a very specific kind, reduplication falls under the remit of the grammar. In addition, reduplication must linguistically encode the function that it is associated with. I tentatively suggested above that reduplication might encode procedures constraining the recovery of either explicit or implicit content. Bearing this in mind, it should now be clear why I do not investigate reduplication proper. Once a speaker of Indonesian, for example, wants to express plurality, he is required to encode this in a particular linguistic form (barring lengthy periphrasis). This plurality is then recovered

from the reduplication (perhaps via a procedure) by the hearer and built into an ad hoc concept communicated by the noun while she identifies explicit content. The speaker's 'decisions' about linguistic form are taken by the grammar, and any function or effect associated with the reduplication is *decoded* from the linguistic semantics. This is completely different from the cases I am interested in. If I wish to communicate a range of weakly communicated stylistic effects, the grammar of English does not require me to use a specific linguistic form. In fact, there are no coded means available for 'delivering' those effects, as I have been suggesting. Instead, I just have to choose from a variety of well-formed utterances which the grammar licenses for production once I have ascertained which utterance is most likely to communicate the desired effects. Reduplications linguistically encode what they are analysed to mean, while the type of cases I am interested in do not linguistically encode their associated effects.

3.3 Ghomeshi *et al*.'s analysis of 'contrastive focus reduplication'

At the beginning of this chapter, I introduced a phenomenon in English which is analysed as reduplication according to Ghomeshi *et al*. (2004): so-called contrastive focus reduplication. Ghomeshi *et al*.'s (2004) work is regarded as the definitive paper on this phenomenon, and it is often referred to as 'the SALAD-salad paper' after the example in (1a), reproduced from above.

(1a) I'll make the tuna salad and you make the **SALAD-salad**.

Below, I outline Ghomeshi *et al*.'s (2004) claim that these X-xs are cases of reduplication which typically encode a prototype. I set out how Ghomeshi *et al*. analyse the morphosyntax of X-x. Their analysis is carried out within Jackendoff's (1997) Parallel Architecture framework. I then address some problems with their findings, making reference to the diagnostics above, which X-x fails. Then, I outline how Ghomeshi *et al*. (2004) analyse the linguistic semantics of X-x and show that X-x does not linguistically encode the interpretations that Ghomeshi *et al*. insist are associated with it.

3.3.1 The parallel architecture

On the surface, X-xs seem to be made up of a particular structure, reproduced strings of sounds, and prosodic stress (which Ghomeshi *et al*. depict in all examples using capitals, but

say little about). This bundle of structural characteristics works together to linguistically encode whatever Ghomeshi *et al*. think X-x encodes. Since Ghomeshi *et al*. think X-xs are made up of these structural components, they need a framework which can account for the interaction between many linguistic interfaces. They (2004, p 336) 'seek an analysis of [X-x] into which it is possible to incorporate its meaning, its syntactic conditions, its reduplicative phonology, and, for those speakers who have them, its prosodic conditions' (by prosodic conditions, Ghomeshi *et al*. seem to mean length constraints). Since some of the examples in (1a-1d) involve repetition above the word level, Ghomeshi *et al*. need a framework which does not view morphology as separate from the syntax, but as a continuation of it. If they don't use a framework which treats morphology as an extension of syntax, Ghomeshi *et al*. would not be able to analyse all cases of X-x as reduplicative. The Parallel Architecture could offer an integrated account of X-x because it links syntactic, phonological and conceptual representations via interface rules (Jackendoff, 1997; Ghomeshi *et al*., 2004). To explain a case of reduplication for a given language, one simply needs to specify the interface rule for a stored item (a word, or a chunked expression). This is linked with how the Parallel Architecture views the lexicon.

Lexical items, as in traditional grammar, are seen as an association of phonological, syntactic and semantic features (Jackendoff, 1997). '[A] lexical item tells the grammar that its three sets of features [phonological, syntactic, conceptual / semantic] can be placed in correspondence in the three independent linguistic components…' (Ghomeshi *et al*., 2004, p 338). Lexical items (and affixes) are treated as mini interface rules which play a role in constructing sentences. This is another reason why it is attractive to the researcher who does not wish to split apart morphology and syntax. Affixes, phonological words and chunked expressions are stored in the lexicon and are associated with their own interface rules. The Parallel Architecture, for these reasons, seems a reasonable choice for someone wishing to analyse reduplication because it aims to posit rules linking linguistic forms with their interpretations. After all, as we saw above, true reduplication involves the production of a specified form that linguistically encodes its meaning or effects. An interface rule linking form with meaning might be one way to model how this happens.

3.3.2 X-x as reduplication within the parallel architecture

Ghomeshi *et al*. (2004, p 341) rightly observe that most work on reduplication concentrates on the phonological *copying* that is characteristic of it. They (2004) think that this focus on phonology neglects discussion of where reduplication sits within the grammar as a whole. For this reason, their analysis focuses on the morphosyntax of reduplication.

Ghomeshi *et al*. note that McCarthy and Prince (1986, cited in Ghomeshi *et al*., 2004, p 341) view reduplication 'as the realization of an abstract affix whose phonological content is an empty (or partially empty) frame; its segmental content is copied from the base in accordance with various phonological constraints'. Presumably because of the substantial analyses of researchers such as McCarthy and Prince, Ghomeshi *et al*. (2004, p 341) also opt to treat all reduplication as involving an ordinary morphological affix. This is in spite of the many cases of X-x that are intuitively unlike affixes of English (e.g., (1c)). Nevertheless, Ghomeshi *et al*. (2004, p 341) say that reduplication should have a totally empty phonological frame or, rather, a frame that contains what they call 'metaphonological content'. By metaphonological content, they mean that some kind of instruction such as COPY X is contained within a phonological frame and X simply specifies the base which is to be reduplicated. Ghomeshi *et al*. (2004) *must* consider, then, that the reduplicative rule for X-x stipulates the base forms that the reduplicative rule can operate on.

Ghomeshi *et al*. (2004, p 342) say:

> '[t]here turns out to be a straightforward way to express COPY X in the parallel architecture. A reduplicative morpheme, like any other affix, has to specify the base to which it attaches. The base will have an index that connects its phonological content to its syntactic content...we can attach this very same index to the phonology of the affix...It is virtually the same as the English plural, except for its phonological content.'

Ghomeshi *et al*. (*ibid*.) also note that some stipulation about base size must be built into their reduplicative rule. However, it is not clear that *size* is the best way to impose a constraint on what can be copied. Sometimes morphological items (for example, clitics and affixes) are copied and sometimes they are not. Ghomeshi *et al*. are right that we need to say something about what bases can be 'copied'. After all, we can only posit a formal rule if we can model its inputs and outputs. However, this is not something that Ghomeshi *et al*. are actually able

to accomplish, as I shortly show. Nonetheless, Ghomeshi *et al*. (2004) note that some speakers seem to have a length preference when judging whether certain X-xs are agreeable. However, this length preference might fall out of the processing load associated with cycling a large string in working memory and, as such, length preferences need not be ascribed to a formal rule.

In summary, Ghomeshi *et al*. (2004) treat X-x as a reduplicative rule that involves affixation in the way that the English plural rule does. The affix is treated as having a phonological frame with metaphonological content. The metaphonological content specifies that COPY X should be carried out. X is a suitable base form and, in the presence of this base form, reduplication is triggered and X is copied to yield X-x. To offer a fuller account, Ghomeshi *et al*. (2004) recognise they should say something about the 'size' of 'copiable' bases and say something about the syntactic categorisation of X-x but are not able to develop generalisations that account for all the data.

3.4 Problems with Ghomeshi *et al*.'s analysis of X-x

3.4.1 Issues concerning 'affixes' and 'copying'

The 'affixation' proposed for X-x seems distinctly uncharacteristic of English affixation. Consider the following derivational and inflectional affixes of English:

> Derivational affixes: -ful (joy-ful) -er (teach-er) -ship (friend-ship) -dom (king-dom)
>
> Inflectional affixes: -s (cat-s) -ing (walk-ing) -ed (jump-ed) -est (bigg-est)

Though I dislike appealing to 'size' when discussing X-x, it should be noted that affixes of English are made up of relatively small strings of phonemes at the surface level. This is not the case for the 'affixes' proposed for X-x. The below examples involve quite substantial strings.

> (4a) So Seb's GONE-TO-BED-gone to bed? Or will I find him playing with his toys in the middle of the night again?
>
> (4b) Laura: I wish Chris would stop being so outrageous.
>
> Dave: I wouldn't say he's OUTRAGEOUS-outrageous. He only wore that costume the once!

Even if X-x can be analysed as having a reduplicative affix, it is not possible to stipulate the 'size' of base that can be 'copied'. By 'size', I mean more than string length. Sometimes, for example, clitics and affixes 'copy' and sometimes they do not. Ghomeshi *et al*. (2004) are unable to model when clitics, and a range of other components, can take part in an X-x, and when they absolutely cannot. Sometimes idiom chunks taking part in X-xs are acceptable and sometimes they are not. Why should the examples in (5a-5f) be ungrammatical (or highly unacceptable, at least) when the examples in (6a-6f) are fine? And why is word-internal irregular plural morphology copied when the 'copying' of the regular plural suffix is disprefered in similar contexts?

(5a) * Everyone was carrying maps and eating ice cream so for your information there were **TOURISTS-tourists** all over Warsaw.

(5b) * I can see blocks of lego poking out, Seb, so your toys ARE under **THE-SOFA-the-sofa**.

(5c) * I wouldn't **DATE-A-date-a** linguist.

(5d) * This is **THE-MIKE-SMITH-the-Mike-Smith**.

(5e) * Julie's got veneers so when you say you like her teeth, it's not really her **TOOTH-teeth** you're talking about.

(5f) * You'll manage to submit your thesis in good time so I wouldn't say your **GOOSE-IS-COOKED-goose-is-cooked**.

(6a) Everyone was carrying maps and eating ice cream so for your information there were **TOURIST-tourists all over Warsaw**.

(6b) I can see blocks of lego poking out, Seb, so your toys ARE **UNDER-THE-SOFA-under-the-sofa**.

(6c) I wouldn't **DATE-date** a linguist.

(6d) Yes, this is **THE-the** Mike Smith.

(6e) Julie's got veneers so when you say you like her teeth, it's not really her **TEETH-teeth** you're talking about.

(6f) You'll manage to submit your thesis in good time so, no, I wouldn't say your goose is **COOKED-cooked**.

There is a perfectly good part-structural and part-lexicopragmatic explanation for what is happening in the above data. For now, it is worth making a few comments about the presence or absence of plural morphology in X-x, as the absence of regular plural morphology provides supporting evidence for claims I make about the syntax of X-x later on. In the right context, *TEETH-teeth*, *GEESE-geese* and *MICE-mice* are all grammatical X-xs. However, *TOOTH-teeth*, *GOOSE-geese* and *MOUSE-mice* are not. Furthermore, *TOURIST-tourists* sounds acceptable to speakers of several varieties of English but these speakers consistently reject *TOURISTS-tourists*. Why should this be?

My analysis of X-x rests in part on the fact the X and x are interpreted using information stored under the same conceptual addresses. Imagine a speaker that wants to communicate a single concept TEETH-TEETH* by uttering *TEETH-teeth,* which features irregular plural morphology. When interpreting the speech stream, the hearer encounters X, *TEETH*, first, and can access and work with the feature of plurality that the concept TEETH-TEETH* needs to have. However, a speaker may want to utter *TOURIST-tourists* and the concept communicated by *TOURIST-tourists* should also feature plurality. However, there is a rule of English that prevents regular plural *-s* from attaching to X, so X *TOURIST* is prevented from encoding plurality in its linguistic form as the attachment of the affix is blocked. The morphosyntactic rules of English prevent the speaker from constructing his utterance in this way. This argumentation later provides supporting evidence for the claim that X-x is a special type of modification.

Ghomeshi *et al*. make an odd prediction about the circumstances under which X-x reduplication is triggered. They (2004) say that, in the presence of a suitable base form, X-x is triggered. We have established above, for example that *teeth* can enter into X-x. But why doesn't the reduplicative process trigger every time someone utters the word *teeth*? Ghomeshi *et al*. do not address this. An obvious answer seems to be that reduplication is not triggered because the speaker does not want to communicate the aspect of meaning that would be decoded from the putative reduplicative construction. This seems reasonable, but, if it is correct, Ghomeshi *et al*.'s (2004) account must be modified in a particular way.

Unfortunately, this modification obliterates its descriptive power. We would have to change our X-x rule to say that a variety of expressions are tagged in the lexicon as suitable base forms for X-x reduplication. However, it would seem that almost any conceptual expression (or complex string that, when pragmatically enriched, delivers a conceptual representation) can be found in an X-x (and thus could be a base) - see (7a) and (7b) below:

> (7a) Sure, Kate Winslet's quite curvy but I wouldn't call her a **VOLUPTUOUS-SEX-GODDESS-voluptuous-sex-goddess**.
>
> (7b) Are you blind? They're in there. I can see the packet at the back. The biscuits are **IN-THE-CUPBOARD-in-the-cupboard**.

Since many strings of varying length, complexity and from different categories can be 'copied', the set of potential bases for X-x is so large and diverse that we cannot extract a rule from the data that has any explanatory power. If we cannot posit a meaningful rule, it is not clear we can claim X-x is reduplication. Above, I set out some diagnostic criteria for reduplication, and explained what reduplication is thought to be. X-x does not seem concerned with word formation in a narrow sense. This is because any 'affixes' proposed for X-x don't look or behave like affixes of English. It is also impossible - on Ghomeshi *et al*.'s account - to specify a morphological rule for X-x that usefully captures the base inputs it can take, and the morphosyntactic category of its outputs. For these formal reasons, it seems X-x is not a case of reduplication. Finally, I'd like to object to claims of reduplication in English on the basis that the English language is not generally considered a hotbed of reduplicative morphology. Bearing this in mind, when a theorist encounters X-x for the first time, it is possible that they consider reduplication as one possible explanation for it. However, given that reduplication would be unusual in English, we should consider ruling out other explanations before settling on a reduplicative account for X-x. Ghomeshi *et al*. (2004) mention the modificational feel of X-x several times in their paper and so had grounds to explore alternative accounts. There are also semantic and pragmatic reasons for saying that X-x is not reduplication.

3.4.2 Ghomeshi *et al*.'s analysis of the semantics of X-x

> 'The use of a word or phrase often leaves open some vagueness, lack of precision, or ambiguity. [X-x] is used as a way to clarify such situations by specifying a prototypical denotation of the lexical item in contrast to a potentially looser or more specialized reading' (Ghomeshi *et al*., 2004, p 311).

When Ghomeshi *et al*. say that the use of a word or expression leaves open vagueness, lack of precision and ambiguity, what they are talking about is the linguistic underdeterminacy discussed in the last chapter. The linguistic semantics of our words radically underdetermines the propositions they express (Carston, 2002). Let us compare (1a) and a slightly amended version of it, (8):

(8) You make the tuna salad, and I'll make the salad.

(1a) You make the tuna salad, and I'll make the SALAD-salad.

The task for the hearer of *salad* in (8) and *SALAD-salad* in (1a) is reference assignment. In neither case does the hearer *decode* a fully-formed, context-specific concept of SALAD*. In both cases, the hearer computes an ad hoc concept of a particular type of referent salad. Ghomeshi *et al*. would consider the second *salad* in (8) to be imprecise or ambiguous in the sense expressed above. There may be many competing potential salad referents that *salad* could pick out, e.g., potato salad, or a Japanese pickle salad. The concept of SALAD* that the hearer recovers on both occasions has to be inferentially determined taking into account the context. However, this is not something that Ghomeshi *et al*. mention in their paper. To aid the hearer in reference assignment, the speaker in (8) utters *SALAD-salad*. For Ghomeshi *et al*. (2004, p 308) '[t]he semantic effect of this construction is to focus the denotation of the reduplicated element on a more sharply delimited, more specialised range...[f]or a first approximation, we characterize this effect as denoting the prototypical instance of the reduplicated lexical expression'. At first blush, this explanation seems to work:

(1a) You make the tuna salad, and I'll make the SALAD-salad.

(9) Ale Aficionado: Fancy a beer?

Lager Aficionado: What, a BEER-beer? (Thanks to Tim Wharton for this example.)

Imagine (1a) is uttered by a speaker of American English in a kitchen in the USA. By uttering *SALAD-salad*, the speaker intends to communicate a very particular concept of SALAD*, namely one which is green-and-leafy and common in Anglo-American food culture. He intends to communicate just this concept of a prototypical SALAD* in contrast to the many other types of salad that exist in the world. *SALAD-salad*, for Ghomeshi *et al*., linguistically encodes this prototypical salad and contrasts it with the discourse-salient tuna salad and the remaining members of the set of potential denotata for *salad*. As a result, the hearer recovers a prototype and understands that the speaker of (1a) will make a bowl of salad containing green leaves as a chief ingredient and *not* any other salad. In the case of (9), uttered in the UK, the Lager Aficionado utters *BEER-beer* to encode a prototypical beer - the affordable, light-coloured, 4% strength lager found in every pub in Britain. The hearer is encouraged to place this prototypical beer in contrast with beers that would not generally be considered prototypical, e.g., dark, fruity ales favoured by the Ale Aficionado.

Since speakers of many X-xs seem to intend their prototypes to be contrasted against other potential referents, Ghomeshi *et al*. (2004) call X-x *contrastive* focus reduplication. However, the authors don't provide an account of contrastive focus. Nor is contrastive stress as a stand-alone phenomenon discussed in detail. Ghomeshi *et al*. (2004, p 317) do say '[t]he notion of a set of alternatives against which an expression is evaluated usually comes up in discussions of contrastive focus, and there is certainly some similarity between [X-x] and contrastive focus'. They (2004, p 317) also note 'the copy has a focus accent just like that of a contrastively focused modifier' but, unfortunately, do not exploit this essentially correct insight. The name *contrastive focus reduplication* suggests X-x linguistically encodes contrastive focus. Moreover, it seems that Ghomeshi *et al*. consider that the accent on X is part of the linguistic form that encodes the prototype *and* the contrastive focus. In other words, Ghomeshi *et al*. treat a focal accent as part of the construction and consider that it is the entire construction that encodes the contrastive focus, as well as the prototype. However, confusingly, it seems that sometimes Ghomeshi *et al*. attribute the focus only to focal stress rather than to the entire construction. To analyse X-x, as I later demonstrate, it turns out one needs to be very clear about what aspect of form is responsible for a given aspect of the interpretation.

3.4.3 X-xs do not always result in prototypical interpretations

Ghomeshi *et al*. (2004, p 308) consider that X-x encodes 'the prototypical instance of the reduplicated lexical expression'. There are several problems with this claim. First, there are X-xs which we don't seem prototypical. These are examples where X-x rather picks out less specialised interpretations as opposed to a 'prototypical' one. Furthermore, Ghomeshi *et al*. do not define what they mean by a prototype and do not discuss how we mentally represent, store, or access them. There is no review of foundational findings in prototype theory, in particular Eleanor Rosch's (1973) work on natural categories. I turn to Rosch's work shortly. But let me first consider cases of X-x which cannot be said to encode prototypes. Consider (1d) and (10a-10b) below:

(1d) John: Have you seen the leak in the bathroom?

Julie: What leak?

John: The **LEAK-leak**!

(10a) Lager Aficionado: Fancy a beer?

Ale Aficionado: What, a **BEER-beer**?

(10b) [A group of friends are discussing what music to play.]

ABBA Fan: I bloody love ABBA.

Metallica Fan: They're awful - what about some **MUSIC-music**?

Classical Music Fan: You call that **MUSIC-music**? Beethoven's **MUSIC-music**!

In the case of (1d), what is a prototypical leak? A leak is just that - a leak. It is not difficult to imagine a context in which prototypicality is not the point of this X-x utterance. Perhaps the only defining feature of a concept of a LEAK may be that some contextually determined amount of some contextually defined entity (a gas, a chemical, a liquid, information etc.) accidentally emerges from some contextually specified environment. It rather seems that *LEAK-leak* communicates something like *the leak we have already spoken about which is really obvious and is currently soaking through to the apartment below* rather than a prototypical leak. On another occasion in another context, *LEAK-leak* would receive an entirely different interpretation. Moreover, John clearly employs *LEAK-leak* in (1d) to communicate his frustration. This is an example of expressive use. *Ghomeshi et al*. (2004) cannot explain how

this expressive, non-propositional meaning is recovered, although many of their examples have a flavour of the expressive about them.

The BEER-beer example of (10a) is a reversal of the dialogue in (9). In (9), a Lager Aficionado uttered *BEER-beer* to encode something that British pub goers would have no problem ranking as prototypical - light, cheap lager. Yet, in (10a), the *BEER-beer* uttered by the Ale Aficionado is intended to pick out a dark, fruity, locally produced ale only. In what sense is this a prototypical beer for most UK English speakers? This is a case of X-x where it seems that it picks out a less specialised interpretation rather than a prototypical one. Finally, consider the differing referents of the *MUSIC-musics* in (10b). If *MUSIC-music* can pick out some of the best Swedish pop, heavy metal *and* German classical music, how can we maintain the claim that X-x linguistically encodes prototypes? Just as *LEAK-leak* above was associated with the recovery of expressive effects, the *MUSIC-musics* in (10b) above might be analysed as playing a role in the recovery of speaker attitudes to different types of music. How can we reconcile that fact that some X-xs do receive a prototypical interpretation with the fact that others do not? Given the discussion in chapter two about concepts and linguistic underdeterminacy, the answer should be no surprise and can be found by reflecting on Rosch's (1973) article *As A Category, It's Natural!* Prototypes are concepts just like any other and, when we do recover them, they are also contextually determined.

Let's think about what people do when thinking about, say, a prototypical beer or salad. People *judge* or *evaluate* things to be a prototypical example of something. Decisions about prototypes involve comparisons and inferences. Until Rosch began conducting her experiments in the 1970s, psychologists adopted a 'classical' view of concepts where the mental categories we had depended on the linguistic words we had for the categories. Rosch explains (1973, p 110) that psychologists thought:

> '...categories exist because we have words for them. For example, we have a category for animals that lay eggs, fly, have feathers, and chirp; the category is "bird". The classical view maintained that if we did not have a word for a bird, the category or concept or bird would not exist.'

Rosch (1973, p 111) challenged the classical view by claiming that categories exist independently of language. The classical theory predicts that all members of a category ought

to have equal status within that category. That is to say, any beer would be an equally good example of beer as any other. However, Rosch has shown that this is not the case. Rosch thought (1973, p 111) that the borders of our categories are fuzzy and that we have to make judgments about whether a particular object is a member of a category or not. We do this by comparing a candidate object to a prototype, which is a kind of good exemplar for a category and exhibits a set of features that potential category members share to a greater or lesser extent (Rosch, 1973). The classic example is that a robin is considered a prototypical bird because it flies and sings. A penguin, which neither sings nor flies, is considered a poor bird examplar because it shares fewer features with the prototypical robin. Thus, when we want to decide if something is a category member, we judge the candidate member against the prototype. But how can such comparisons show that the prototypes themselves are contextually determined?

If a speaker of British English is asked to name a prototypical fruit, he or she is highly likely to say *apple* or *banana* rather than *kumquat* or *sharon fruit*. This happens rapidly. In these cases, it's hard to claim the prototypical fruit is contextually determined. After all, we think about fruit quite often and our concepts BANANA and APPLE might become highly accessible and stabilised over time. The trick is to use Rosch's methodology but to ask people to think about something for which they do not commonly access a prototype. Ask someone to give you a prototypical kitchen implement or a prototypical item of stationery. The informant would have to consider what constitutes a good exemplar for these categories. They would have to decide what features a prototypical kitchen implement or item of stationery should exhibit. In other words, for these less-commonly considered objects, the prototype itself would have to be constructed in context. Since our mental representations of can-openers and biros are, after all, just another kind of concept, it is no surprise that they should be contextually determined. Any conceptual constituent in a mental representation has to be constructed in context. What we also learn from this line of thought is that people's prototypes can be extremely idiosyncratic and they depend on the assumptions and the conceptual addresses that are readily accessible to individuals. A reluctant cook might judge a can-opener or a fork to be prototypical kitchen implements. A trained pastry chef might be more inclined to name a pallet knife as a prototype, for example.

Rosch's work supports the view that prototypes are concepts like any other and are pragmatically determined in context. Ghomeshi *et al*. do not mention pragmatic inference in their discussion of what X-x means - yet, for me, X-x is clearly a matter for lexical pragmatics. As such, Ghomeshi *et al*. cannot explain why a pastry chef would name a pallet knife as a prototypical kitchen implement. Nonetheless, I can now say why many X-xs appear to receive a prototypical interpretation, post hoc. During reference assignment, a hearer follows the relevance-theoretic comprehension procedure, entertaining interpretative hypotheses about potential referents in order of accessibility, and stopping when her expectation of reference is satisfied. In a context where she is asked to think of a prototypical fruit, she will hit upon a referent which is highly accessible such as an apple. It is highly accessible because it is often retrieved since we encounter it commonly in our food culture. An apple is optimally relevant because an apple is prototypical enough for the hearer to abandon evaluating other fruits in this context. She will not hypothesise about kumquats, lychees and kiwi fruits, as this will yield no cognitive effects. An apple will do. The 'prototypical readings' associated with X-x are simply a result of accessibility of information in conceptual entries.

3.4.4 X-x does not encode contrastive focus

Not only are there examples of X-x which do not encode prototypes, but there are also cases of X-x which don't involve contrastive focus. It is theoretically always possible to contrast the concept recovered from X-x with another entity. However, it is not always clear that the hearer is intended to do so. A hearer should only entertain such a contrast if the speaker gives her grounds, by virtue of his ostensive behaviour, that this will yield some cognitive effects. Some contrasts are relevant and are intended by the speaker to be attended to and developed, and others are not. Not every contrast is relevant and to pursue each potential contrast would put hearers to processing effort that is not offset by cognitive effects.

Clearly, some speakers do intend the concepts communicated by their X-xs to be placed in an optimally relevant contrast with another entity. Let us return to (9):

(9) Ale Aficionado: Fancy a beer?

Lager Aficionado: What, a **BEER-beer**?

This time, imagine a change to the context. The Lager Aficionado wants to question the choice of beer that the Ale Aficionado wants them to drink. It is the Lager Aficionado who intends to communicate a concept of BEER* which is the fruity, dark beer that the Ale Aficionado enjoys and is *exactly not* like the light, weak lager he drinks himself, and which he manifestly does not want to drink at the time of uttering. It is optimally relevant for the Ale Aficionado to draw a contrast between the two types of beer. It is an optimally relevant contrast because its computation can yield a wide range of cognitive effects. The Ale Aficionado can recover many assumptions from such a contrast. He can infer that the Lager Aficionado dislikes ale and has no intention of drinking any that evening. He might even infer that he does not have as much in common with his friend in terms of taste as he previously thought.

This contrasts with cases like (1d) below, which feels distinctly non-prototypical.

(1d) John: Have you seen the leak in the bathroom?

Julie: What leak?

John: The **LEAK-leak**!

It's certainly possible to place the concept communicated by *LEAK-leak* in (1d) in contrast with other leaks. Perhaps there are other leaks in the bathroom in question. Perhaps body lotion has leaked onto a shelf or contact lens fluid has escaped from the lens case. There are certainly other leaks out there in the world: gas leaks, chemical leaks and information leaks, say, about government expenses or offshore bank accounts, for example. Why is it not relevant for Julie to place the concept communicated by *LEAK-leak* in (1d) in contrast with other leaks? Because it would put her to processing effort that would not be offset by sufficient cognitive effects in this context.

It seems that when Julie hears *LEAK-leak* in (1d), she is intended to recover a wide range of more-or-less weakly communicated assumptions. A simplified set of assumptions recovered by Julie might look like (11):

(11)

[John is speaking about the leak that he has told me about twice today already]

[The leak that John has told me about today twice already is getting worse]

[We really need to do something about the leak that John has already mentioned to me twice today]

[John is cross with me because he has had to remind me several times about this worsening leak]

[John thinks I am stupid because he always tells me that he finds forgetful people stupid]

The set of weakly communicated implicatures that Julie recovers play a role in inferences that she makes about John's state of mind and, in particular, John's feelings of anger towards her. However, crucially, none of the above implicatures would be recovered by placing the concept communicated by *LEAK-leak* in contrast with another type of leak. It seems that something encourages Julie to construct a very specific concept LEAK-LEAK* that helps her to access the weakly communicated implicatures in (11) above. What provides extra help to Julie in constructing this concept? The prosodic packaging of *LEAK-leak*. Suppose *LEAK-leak* is uttered in a drawn-out manner, with pitch accent on *LEAK*, exhibits an exaggerated pitch contour, and is recognisably louder in relative terms than the rest of the utterance that contains it. This would *highlight* X (see Wilson & Wharton, 2005), encouraging Julie to pay attention to it and search for additional cognitive effects. We may also be looking at a case of expressive use here. I return to both points shortly.

Ghomeshi *et al*. (2004) treated X-x as a construction with three specific ingredients: a particular morphosyntactic structure, 'reduplicative' phonology, and focal stress. They seem to bundle this stress together with the rest of the construction for most of the time, but sometimes seem to analyse it as having its own function. Though all attested examples of X-x do exhibit some kind of accent on X, it is not clear that this accent should be treated a requisite component of any 'construction' or that its presence always results in the same interpretation or effects. In fact, it is possible to communicate everything that *LEAK-leak!* communicates by simply uttering *the LEAK!* with more-or-less the same tone-of-voice as described above.

Nevertheless, Ghomeshi *et al*. do say (2004, p 344) there are interface rules that automatically correlate stress and information structure and, so, they don't see any reason to investigate

accent, or any other aspects of 'prosodic packaging'. This suggests that they *do* perhaps attribute contrastive focus to prosodic factors. However, an implicit consequence of this approach is that wherever you have such an accent, you must always recover contrastive focus. However, as Wilson and Wharton (2005) have shown, similar prosodic forms may result in different effects on different occasions. Furthermore, prosodic features such as pitch accent or pitch contour height, rather than expressing certain propositions or concepts themselves, can instead simply guide the hearer to one of a range of salient processing strategies (Wharton & Wilson, 2005). This means that, for some X-xs, the pitch accent on X may encourage the development of a contrast between a referent and another contextually determined entity if this is optimally relevant. This seems to be true for examples like (1a) with *tuna salad* and *SALAD-salad*. However, contrastive focus won't be recovered in contexts where there is not a salient competing referent, as in (1d). This means that X-x cannot linguistically encode contrastive focus.

If the pitch accent on X does not always result in contrastive focus, what *does* it do? It is hard to isolate what pitch accent contributes to an interpretation on any one occasion. In fact, it is possible that pitch accent could play a different role (*or more than one role*, as I show in chapter four), depending on the context. Indeed, pitch accent might serve to guide reference assignment, but, as part of tone-of-voice, it might also provide evidence that leads a hearer to calibrate a speaker's emotional state. It can be expressive. In any case, the important point here is that the pitch accent of X-x does not linguistically encode one single aspect of meaning, least of all contrastive focus.

While not all X-xs result in contrastive focus, I must explain why many X-xs do result in the recovery of such contrasts. To explain why, I introduce something that forms part of my own analysis. I have already said that I treat X-x as a type of modification - an idea that Ghomeshi *et al*. also mentioned in footnotes. If we suppose that X is a modifier and combine an assumption about the semantics and pragmatics of some modifiers with the claims of Wilson and Wharton (2005), I can explain why contrastive focus emerges in the construction of ad hoc concepts. Sedivy *et al*. (1999, p 116-119), in their paper on incremental semantic representations, review some of the semantics literature on different types of adjectives. Two broad points of view emerge: either adjectives linguistically encode any associated contrast, or adjectives give rise to a presupposition of contrast. After all, speakers use modification to

help hearers narrow down a set of possible referents. It sounds odd to utter *pass me the blue book* in a context where all the books are the same colour, for example. I adopt the view that the presence of modifiers alerts hearers to the possibility that there are other potential referents they could entertain. Sometimes, it might be relevant for hearers not just to assign reference but also to develop the contrast between the intended referent and other possible referents, while sometimes it is not. A speaker could use prosodic highlighting to indicate which of a range of salient processing paths the hearer should take, as Wilson and Wharton (2005) said. If a speaker of an X-x wants to make it manifest to his hearer that she should indeed contrast the concept communicated by X-x with another entity, prosodically highlighting X-x with pitch accent would be one way for the speaker to encourage his hearer to infer she should do so.

3.4.5 No role for pragmatic inference in Ghomeshi *et al*.'s account

I identify the following semantic issues with Ghomeshi *et al*.'s analysis X-x:

> 1) X-xs do not linguistically encode prototypes. Prototypes are contextually determined via pragmatic inference, as are all concepts.
>
> 2) X-xs do not linguistically encode contrastive focus. Where contrastive focus is associated with X-x, it is contextually recovered via pragmatic inference when optimally relevant to do so.
>
> 3) The focal stress that Ghomeshi *et al*. count as a component of X-x does not linguistically encode contrastive focus by itself. In the case of X-x, there is no one-to-one correspondence between the 'focal stress' and its contribution to the interpretation. Instead, the accent in X-x highlights the X, in turn suggesting an interpretive strategy the hearer should take, and this is also determined via inference.

There is a common thread linking these issues and it is that no role is afforded to pragmatic inference in the recovery of whatever a given X-x communicates. The word *pragmatics* is mentioned but a handful of times in Ghomeshi *et al*.'s (2004) paper and, even then, it is found mostly in footnotes. Certainly, Ghomeshi *et al*. neither define this term nor set out what they mean by semantics. Pragmatics has traditionally been defined by whatever formal semanticists could not explain (see Blakemore, 2002; Carston, 2002). Unrelated problems

that could not be accounted for in formal models were simply thrown into 'the pragmatics wastebasket' (Bar-Hillel, 1971). 'Pragmatics' was a bundle of ragtag problems that formal theories could not model. This is something that relevance theorists object to. The domain of pragmatics is not that which formal semantics cannot account for but is something that can be defined in its own right. As we saw in the previous chapter, relevance theorists define (linguistic) semantics as the domain of linguistic decoding and pragmatics as the domain of inference. However, if only by allowing things to fall into the wastebasket, the theoreticians who make use of the pragmatic wastebasket do have at least some notion of pragmatics and allow it to play at least a small role in their accounts of meaning phenomena. Ghomeshi *et al*., however, do not seem to even have a pragmatics wastebasket. Every aspect of meaning associated with an X-x (be it a concept, a sense of contrast, or some weak effects) is attributable only to the linguistic semantics of X-x.

Not affording pragmatic inference even a small role in the interpretation of X-xs leads Ghomeshi *et al*. to overburden its linguistic semantics. This is the chief cause of the issues outlined in points (1-3) above, and, in other words, as I suggested above, is a result of not paying attention to distinctions made in chapter two of this thesis.

3.5 A relevance-theoretic analysis of X-x modification

3.5.1 X-x as modification involving repetition of linguistic form

In a structure of the form X-x, I analyse X as a modifier which plays a role in the hearer's identification of the concept that is communicated by the head x, and so contributes to reference assignment. The hearer of an utterance containing *SALAD-salad* is expected to use the interpretation of *SALAD* in the identification of a specific kind of salad. However, the interpretation of this modifier varies from context to context. That is, the concept that it communicates is pragmatically determined. This suggests that we need an account of how the context is used to bridge the gap between the linguistically encoded meaning of a word and the (often very fine-grained) concept it communicates on a particular occasion. Such an account has been developed by Robyn Carston (2002) (see chapter two, §2.5.3), and I draw on this account here.

However, this cannot be the whole story. X (*SALAD*, for example) is no ordinary modifier. In contrast with, say, *GREEN* or *MEAT*, it is a modifier which is phonologically similar to the head. This raises the question of what kind of point is served by the choice of this kind of modifier as a means of communicating the intended concept. I argue that the decision to produce a modifier-head sequence in which the modifier is phonologically similar to the head has a two-part explanation. On the one hand, the phonological similarity cuts down the processing effort entailed by reference assignment and contributes to the establishment of a sense of cognitive mutuality between speaker and hearer. On the other hand, the phonological similarity between X and x can serve a secondary purpose: in the right circumstances, it *shows* the linguistic form itself that has been repeated to indicate to the hearer that she should collect extra effects over-and-above those on offer via mere reference assignment.

3.5.2 X-x has modificational syntax

Consider the following utterances:

> (12a) [Twitter, 04/2013, from a discussion about the welfare of someone's neighbours.]
>
> User: Are you sure they're not injured / dead? Is 311 like a not-quite-emergency emergency?
>
> (12b) [Twitter, comedian Susan Calman, 07/2013, about a French TV show.]
>
> Susan Calman: Almost time to watch the Returned. It's exciting because there ain't no party like a "FrenchTVShowAboutTheResurrectedWhoMayBeEvil" Party
>
> (12c) [Twitter, 03/2013, from a discussion about Easter confectionery.]
>
> User: Welcome to that time of year when you have Cadbury's creme eggs for breakfast, as permitted by the "What? It's an egg!" loophole.

There are important similarities between these examples and cases of X-x. The former shed light on how the latter are interpreted. X-xs and the cases above are examples of modification, albeit a *marked* kind. By marked, I mean that the above examples may require more processing, and, thus, entitle a hearer to expect more or different effects. In other words, the linguistic form of the proposed modifier stands out or attracts attention in some way.

Compare (13) with (12a). We have good grounds to suppose that the strings occurring before the head noun *emergency* behave as modifiers.

(13) User: Are you sure they're not injured / dead? Is 311 like a normal emergency?
(*normal* emergency vs *not-quite-emergency* emergency in (12a))

In (13), the head noun *emergency* is modified by *normal*. This adjective replaced *not-quite-emergency*, but the original meaning of the utterance has largely been preserved. Since *not-quite-emergency* and *normal* are inter-substitutable in these contexts, they occupy the same syntactic position as each other. If we want to say that *normal* is an adjective modifying *emergency*, *not-quite-emergency* must also be a modifier. Finally, it is clear that the concept of EMERGENCY* that the speaker of (13) wants to communicate is in some way shaped by *not-quite-emergency*. It is *not-quite-emergency* that helps the hearer distinguish this particular type of emergency from other more serious ones, for example. The source of the distinctiveness that separates this type of emergency from other types is found under the conceptual addresses that the string unlocks.

All of the examples above, and all cases of X-x, are instances of what Ghomeshi *et al*. (2004, p 308) jokingly refer to in footnotes as the 'you-can-put-anything-you-want-before-the-head construction'. However, note that for Ghomeshi *et al*. (2004, p 308), X-x and the cases in (12a-12c) appear to constitute *two different 'constructions'*, as seen in (14) below:

(14) LIKE-'EM-like-'em? Or, I'd-like-to-get-store-credit-for-that-amount like-'em?
(From Ghomeshi *et al*., 2004, p 308)

Ghomeshi *et al*. (2004, p 308) say '[t]his example also contains an instance of *another construction*, the 'you-can-put-anything-you-want-before-the-head' construction' (my italics). However, if we are to say that X-x is one type of construction but the cases in (12a-c) are another type, then we would need to posit different morphosyntactic rules for each proposed construction, and I am not sure this is necessary. Cases like *not-quite-emergency emergency* are clearly modificational. I now present evidence that suggests X-xs are also modificational, even if their form looks atypical of English pre-modification.

We can exploit a syntactic constituency test - *one*-replacement - to show how X behaves like a modifier syntactically. Consider (15a-16e):

(15a) Pass me the BLUE book.

(15b) Pass me the BLUE one.

(15c) Pass me the SALAD-salad.

(15d) ? Pass me the SALAD one.

(15e) Pass me the SALAD-y one.

In *one*-replacement, a bar-level projection is replaced with *one*. Instead of uttering *pass me the blue book*, in the right context, you may utter *pass me the blue one*. This test provides evidence for syntactic constituency. The test alone doesn't confirm whether or not X is a modifier. But, with certain X-x examples, the *one*-replacement test yields an unacceptable utterance. It *sounds* odd. How do we improve the acceptability of the utterance? We add the adjectivising / adjective-preserving suffix *-y*. Acceptability is restored, as we see in (15e). This shows that X is either already a modifier or is converted to one at some point. Of course, *one*-replacement cannot be carried out on every utterance because it does not apply to all syntactic categories but it does yield indirect syntactic evidence for treating X-x as modification.

Finally, recall that above in example (5a) I showed that the regular plural suffix *-s* cannot attach to X. Adjectives of English do not take plural morphology (*blue books* vs. * *blues books*, and **TOURISTS-tourists* vs. *TOURIST-tourists*). This provides further supporting structural evidence that X is a modifier.

There are clear semantic and lexical-pragmatic grounds for treating X-x as modification because X is the source of the distinctiveness that sets apart the concept that x is intended to communicate from any competing concepts that x could communicate. The presence of X encourages the hearer to construct a concept of X-x in a particular way. After all, when a hearer processes an X-x, she does not end up with two separate concepts X* and x* in her mental representation. The hearer just recovers one concept, but it is one that is shaped taking into account the presence of X, among other things. If you simply utter *pass me the salad*, I may pass you exactly the right salad. However, in a context where there are many salads that might do, it makes sense to modify *salad* in some way so as to help me narrow down the set of referents that I entertain. You might put accent on *salad* and utter *Pass me the SALAD*. You might equally just point at the dish you want as you speak. By the same token,

however, you might just use a modifier to aid me in reference assignment. It simply happens that, in English, you can use a modifier that is phonologically very similar to its head. There is no grammatical reason why the phonological form of a modifier X and the phonological form of a head x cannot be similar. There is no rule of English (barring phonotactic constraints) that says we cannot have similar-sounding words adjacent to one another. We can utter modifiers that are similar to their heads if we want to, but this might still be constrained in some way. What restrictions *are* placed on the form of X? As we saw in (5a-5f), not everything X-x is grammatical or acceptable.

3.5.3 What can acceptably enter into an X-x?

Ghomeshi *et al*. could not find a way to fully explain what could be 'copied' in their reduplication analysis. Recast, that problem translates to asking 'does anything restrict what you can put before the head in English pre-modification'? Indeed, there is a grammatical explanation for what can and cannot fill X. Compare (16a) with (16b):

(16a) * I wouldn't DATE-A date-a linguist.

(16b) I wouldn't DATE-A-LINGUIST date-a-linguist.

What comes before the head has to be a complete syntactic constituent. The ungrammatical X-x in (16a) and (16b) do not provide the complete syntactic/semantic skeleton required for the fleshing out of a concept. *DATE-A* lacks its object, for example. The ungrammaticality or unacceptability of such cases is a result of incomplete argument structure. However, this is not the full story. Any old constituent simply will not do. Thus, it is not just grammar that has a say in the form of X-x. Compare (17a) and (17b):

(17a) * I can see them poking out. Your toys are under THE-SOFA-the-sofa.

(17b) I can see them poking out. Your toys are UNDER-THE-SOFA-under-the-sofa.

The X-xs in (17a) and (17b) above both involve complete constituents. [the sofa] is a determiner phrase. [under the sofa] is a prepositional phrase. Why aren't both X-xs acceptable here? There are communicative and lexico-pragmatic reasons for what X can be. The form of X-x must indeed consist of a syntactic constituent but the shape the licensed constituent takes depends entirely on the concept intended for communication. In (17b), what is 'at issue' is where the toys *are*. In fact, we might want to say that the hearer is

intended to draw a contrast between where the toys *are* and all the places where the toys patently *are not*. To do this, the hearer needs to construct a concept of the location of the toys UNDER-THE-SOFA*. What linguistic material will give her access to this concept? A constituent that contains a preposition - in this case *under*. Thus, although *THE-SOFA* in (17a) is a constituent, it is not a constituent that will unlock the conceptual addresses intended for use by the speaker in the identification of the explicit content of the utterance, more specifically a concept of where the toys *are*. The form of X-x, then, also depends on the concept that we want to communicate. This is something that Ghomeshi *et al*. (2004) fail to recognise. The only other constraint acting on what can end up in X is string length. However, this is a function of the speaker's working memory and is not determined by a formal rule.

There is nothing from a grammatical point of view to rule out the possibility that a modifier is phonologically similar to its head. However, this still leaves us with the question of why a hearer should use this particular structure rather than another. Why utter *SALAD-salad* instead of *GREEN-salad*? What can a speaker gain from deliberately employing a modifier which is similar to its head? This is clearly a question for pragmatics, and, I suppose, for pragmatic stylistics because the choice of modifier depends on *how* the concept that the speaker wishes to communicate is communicated. I draw on Carston's (2002) lexical pragmatics to show why the choice of an X-x structure (rather than what I will call a Y-x one) provides a different sort of route to the assignment of reference - a different route because X-x requires the exploration of a smaller number of conceptual addresses than Y-x modification. This, in turn, assumes and, indeed, communicates a greater sense of mutuality than the route provided by (for example) *GREEN salad*. Secondly, I shall argue that because this X-x requires the hearer to richly search the conceptual entry for X and 'dig out' the information required for reference assignment very much on her own responsibility, X-x represents a potentially much weaker case of communication than Y-x modification does.

I compare reference assignment for *green salad, GREEN salad, and SALAD-salad*. I do this to show the particular role that the prosody of X-x can have in highlighting, on one hand, and in the communication of its own (often expressive) effects, on the other. I also contrast GREEN salad with SALAD-salad to show the latter is equally just modification, as I have suggested.

3.5.4 Modification with deliberately repeated linguistic material: the semantics and pragmatics of *green salad*, *GREEN salad* and *SALAD salad*

(18) [Two siblings discuss what to have for lunch. There are a variety of dishes on the table but only one salad.]

Laura: What would you like?

Chris: A bit of green salad, please.

I take Chris to have communicated a concept GREEN SALAD* which is made mainly of iceberg lettuce, and contains little bits of tomato and cucumber, and for the purposes of the discussion I assume that most readers agree with this interpretation.

In (18), Laura accesses her conceptual addresses for *green and salad*. Under her conceptual address for *green* are stored assumptions about the colour green, memories and experiences of green, and perhaps scientific knowledge of and impressions of the colour green. Following the relevance-theoretic comprehension procedure, Laura takes just the information stored that satisfies her expectation of relevance. In this context, guided by what she can see, she likely entertains shades of green and assumptions about green that are similar to those that might be accessed by mentally representing the salad that is on the table. She doesn't access assumptions about greens that are closer to blue, or the green hue her complexion took on last time her sister had a hangover. At the same time, Laura takes optimally relevant information from her conceptual address for *salad* and constructs an ad hoc concept GREEN SALAD*. She will access assumptions about salads like the one she can see - salads with plenty of lettuce leaves. She will not entertain thoughts about pasta salads or meat salads, as these lines of thought would yield no cognitive effects in this context. Her complex concept GREEN SALAD* has a compositional semantics. It features material stored under the conceptual entries for both *green* and *salad*. Guided by the contextual clues around her and her search for optimal relevance, GREEN SALAD* is a concept of a green-and-leafy salad that is typical of Anglo-American food culture - a salad just like the one she can see.

Note two things. Firstly, reference assignment in this context involved the activation of two *separate* conceptual addresses. Secondly, a particular green salad is mutually manifest to both Chris and Laura by virtue of them both being able to mentally represent the salad on the

table in front of them if required to do so. For this reason, *green* in *green salad* receives no accent. There is no need to provide extra prosodic guidance to aid Laura in reference assignment. To highlight *green* with accent would encourage Laura to look for extra or different cognitive effects, which, in this context, are not on offer. Both observations turn out to be useful for understanding how reference assignment is carried out for *SALAD-salad.*

I now compare the interpretation of (18) with that of (19). In particular, the interpretation of (19) is instructive concerning the role of pitch accent in reference assignment with modifiers.

> (19) [Two siblings are discussing what to make for lunch. There are no dishes that are already prepared.]
>
> Chris: What are we going to have for lunch?
>
> Laura: Can we just have salad?
>
> Chris: Well, I suppose I could make a GREEN salad.

In (19), by uttering *GREEN salad*, Chris also intends the recovery of a concept of a salad that is green-and-leafy - just like the one in (18). However, there is no salad present in the visual context. Bearing this in mind, why does reference assignment differ between *green salad* and *GREEN salad*? What role does prosody play here?

Certainly, a good deal of the analysis of *green salad* applies to *GREEN salad*. Just as with *green salad* above, when Laura hears *GREEN salad*, she accesses her conceptual address for *green* and selects optimally relevant information which is then used along with optimally relevant information stored under the conceptual address for *salad* to create a complex concept GREEN SALAD**. However, in (19), prosody plays a particular role in narrowing down the set of potential referents and contrasting the intended referent with the rest of the set of potential referents. This is because there is now a set of potential competing salad referents that are *green* in some sense. Chris wishes to communicate the concept of a salad that is green-and-leafy *as opposed to* any other green salads that he could make for lunch. Salads can be chiefly composed of green ingredients that are not lettuce, e.g., asparagus salads or Japanese salads containing seaweed. In fact, a salad could be *green* in the sense that its ingredients have been responsibly sourced and may contain no green-coloured ingredients at all. A speaker aiming at optimal relevance would be justified in modifying the form of his

utterance so as to aid the hearer in determining what type of green salad is intended as a referent on this occasion.

In (18) there are a range of dishes on the table in front of Laura and Chris, but the salad on the table is the only member of the set of dishes on the table that is a salad. For this reason, *green* is unaccented as it is not necessary for distinguishing that salad from the rest of the dishes that are on the table. In (19), however, Laura is required to settle on a referent that is a member of the set of green salads - none of which are present in the visual context. In order to construct the concept GREEN SALAD** that is intended for recovery on this occasion, Laura must use *GREEN* to guide her in modifying *salad* in the way Chris intended. Laura must identify only the member of the set of all green salads that will satisfy her expectation of relevance. The trouble is, as I showed, there are many assumptions and much encyclopaedic information stored under the conceptual entry for *green* that could be used to this end. How should Laura work out which information stored under *green* will help her construct a concept GREEN-SALAD** that is sufficiently like the one intended by Chris for successful communication?

By accenting *GREEN*, Chris draws attention to, or highlights, that constituent. Pilkington (2000) considers that prosodic emphasis can slow down processing during utterance interpretation, encouraging hearers to explore a conceptual address more thoroughly in order to identify just the information stored there that is intended for retrieval by the speaker. Laura is encouraged by the accent on *GREEN* to identify just the information stored under the conceptual address for *green* that allows her to distinguish the intended salad referent from other green salads. That is to say, Laura is encouraged by the accent to look for optimally relevant information about green-and-leafy salads as opposed to, say, salads that are green in the sense that they are produced sustainably. By Chris drawing Laura's attention to *GREEN* and, arguably, slowing down utterance interpretation in such a way as to encourage Laura to richly explore the conceptual address for *green*, she can identify the intended information stored under *green* and use just this information to modify *salad* in the way Chris intended. Since Laura is trusted to search through the conceptual entry for *green* and pick out the optimally relevant information stored there on her own responsibility, the concept communicated by *GREEN salad* is weakly communicated - much more weakly communicated than the concept communicated in (18).

(18) and (19) demonstrate that any pitch accent in modificational structures can provide guidance to a hearer about which information stored under a Y modifier's conceptual entry should be taken up to modify the concept communicated by the head x. The typical context for this is one where, in the absence of appropriate contextual clues, there are competing potential referents. Pitch accent can *also* perform this function in X-x modification. Thus, any difference in how modification is achieved via Y-x structures and X-x structures should not be attributed to any pitch accent that can be found for both types of modification. Finally, since Y-x and X-x modification are the same syntactically, any difference between the two types of modification cannot be accounted for in syntactic terms. I now show below that a key difference between Y-x modification and X-x modification lies in how their linguistic form exploits how speakers go about decoding and developing the explicit content of an utterance. Let us return to example (1a):

(1a) You make the tuna salad and I'll make the **SALAD-salad**.

In (18) and (19), Laura had to access two conceptual addresses in order to construct an ad hoc concept in reference assignment. An ad hoc concept of the type Y-X* was constructed which had compositional semantics as the Y-xs resulted in concepts featuring conceptual material from two conceptual addresses. It is in this respect that the linguistic form of X-x sends the hearer down a different processing route from the one involved in Y-x modification.

In Y-x modification such as *GREEN salad*, the hearer must activate *two* conceptual addresses and work on them both simultaneously in order to assign reference. The hearer of *GREEN-salad* 'juggles' two conceptual entries at once and searches through both conceptual entries at the same time. This is patently not the case for a hearer interpreting *SALAD-salad*. With *SALAD-salad*, only the individual conceptual entry for *salad* is activated. In (1a), the hearer only works with one conceptual address. In the case of *SALAD-salad*, because the concept intended for communication is recovered solely on the basis of a single conceptual address, it means that whatever information the speaker intends the hearer to build in to her concept SALAD-SALAD* to distinguish it from other salads, that information must be found under the conceptual address for *salad*. In the other examples, the information that allowed a *green salad* or a *GREEN salad* to be distinguished from other potential referents was found under the conceptual address for *green*. *Green* was the source of the distinctiveness. However, with

SALAD-salad, the source of the distinctiveness that will allow a *SALAD-salad* to be distinguished from other types of salad has to be found under the conceptual address for *salad*.

The concept of SALAD-SALAD* that the hearer wishes to communicate in (1a) certainly ought to contain the same aspects of green-ness and leafiness that the salad concepts in (18) and (19) do. However, these are not the only features that the speaker intends the hearer to include in her concept SALAD-SALAD*. In fact, as we will see, the information that is recovered while computing the intended concept is so weakly communicated and context-specific that it is information that would not feature in or result from some sort of standard conception of SALAD*. Imagine for (1a) that Chris (speaker) and Laura (hearer) are discussing getting started on making the dinner. They have already discussed what kind of salads they will make. It would be reasonable to expect Laura to remember what she agreed to make - a green-and-leafy salad. One does not typically forget what one has arranged to make for dinner. Thus, Laura constructs a concept of SALAD-SALAD* that requires or leads to the recovery of something like the following (simplified) set of assumptions:

(20)

[I agreed to make a green-and-leafy salad]

[The salad I make should resemble the salad we discussed in exactly this respect]

[The salad I will make will be unlike the salads we already ruled out]

[The salad I make will not contain meat because we already ruled this out]

[The salad I make will not contain pickles because we already ruled this out]

[The salad I make will not be a tuna salad because Chris is going to make one]

Where would Laura most readily be able to access the above assumptions? All of the above are assumptions about salads. All other things being equal, these assumptions would normally be most easily accessed under the conceptual entry for *salad*. Certainly, it might be possible to access some of the assumptions in (20) via the conceptual entry for *green* but this would be much more costly in processing terms - it would involve the activation of more conceptual addresses, and involve more inferential steps. This is because Laura would have to search through her conceptual addresses for *green* and for *salad* at the same time in order to identify

two sets of assumptions (some about green things, some of which would be salads, and some about general salads, some of which would be green) that can be put together to form the set in (20) above. This would be effort that would not be rewarded with extra or different cognitive effects in the context. The same assumptions might be recovered via, say, *GREEN salad* but at a higher level of processing cost, reducing the overall relative relevance of the utterance. This is one part of the explanation of why a speaker decides to use X-x modification instead of Y-x modification. A speaker of an X-x cuts down the relative amount of cognitive effort, in a sense, that his hearer should expend in reference assignment. X-x modification facilitates reference assignment by reducing the number of conceptual entries that must be activated and 'juggled' during ad hoc concept construction. Reduction of processing effort in this respect is not the whole story for X-x modification, however.

Information under conceptual addresses is retrieved according to accessibility. Assumptions which are commonly retrieved become more readily accessible the more they are accessed. When a hearer accesses her conceptual address for *salad*, there are some assumptions that spring to mind more readily. Though I do not wish to suggest that assumptions are stored in a permanent ranked list, the hearer of any *salad* would normally recover those most accessible assumptions first in the absence of encouragement to do otherwise. However, in the case of (1a), it is not (just) those readily accessible assumptions that Laura needs. Laura should also recover something like the much less accessible assumptions found in (20). However, as the assumptions in (20) are recovered almost entirely on Laura's own responsibility, we have to say that (20) represents a case of weak communication. SALAD-SALAD* is weak communication indeed. No phonologically distinct modifier has been provided to help the hearer settle on the intended referent. However, the speaker *does* accent *SALAD*.

The accent on *SALAD* is perceptually salient. Following the relevance-theoretic comprehension procedure, Laura pays attention to this perceptually salient word and will subconsciously ask herself why the speaker accented it, forming hypotheses about Chris' communicative intentions on the basis of the accent. Guided by contextual clues and an expectation of optimal relevance, Laura infers that she should explore the conceptual address for *salad* more thoroughly – in the same way as in (19). Laura can also infer, from the fact that no phonologically distinct modifier was used, that she has most of the responsibility for

constructing this weakly communicated concept, SALAD-SALAD*. In this way, Laura is led to the conclusion that she is being *trusted* to search through her conceptual entry and find the information that she needs to construct a concept of a salient salad as opposed to other types of salad. In this case, she is trusted to recover something like the range of assumptions in (20). X-x modification, in this way, leads to more cognitive mutuality than Y-x modification. The concepts communicated by X-xs can be very fine-grained and nuanced, and it is the computation of these which establishes a greater degree of cognitive mutuality than might obtain via Y-x modification. However, a substantial part of the greater degree of cognitive mutuality results from the trust that is placed in the hearer to come up with the concept SALAD-SALAD* on the basis of just a single conceptual entry. Laura realises, subconsciously, that she has been trusted to construct this concept herself. There is a sense of complicity, as Sperber and Wilson (1995, p 217) put it, between the speaker and hearer. Since it is mutually manifest to speaker and hearer that they are complicit in this endeavour, the degree of cognitive mutuality at this time is increased.

Considering what I have said about conceptual addresses, accessibility and accent, I can now sketch - in pragmatic terms - why Ghomeshi *et al*. are correct to claim many X-xs result in a prototypical interpretation. Prototype formation exploits accessibility of assumptions, and the extra time that X-x provides for rich exploration of conceptual entries. Moreover, it seems that prototypes are an excellent way of establishing even more cognitive mutuality between speaker and hearer than would be on offer with a 'non-prototypical' concept.

If it is mutually manifest to you and me that we are similar individuals who share a physical environment and share many cultural assumptions, should you ask me to think of a pencil, there is a good chance that my concept PENCIL* is something like your concept PENCIL*. I simply access my conceptual address for *pencil* and access the images, memories, and other aspects of encyclopaedic knowledge needed to construct PENCIL*. Since it is mutually manifest that our concepts of PENCIL* overlap to a good degree, we can say that some degree of cognitive mutuality is established between us. However, imagine if you are dissatisfied with my description of a pencil, or I am slow to answer. You may say *Come on! Think of a PENCIL-pencil'* with the mutually manifest intention that you want me to think of a good example of a pencil or, in other words, something like a prototypical pencil. Encouraged by the accent on *PENCIL* and facilitated by the reduction in processing effort from accessing only one

conceptual address, I have extra time and good grounds to compute a prototypical concept PENCIL**. The lack of a phonologically distinct modifier also communicates that I am being given the responsibility for finding the information that makes a *PENCIL-pencil* distinct from any other type of pencil. As a result of my understanding that I'm responsible for the bulk of the pragmatic work required to recover the prototype-like concept, greater cognitive mutuality is achieved than in the case of the ordinary pencil. But, prototypes can be used to establish *even more* cognitive mutuality than is outlined here.

When we construct non-prototypical concepts, we mainly construct only the concept which is playing an inferential role in the current interpretation process. That is, we construct a particular ad hoc concept for the purposes of achieving relevance in that context. There may, of course, be some incidental adjustments to our cognitive environments as a result of the new concept. For example, British children do not usually know that black swans exist. Computing an ad hoc concept BLACK SWAN* for the first time may lead a child to delete the assumption that all swans are white. However, the changes brought about by computing a prototype might be more subtle and wide-reaching than the case involving assumptions about swans. According to Rosch (1973) our prototypes are good exemplars of a category. Furthermore, the other category members are *ranked around the prototype* (Rosch, 1973). The prototype and the other category members are ranked by virtue of the number and type of features they share. If a hearer of an X-x recovers something like a prototype from it, then she has not just constructed the concept of a prototype. She has also re-ranked and reconsidered other category members. Why? Because the newly activated prototype has to share relatively more features with better exemplars of its category than worse ones. Existing connections between mental spaces with be strengthened (due to continued activation) and new connections between mental spaces will be created. The extra time an individual would have to explore the conceptual address for an X facilitates this. However, it is likely, in cases where it is manifest that speaker and hearer are intended to come to entertain the same prototype, the hearer ends up rearranging her categories and connections in a way similar to the speaker. Thus the cognitive consequences of prototype computation in X-x are wider-reaching in terms of cognitive mutuality than non-prototype cases of X-x. This is another way that relatively more cognitive mutuality can be achieved through a prototype. Since these type of X-xs seem to result in a lot of cognitive mutuality which, according to Sperber and

Wilson is the name of the game in communication, it's no surprise that X-x resulting in prototypical interpretations crop up frequently.

3.6 Is X-x modification a form of stylistic repetition after all?

3.6.1 X-x and style

The X-x phenomenon is an excellent example of how speakers formulate their utterances not only to communicate a concept but also to optimise the amount of cognitive mutuality between speaker and hearer. Humans like to share impressions, ideas, and jokes. It is advantageous (and often pleasurable) for us to be able to do so. We have to communicate in order to live together and survive, and to create the relationships necessary for this. If we did not need or want to increase cognitive mutuality between ourselves and other humans, we should not bother to communicate at all. Sperber and Wilson's (1995) approach to style suggests it cannot be restricted to its decorative or aesthetically pleasing aspects. Their view departs from the classical approach to style adopted by Aristotle (1963), who argued that while literal language is there to give us clarity, figurative language is there to prevent communication from seeming prosaic.

In fact, Sperber and Wilson's (1995) notion of style is not restricted to literary or poetic language; style concerns *any* decision one makes about the form of one's utterance, whether this be in considered writing or (less considered) spoken communication – in conversations about making salads, for example. These decisions are oriented towards the search for relevance and reflect (and sometimes convey) speaker assumptions about the hearer's contextual and processing resources. The result of such decisions may not be a specific assumption or even set of assumptions but, nevertheless, contributes to cognitive mutuality by marginally increasing the manifestness of weakly communicated assumptions. X-x increases cognitive mutuality by giving hearers the responsibility for the derivation of finely-nuanced ad hoc concepts which are sufficiently like the finely-nuanced concepts the speaker has in mind, and in this way its use communicates a greater level of cognitive mutuality than the use of Y-x modification, which leaves less responsibility to the hearer.

Affective mutuality and a sense of complicity are not the only things to emerge from processing an X-x. There are cases of X-x where their prosodic elements can give rise to weakly

communicated non-propositional effects which result in mutual affect between co-communicators. However, there is still more to be said about the recovery of non-propositional effects from an X-x than this. It is possible that in some of the cases of X-x speakers *exploit* the reproduced form that they use to communicate a concept to encourage the hearer to derive even more extremely weakly communicated effects. This is explained with reference to the notion of *showing* (Sperber & Wilson, 1995, Wharton, 2009), or, rather, *displaying*. However, because of the prosodic packaging of X-x, in order to test out this idea, I must assess the extent to which it is the prosodic aspects of X-x which are responsible for the recovery of weakly communicated non-propositional effects.

3.6.1 X-x and the non-propositional effects of expressive prosody

Recall example (1d):

(1d) John: Have you seen the leak in the bathroom?

Julie: What leak?

John: The LEAK-leak!

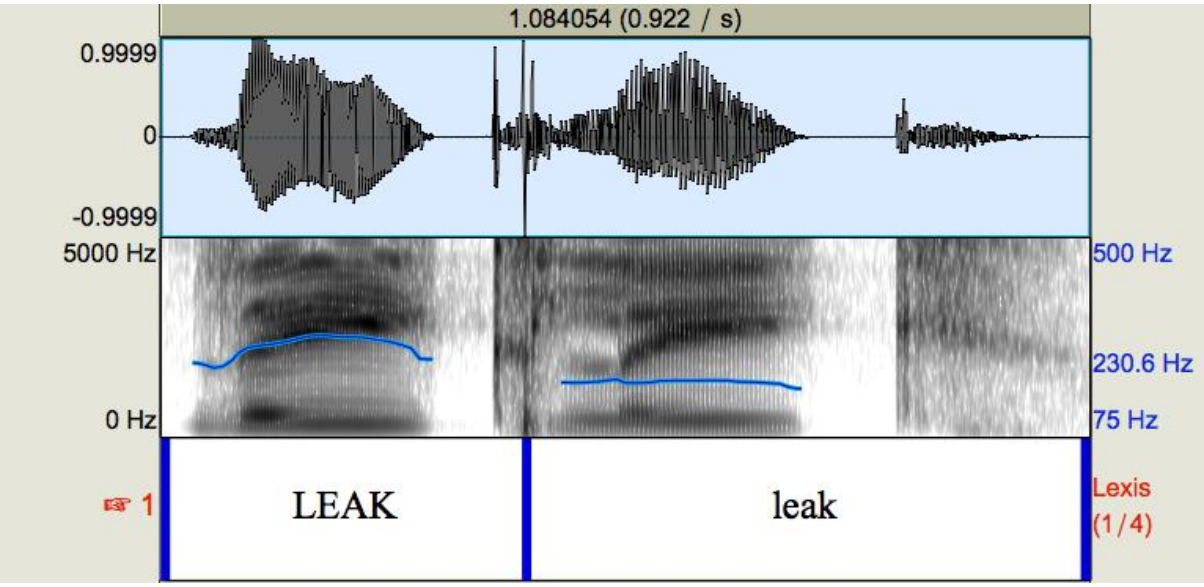

Figure 3 Pitch contour for 'LEAK leak' showing expressive prosodic features

I have already made the argument that the accent on X can, coupled with the hearer of an X-x being restricted to one conceptual address, lead to the communication of a very weakly communicated concept. However, the prosody of an X-x can lead to the recovery of non-propositional effects. The 'prosodic packaging' of an X-x is often expressive - the prosodic

make up of an X-x can indicate rather than describe the emotional state of the speaker in conceptual terms (see chapter two, §2.8.3). If we look at figure 3 above, we see that X is uttered with a very salient pitch accent and high intensity. The utterance was also quite loud. These features would, working together, activate a procedure or procedures in the domain of emotion-reading, which would allow Julie to recover and calibrate a representation of John's emotional state. John's prosodic behaviours are natural signals which communicate something about how he is feeling. I would also like to note that I think most people would be happy to say that John is being emphatic here. What is it about his behaviour that is emphatic? Certainly, John's prosodic behaviour is highly salient, and it is highly salient because he is extremely ostensive. The ostensive behavioural evidence that John puts forward for what he wants to communicate is dramatic, and attention-attracting. This might suggest that emphasis has something to do with the level of ostension ascribed to *showing* behaviours in ostensive-inferential communication. What is John *showing*? In one sense, his emotions. In another sense, his prosodic behaviours certainly highlight (and, thus, *show*) X *LEAK*, i.e., they focus Julie's attention on the word *leak*. However, there is another way that John *shows* something. I am going to argue that John, fairly subtly, further *shows* his X-x.

3.6.2 X-x as repetition that *shows* to guide the recovery of non-propositional effects

Above, I showed that some of the stylistic effects communicated by a given X-x can be attributed to prosodic factors. On one hand, the accent can guide a hearer in recovering a particular range of very weakly communicated weak implicatures, or, non-propositional effects. On the other hand, the expressive prosody can play a role in the recovery of weak non-propositional effects with a more emotional flavour. I want to say that there are also non-propositional effects communicated by the repetition of form in an X-x. The repeated form of X-x could play a role in the communication of a range of more very weakly communicated implicatures, or effects of the type that Sperber and Wilson (1995) call poetic effects. However, there are two points worth making here. First, it's not clear that the main or the sole reason for using X-x is just to communicate these effects. There is good reason to think they X-xs are motivated mainly by a desire to cut down processing effort on one hand, and to increase affective mutuality on the other. While a range of weakly communicated effects are recovered, this may just be a by-product of the speaker's desire to form his utterance so as to achieve what I just mentioned. Second, however, if this is all there is to be

said about how X-xs communicate non-propositional effects, then this analysis of X-x - apart from its prosodic overlay - would have nothing to do with *showing*, and I have been harbouring a suspicion that the repetitions I examine might *show*. Words are, almost always, treated as cases of *saying*. Utterances of linguistic materially typically provide very indirect evidence for what the speaker wants to communicate. However, I want to say, as I have already hinted, that words can be *shown* themselves. I think X-x is a subtle case of *showing* that can lead to the recovery of additional and very indeterminate poetic effects. It will be a number 9 on Sperber and Wilson's (2015) grid - a case of *indeterminate showing*. I would like to revisit the case of onomatopoeia to assist my analysis of X-x as repetitions of form that also *show*.

Although quite stabilised onomatopoeia have a conceptual semantics, they still retain their original resemblance to the sounds which they resemble, and, as such, provide relatively direct evidence for the information that the communicator wants to communicate (Wharton, 2009, pp. 99-101). The question which onomatopoeic expressions raise is what role ostension plays in any *showing* which is associated with them. Compare (21a) with (21b):

(21a) Brian **played** the guitar strings noisily.

(21b) Brian **twanged** the guitar strings noisily.

(21a) and (21b) might reasonably be said to express the same proposition. However, there are weak non-propositional effects that are recovered from (21b) that are lost in the paraphrase – some of these effects are to do with the quality of the sound.

There are two things that the hearer of (21b) might be justified in noticing as part of utterance interpretation. Firstly, since there are other words that can communicate a similar concept (such as *play*), the hearer can use the fact that the speaker of (21b) did *not* choose to use another available expression in his utterance as evidence in hypotheses about the speaker's intentions. The speaker might even deliberately make this assumption slightly more manifest through his ostensive behaviour, e.g., prosodically highlighting *twanged*. Secondly, the resemblance of the phonological form /twæŋ/ to the sound that a guitar string produces makes *the phonological word* itself even more manifest - it doesn't just *show* the sound which inspired the original onomatopoeic coinage. The hearer, again, might ask herself why that linguistic form was uttered rather than another. What is happening in this case is that the

relatively rare and, in some sense *marked*, onomatopoeic phonological form of the word is also being *shown*. The speaker could be thought of as deliberately and openly *displaying* this (choice of) linguistic form, and exploiting the fact that the onomatopoeic form that he has chosen for effect is already marked anyway. The speaker 'plays on' the onomatopoeic form he has chosen instead of *played*, drawing further attention to it or, albeit not very ostensively, *showing* it. Why would the speaker of (21b) *show* the onomatopoeia that he has already chosen for other communicative reasons?

Let's be clear that the *showing* I am talking about here is not that which is normally associated with onomatopoeia and which involves exploiting a relationship of resemblance between the onomatopoeic expression and a sound (Wharton, 2009; Sasamoto & Jackson, 2015). I'm talking about the speaker further *showing* the (choice of) form of his utterance. I think he does this to encourage the hearer to attend to the form of the utterance, ask why he uttered what he did instead of something else he could have uttered, and 'wring out' yet more very weak effects from its processing. Note that the hearer has to infer what is being *shown*, and how and why as part of this interpretation process. The hearer has to reason from the linguistic form she is *shown* to the information that the speaker wants to point out, and I would say that this information constitutes an extremely tiny increase in the manifestness of a large number of assumptions. The communication is very weak, and the range of effects to be recovered is extremely indeterminate - the type of thing that is often better off *shown* than *said*. If this is *showing* (or, perhaps, *display* is a better term), it is very, very subtle as far as ostensive speaker behaviour goes, as are its extremely nebulous effects. A picture is forming in this thesis of acts of *showing* that seem to differ in terms of their subtlety and salience. I think most people would agree with me when I say that the speaker of (21b) is categorically not being very emphatic, even if he is *showing*, for example.

With this in mind, let us return to X-x. A speaker of an X-x never *has to* utter it. His decisions of style depend on what he wants to communicate, and what he wants to make manifest about his relationship to the hearer by the deliberate choice to use X-x modification instead of Y-x modification. A speaker of an X-x is not required to use X-x by the grammar - it's a stylistic choice. And, simply put, that's why it can't be true reduplication. However, a speaker of an X-x is also well aware that the X- x he chose to communicate a concept is also *marked* because of it is made up of two adjacent phonologically similar expressions. Speakers of X-xs

can absolutely exploit this and may use it to their advantage in certain contexts. The speaker can 'play on' the fact that he deliberately *did not* use Y-x modification in a linguistic context where he could have and, indeed, might have been expected to do so. He can make the form that he has chosen a little bit more prominent, and exploit the existing markedness of the X-x over the Y-x he could have chosen. What is made manifest by deliberately and openly using repeated forms adjacently is the choice to use X-x itself rather than uttering a phonologically different form. This could be further highlighted by accenting the X prosodically (and prosodic highlighting is undoubtedly present in all X-s). By making the choice to utter X more manifest, the speaker intends the hearer to attend to the repeated form and ask why the speaker made that decision. She will infer she should 'squeeze out' even more weak effects from the bundle of communicative behaviours in an X-x that she has to process. As we have seen, the prosodic packaging of many X-xs gives rise to a wide array of weak, non-propositional effects, but what the speaker of an X-x can do over-and-above this is deliberately and openly display his choice of linguistic form X-x by ostensively not uttering something less marked. This, in turn, encourages the speaker to extend the array of effects that she is looking for. Since it would not take many inferential steps for the hearer to conclude this is what she ought to do, the evidence provided for pointing out the first layer of information is relatively direct. Thus, depending on speaker intentions, many X-xs can also be cases of *showing*. This will have consequences for how we understand *showing* in Relevance Theory. Linguistic forms are not visual. Ignoring writing, which is not natural to us as a species, phonological forms are transient sounds that disappear after we utter them. We cannot see the words we utter. X-x is another reason to consider extending the notion of *showing* to modes other than the visual.

Above, I suggested that X-x would appear in the vicinity of 9 on Sperber and Wilson's (2015) two-dimensional space, making them cases of indeterminate *showing*. Why? Recall what I said in chapter two, §2.8.2:

> Position 9: **Indeterminate showing**. The example provided here for imports clustering around 9 is simply showing someone a photograph of, say, your children. Certainly, no proposition would capture the import of the communicative act here. And, again, though Sperber and Wilson do not make this explicit here, I think there is inference involved concerning what is actually *shown*, and more so than for position 8. Are you shown the whole photograph? A portion of it? One particular child? (Sperber &

Wilson, 2015, p 124). The evidence here seems much less direct than that pulled directly from a clock face.

I suggest there are two reasons to think X-x represents a case of indeterminate *showing*. First, the hearer really has to work out what it is that is being *shown*, or further *shown*. For me, it's clear that inference is involved in working out that the speaker is very subtly displaying the form he has chosen on the basis of the hearer's ever-so-slightly ostensive behaviour. The evidence presented is still direct enough to be *showing*, but is much less direct than, say, a time on a clock face. Very crucially, however, nothing remotely propositional can capture the import of subtly showing your X-x. The import is extremely indeterminate. X-x makes a hearer pay more attention to itself, and wring out even more effects; I would argue that as it is not extremely ostensive, it serves to marginally increase the manifestness of a wide range of weakly communicated assumptions.

The analysis of X-x has showing suggests two points to bear in mind going forward:

(1) We do have a case of stylistic repetition which seems to be a case of *showing*, and, moreover, involves the *showing* of words. This is a case of *showing* which is not restricted to the visual.

(2) This case of showing would fall in the vicinity of 9 on Sperber and Wilson's (2015) grid. It is possible that any other stylistic repetitions that *show* will fall there also.

3.7 Concluding remarks – X-x has a place in this thesis after all

Reduplication proper has been shown to be a grammatically mandated process and, as such, has no place in this thesis. I have made a convincing case that X-x is NOT reduplication based on the conception of reduplication I set out in §3.2 of this chapter. If it were the case that the stylistic effects associated with an X-x are recovered on the basis of prosodic characteristics alone, I would exclude the repetition of form in X-x from this thesis. Certainly, the prosodic packaging of an X-x might lead to the recovery of some weak effects, but the use of a phonologically similar modifier and head also encourages a hearer to extend the range of effects on offer to her - of course, entirely at her own responsibility. We are talking about extremely weak communication indeed. I proposed that having two adjacent phonologically similar phonological words is unusual or noticeable in some way and thus is the type of thing

that a speaker might be able to exploit in order to encourage a hearer to derive extra effects during utterance interpretation. This was done using the notion of *showing*. In some contexts, a speaker of an X-x can deliberately and openly *show* his (choice of) modifier. What contexts? Contexts where the speaker could have and might reasonably have been expected to have uttered something else. By drawing attention to his choice of word (or choice of construction in longer strings), the speaker invites the hearer to hypothesise on the motivation for this and by doing so, she is intended to reason to the optimally relevant conclusion that she should recover a wider range of very weakly communicated non-propositional effects than if the (choice of) linguistic form had not been *shown*. If this analysis is correct, then we have to say that X-x does after all have a place in this thesis, and its analysis turns out to provide an important clue to how the study of ostensive and stylistic repetitions might achieve relevance. Though a speaker of X-x chooses a particular linguistic form to communicate a particular concept, he may also 'play on' or exploit that form for other communicative ends. He can *show* his choice of form to encourage a hearer to subconsciously reason from that to the first layer of information that he wants to point out. This might be how other stylistic repetitions also achieve relevance.

The analysis of X-x set out here ties in nicely with the relevance-theoretic conception of style that underpins this thesis. Decisions of style are multifaceted and complex. There is never just a single reason for uttering what we utter. In the end, we have to conclude that the repetition of form found in X-xs is, in a great many cases, both communicative *and* motivated at least in part by a desire to communicate extra weak non-propositional effects without linguistically encoding them. In other words, X-x is exactly the type of phenomenon I am interested in explaining, and the decision to focus this chapter on it turns out to be wholly justified.

Chapter Four: Repetitions in Adjacent and Non-Adjacent Intonation Groups

4.1 Repetitions of material within adjacent intonation groups

In the study of repetition, an important topic is epizeuxis, which is defined as the immediate repetition of a word or phrase (Pilkington, 2000, p 123). In Relevance Theory, the discussion of epizeuxis is part of a wider discussion of style (Sperber & Wilson, 1995). Sperber and Wilson discuss cases of epizeuxis to show that the use of a particular device is a consequence of how utterances are processed for relevance, and there are no one-to-one correspondences between stylistic devices and guaranteed effects (Sperber & Wilson, 1995, pp. 219-222). As such, the focus for Sperber and Wilson was not to give a full treatment of repetition, and this explains why their account ranges over a smaller subset of cases than that addressed in this thesis. However, their analysis provides a natural account of stylistic effects. For this reason, it provides an important starting point for the analysis of a wider range of repetition phenomena, which is addressed in this chapter and the next, including repetitions of material which are not repeated adjacently, or are repeated within the same intonation group rather than in separate intonation groups.

4.2 Sperber and Wilson's (1995) account of repetition

Sperber and Wilson (1995, p 219) present the following utterances for discussion:

(1a) Here's a red sock, here's a red sock, here's a blue sock.

(1b) We went for a long, long walk.

(1c) There were houses, houses everywhere.

(1d) I shall never, never smoke again.

(1e) There's a fox, a fox in the garden.

(1f) My childhood days are gone, gone.

All of these examples involve the immediately adjacent repetition of a word or phrase, separated by an intonation boundary.

According to Sperber and Wilson (1995, p 219), the repetition in (1a) encourages the hearer to construct a mental representation of three sock referents, which are similar to or different from one another by virtue of their colour. In the case of (1b), Sperber and Wilson (1995, p 219) say that the repetition of *long* encourages the hearer to assume that the walk was longer than expected. Likewise, in (1c), the repetition of *houses* encourages the hearer to infer that there were more houses than expected. In other words, Sperber and Wilson show that there are repetitions which achieve relevance by contributing to the recovery of the proposition expressed. This is not the only way that repetitions of the above form can contribute to the proposition expressed, however. Consider (1d):

(1d) I shall never, never smoke again.

The speaker of (1d) is an avid smoker of thirty years. She has just been told that she must give up smoking, something that she is reluctant to do, and laments this. In this case, the speaker repeats *never* to communicate her attitude towards the proposition expressed, and perhaps her degree of commitment towards it. As such, this repetition triggers the recovery of a particular higher-level explicature, the embedding of a lower-level proposition under a propositional attitude description.

Not all of the repetitions in Sperber and Wilson's dataset achieve relevance just by guiding the hearer's identification of explicit content. Consider (1e) and (1f), where the function of the repetitions is to contribute to the extent and nature of the emotions communicated by the speakers – the excitement felt on seeing a fox in (1e), and some kind of regret or sadness in the case of (1f). The range of effects communicated by (1e) and (1f) are extremely hard to pin down in propositional terms. Since the repetitions in (1a-1f) achieve relevance in very different ways, this raises the question of whether their interpretation can be accounted for in terms of a single-purpose interpretive rule or mechanism.

4.2.1 Interpreting repetitions with a 'catch-all' rule

The range of interpretations resulting from (1a-1f) cannot be explained in terms of a 'catch all' rule. This is not an economical way for humans to process repetitions, in any case. Such

an approach would require us to posit a formal interpretive rule which draws on a taxonomy of repetition types. On this view, a hearer must identify the right type of repetition from the taxonomy and interpret that repetition type in the way that the rule specifies. We have to ask whether it could even be possible to account for the type of repetitions found in (1a-1f) in terms of a rule that looks to systematically relate linguistic forms to their particular interpretations. In the cases of (1c) and (1e), for example, such a rule would not make the right predictions enough of the time. In (1e), although the speaker utters *a fox* twice, the speaker is not communicating that there are two foxes in the garden. Likewise, in (1c), although the speaker utters *houses* twice, he is not communicating that there are two groups of houses, but communicates that there is one group of houses which is larger than expected. In any case, even if we could posit some kind of interpretive rule to deal with these cases, this engenders further problems.

Firstly, such an interpretive rule must be built into the grammar and, as such, requires an amount of storage space, particularly when such a rule needs to link up with a potentially large taxonomy of repetition types. Mental resources, however vast, are finite and the wide range of interpretations which can be communicated by repetitions would require a relatively large amount of storage - if they could be stored at all. Secondly, storing a taxonomy of repetition types raises the questions of how this taxonomy would be acquired, and what would happen if a hearer encountered a repetition type that she had not encountered before. Surely, on this view, communication involving repetition of linguistic form would break down more often than it actually seems to. Thirdly, having a taxonomy forces us to posit an additional mechanism on top of the rule *and* the taxonomy – a mechanism that allows the hearer to judge which repetition type she should select from the taxonomy in the first place. It is not clear what the nature of this mechanism is, or when it would be permitted to kick in during the interpretation process. How is it, then, that hearers largely *do* manage to interpret repetitions effortlessly and, for the most part, as intended by speakers?

The key is not to assume that repetitions of linguistic material are taken care of by a formal linguistic rule. Instead, the interpretation of repetitions is dealt with by general principles and mechanisms that govern the interpretation of all communicative behaviours, whether verbal or nonverbal, and, importantly, whether involving the repetition of linguistic form or not. In other words, repetitions are processed in line with the communicative principle of relevance,

and the expectation is that the speaker of a repetition has aimed at optimal relevance, just as for any ostensive utterance. Faced with a repetition of linguistic form, the hearer should simply assume that any extra effort she has been put to by virtue of having to process the repetition will be offset by extra or different cognitive effects, and she must recover just the optimally relevant effects which satisfy her expectations of relevance in that context (Sperber & Wilson, 1995, p 220). This is a cognitively plausible proposal, as it does not require the storage and execution of special-purpose rules.

Let us see how this strategy plays out when applied to the data in (2), for example.

(From Sperber & Wilson, 1995, p 221)

(2) My childhood days are gone, gone.

Sperber and Wilson (1995, p 222) say that the first *gone* encourages the hearer to derive a range of implicatures about childhoods that she could reasonably have been expected to derive in that context. The second *gone*, however, encourages the hearer to consider that there may be a further range of implicatures that the speaker would like to back to some degree. In other words, the optimally relevant way to interpret the repetition of *gone* here is for the hearer to *expand* the context for interpretation and continue to *extend* the range of implicatures that she recovers. The hearer must *expend more processing effort* in doing so, and, as this is done at her own responsibility, this effort pays off with a rich range of weakly communicated implicatures. As Sperber and Wilson (1995, p 222) put it, the hearer then recovers a torrent of memories and thoughts about childhood. Through entertaining the weakly communicated implicatures recovered about childhood (childhood is a precious time, childhood is all too short, childhood is a time of innocence, etc.), the hearer may come to experience an emotional response of some sort. Sperber and Wilson suggest (1995, p 222) that we can say something about how this emotional response comes about. They (1995, p 222) say '[w]hat look like non-propositional effects associated with the expression of attitudes, feelings, and states of mind can be approached in terms of the notion of weak implicature…'.

4.2.2 The weak implicature account of repetition

For Sperber and Wilson (1995), when faced with a repetition such as the one in (2), the optimally relevant way to interpret it is to assume that it suggests a particular processing

strategy to the hearer. The repetition is an encouragement to the hearer to *expend* more effort in *extending* the context for interpretation. The hearer should access a wide range of weakly communicated implicatures, the search for which triggers the recovery of non-propositional effects in the hearer during the process. We can imagine that the hearer of (2) accesses a wide range of assumptions about her own childhood which, in turn, bring about an emotional and quite possibly sentimental response in her, a response which likely mirrors the feelings and sentiments that the speaker could be said to intend to communicate. Note that, on this view, Sperber and Wilson rely heavily on this notion of weak implicature to account for how non-propositional effects are triggered by a repetition like that deployed in (2). I will call the treatment of repetitions of the kind seen in (2) *the weak implicature account of repetition*.

In discussing the various effects of (1a-1f), Sperber and Wilson (1995, p 219) show that the effects of a given device are not constant because the resulting interpretations emerge from repetitions being interpreted in optimally relevant ways suggested by the context. Sperber and Wilson (1995, p 219) also note that 'the emphatic effects of repetition are worked out in different ways' and, where *emphatic effects* are recovered, it should also be expected that this happens in different ways given that repetitions are processed in the manner suggested by the context. However, what is important here is that Sperber and Wilson suggest that all repetitions are necessarily associated with the recovery of *emphatic effects*. We can see that, at least for some of Sperber and Wilson's cases, we do seem to recover some sense of emphasis. For example, in the case of *a fox, a fox*, we might suppose that the speaker is emphatic in his excitement about seeing a fox. The assumption that repetitions are, in some sense, emphatic is in line with what others have said about repetition (see chapter one, §1.6, and Pilkington, 2000).

4.2.3 Advantages of Sperber and Wilson's (1995) analysis

Sperber and Wilson's (1995) account of repetition is more cognitively plausible than a solution that requires rules and taxonomies for interpreting different repetitions. In fact, their account assumes a *natural* linkage between repetitions and their resulting interpretations in that effects follow from the general expectations of relevance that are raised and, indeed, are communicated by an overtly communicative act (linguistic or non-linguistic) – style arises in

pursuit of relevance (Sperber & Wilson, 1995, p 219). I understand *natural* to mean: following what the interpretive systems have evolved to be set up to do. In other words, a phenomenon receives a natural explanation if its interpretation can be accounted for by principles and expectations that drive the interpretation of communicative behaviours generally.

To see how a natural treatment is given to another case, consider how Sperber and Wilson (1995, pp. 222-224) treat gapping, a phenomenon in which a particular syntactic phrase is not phonologically realised in an utterance. Gapping could be thought of as a purely linguistic phenomenon (e.g., Waltraud, 1999; Hernandez, 2007), and on such an approach it would be dealt with using formal rules. However, gapping need not be accounted for in this way.

> (3) Diane lives in a beautiful house, Akis lives in a gorgeous apartment, and Phil in a bedsit.

(3) seems to be associated with the recovery of humorous stylistic effects which are hard to paraphrase in propositional terms, and which are lost when the omitted VP is re-inserted. Whilst the VPs in (3) are all syntactically similar, there is a semantic mismatch between the first two VPs and the missing one. The overt VPs concern desirable residences while the omitted VP concerns an undesirable place to live. Suppose that the omission of the final VP is ostensive enough to attract the hearer's attention and leads her to consider why the speaker has omitted the VP. Sperber and Wilson (1995, pp. 222-223) suggest that the omission of the VP is salient enough in the context to trigger a search for a context in which all VPs can have parallel contextual effects, i.e., be processed together for relevance. The hearer must 'bring together relatively unrelated encyclopaedic entries and construct non-stereotypical assumptions' (Sperber & Wilson, 1995, p 223) to arrive at an interpretation that satisfies her expectation of relevance. The optimally relevant interpretation here is one in which humorous effects are recovered about nice places to live and not-so-nice places to live. The omission of the VP is sufficient to trigger the hearer's natural tendency to look for parallel contexts, and parallel cognitive effects. Gapping can be accounted for in terms of how utterance interpretation is approached generally. No special linguistic rules are needed, even if the phenomenon exploits the manipulation of linguistic structure. The effort the hearer is put to in resolving the gapping is offset by optimally relevant stylistic effects.

In the same way, Sperber and Wilson's account of the repetitions in (1a-1f) is a natural account. The fact that something has been repeated is enough to suggest a particular processing strategy to the hearer which will yield optimally relevant cognitive effects in that context, although more needs to be said about how the hearer reasons from the ostensive repetition to the processing path she then undertakes. In any case, a natural account of repetition is of course advantageous for theorists with reductive tendencies. Less can be stored in the grammar, which cuts down on storage requirements and processing power. However, the *real* advantage of a natural account of repetition lies in the fact that it extends the range of repetition data which we can account for considerably. If Sperber and Wilson are right that certain stylistic devices are best given a natural treatment, the fact that we apply communicative principles rather than special-purpose linguistic rules to their interpretation means that we can extend the investigation to phenomena which are borderline linguistic or, indeed, clearly non-linguistic. Relevance theorists (Sperber & Wilson (1995) and (particularly) Wharton (2003, 2009, 2015)) have shown that both linguistic and non-linguistic behaviours are used in ostensive-inferential communication. These non-linguistic behaviours can also be used to communicate the stylistic and non-propositional effects of the type this thesis is concerned with. I extend Sperber and Wilson's (1995) natural approach to repetition to borderline linguistic and non-linguistic phenomena in chapter five.

Another advantage of Sperber and Wilson's (1995) account is that they have shown that the stylistic effects of repetition can be recovered both at the explicit level of utterance interpretation and at the implicit level of utterance interpretation. This is in contrast with accounts of stylistic phenomena that have assumed that stylistic effects always fall on the implicit side of utterance content (see Grice, 1989). Moreover, Sperber and Wilson's (1995) account of repetition also meshes well with the current relevance-theoretic understanding of the relationship between form and interpretation, or, the relationship between language and thought. In chapter two, §2.7, I showed that there is no need to assume that a constituent of an utterance corresponds to a constituent of a thought that a speaker wants to communicate. Since hearers of repetitions are simply to work out how they are to be optimally processed so as to preserve expectations of relevance, there is no need to assume that a repeated utterance constituent corresponds to a repeated constituent at the level of thought. In cases such as *my childhood days are gone, gone*, the repetition serves to trigger the recovery of a

wide range of weakly communicated implicatures that the hearer can assume is a faithful representation of the speaker's thoughts. The point of the repetition in (2) is not to trigger the recovery of a single thought that contains two mental constituents each of which corresponds to an instance of *gone*.

A final advantage of Sperber & Wilson's (1995) account of repetition is that it could be extended or adapted to the analysis of other stylistic phenomena. I think here of the appositional constructions addressed by Blakemore (2008) in her paper on apposition and affective communication. Consider the form of (4a-4c) below:

(4a) Chris was **depressed, flattened**.

(4b) The evening's been **spoilt, ruined**.

The appositions in bold resemble Sperber and Wilson's cases in two important respects. First, both the repetitions and the apposition cases involve the interpretation of immediately adjacent linguistic material, separated by an intonation boundary. Second, pre-empting my own analyses a little, the repetitions in (1a-1f) unlock a particular single conceptual region or fixed number of regions, while the appositions in (4a-4b) unlock two proximate, related, and, possibly, overlapping conceptual regions - thus, also a fixed number. A detailed examination of the form and interpretation of the type of repetition data presented so far in this chapter should be useful in analysing how appositions are also interpreted. I think here, in particular, of the role that intonation boundaries play in guiding utterance interpretation, and of the number of conceptual addresses accessed during utterance interpretation.

4.2.4 Questions raised by Sperber and Wilson's (1995) account of repetition

Despite the advantages of Sperber and Wilson's (1995) account of repetition, there are a number of questions which this account raises.

Sperber and Wilson consider that a sense of emphasis of some kind is associated with all repetitions. As we saw in the introduction to this thesis, other theorists have appealed to notions of *emphasis* or *being emphatic* in the discussion of data which have been labelled as repetition of one kind or another. In terms of pragmatics research, I have not found any account of what it means to *emphasise*, or to *be emphatic*, or any definition of what *emphatic effects* might be. There exist no cognitively-motivated accounts of these terms within

Relevance Theory, or any thorough treatment of the *emphatic effects* that are associated with other phenomena, e.g, some of Blakemore's (2008) appositions. Furthermore, in the introduction, I suggested that there has been some conflation in the repetition literature regarding the form and the effects of stylistic repetition. I have already suggested that emphasis is a property of speaker behaviour, and not an effect. As such, I cannot follow Sperber and Wilson in saying that all repetitions have emphatic effects, as the effects they propose are not emphatic.

A further question concerning Sperber and Wilson's (1995) account is raised by the sort of phenomena found in (5a) and (5b). In these examples, the items which have been repeated are, in contrast with the ones in Sperber and Wilson's (1995) original examples, non-conceptual and either borderline or non-linguistic. How should repetitions of such material be interpreted?

(5a) [A response to stubbing a toe.]

Ow, ow, ow!

(5b) [A text message from a friend to cancel a drinks arrangement.]

Friend: ☺ you're going to kill me, really ☺ ☺

The repeated items in (5a) and (5b) either encode, or activate procedures in varying domains of cognition. *Ow* is an interjection and is analysed as borderline linguistic (Wharton, 2009). Emoji and emoticons are (currently still!) non-linguistic. Sperber and Wilson's (1995, pp. 219-222) *weak implicature account of repetition* assumes that a repeated constituent generally encodes a concept which is an address leading to encyclopaedic assumptions. The effects of repetition emerge from the way it encourages the hearer to re-visit and re-search the same conceptual address in order to recover more and more weakly communicated implicatures (Sperber & Wilson, 1995, pp. 219-222). However, since the repeated material in (5a) and (5b) do not encode concepts, this leaves us with the question of how we account for such repetitions. I address this question in chapter five.

Sperber and Wilson's (1995) account of repetition only addressed cases of epizeuxis where conceptual items are repeated *adjacently* and separated by an intonation boundary indicated by commas. However, there are other stylistic repetitions of conceptual material which

Sperber and Wilson did not analyse. There are cases of stylistic repetition where the repeated constituents are not separated by an intonation boundary and which seem to be interpreted in a different way from the cases where the repeated constituents *are* separated by one, as in (6a-6c) below

(6a) | We went for a **long long** WALK |.

(6b) | He's a **big big** GUY |.

(6c) | I'm **very very** TIRED|.

Thus, repetitions can occur both within intonation phrases, and between adjacent intonation phrases. This raises the question of whether intonation structure plays a role in the sort of effects which are communicated in each case. In (6a), we might want to say that the speaker is communicating something about the actual length of the walk – just that it was very long. This is a little different to (7) below.

(7) We went for a long, long walk.

In (7), the speaker does more than communicate something about the length of the walk. I think he also communicates his attitude to the walk, and some poetic effects. The hypothesis which I investigate in this and the following chapter is whether intonation boundaries are used to indicate where the effects of repetition lie - either on the explicit side of communication or the implicit side. In cases such as | *long long* | and | *big big* |, the repetitions seem to contribute to the proposition expressed. In cases where an intonation boundary separates repeated constituents, it seems that the repetition is intended to contribute to the hearer's recovery of implicit content. The cases of repetition in which intonation boundaries play a role in the identification of explicit content are addressed in chapter five and are not discussed further here.

Apart from the cases in (6a-6c), there are other repetitions not considered by Sperber and Wilson. For one, their cases involve the adjacent repetition of quite 'small' syntactic constituents, e.g., repeated determiner phrases. Their examples do not feature any repetitions at or above clause level. For another, there are cases in which material *is* repeated in a separate intonation group, but where that repeated material is not immediately adjacent to the original. In other words, the repetition occurs at some 'distance'. Such cases also seem

to contribute to the recovery of implicit content – in the case of (8) below, these are weakly communicated humorous effects.

> (8) [From 'Catherine Tate's Nan', BBC1, 01/14, 'Nan' is sitting next to two women in burkas on a visit to the council offices. She speaks to the nearest woman.]
>
> Nan: Innit **hot**?...Innot **hot**, hey? Oooh, I am **hot**…You **hot**, love? You gotta be **hot**. I mean, it stands to reason, dunnit? I mean, if I'm **hot**, **you gotta be ROASTing**, ain't ya? **You gotta be ROASTing**. I don't know how you do it, darling. I really don't. I'd be **SICK** I'd be so **HOT**. I would. I'd be sat here, spewing. I'd be **SICK | SICK | SICK | SICK |**! I'd have a bilious attack and then I'd be **SICK** I'd be so **HOT**.

In (8) there is repetition of linguistic material in fairly proximate intonation boundaries with some intervening linguistic material between repetitions. However, we also find the repetition of material over much longer stretches of text. In other words, there is a greater amount of intervening linguistic material between repetition tokens. Consider (9):

> (9) [From 'Anthem', by Ayn Rand (2010). The speaker describes his various living quarters.]
>
> Page 5: The sleeping halls are white and clean and bare of all things, save one hundred beds.
>
> Page 6: The sleeping halls are white and clean and bare of all things, save one hundred beds.
>
> Page 14: The sleeping halls are white and clean and bare of all things, save one hundred beds.

It is clear that analysis of the data in this section requires consideration of the role of intonation boundaries. The range of data here also requires consideration of the notion of 'distance' between repetition tokens. What purpose is served by a speaker or writer producing repetitions over longer stretches of text or discourse? Do all such repetitions have intended stylistic effects? How is it is that hearers manage to identify that something has been intentionally repeated for effect in the first place? Sometimes, such repetitions involve only a handful of repetitions stretching across entire novels, films, or programmes, and we need

to be able to say how audiences notice that these have been intentionally repeated in the first place. Since the data in (8) and (9) also play a role in the communication of implicit content, the remainder of this chapter is dedicated to examining how repetitions associated with the recovery of implicit content achieve relevance. I think the best way to embark on this endeavour is to consider how intonation boundaries interact with adjacent repetitions to result in the recovery of non-propositional effects. The first question I address is whether the placement of intonation boundaries in (spoken) epizeuxis provides an indication to the hearer as to how the repeated material should be processed. The picture that emerges is a complex one where intonation boundaries in (spoken) epizeuxis can also play a role in further *showing* aspects of repeated form, and where more than just lexical form is repeated for effect.

Before proceeding, I wish to raise one more point regarding Sperber and Wilson's (1995) treatment of repetition: There is a final issue concerning the nature and recovery of non-propositional effects which arises from Sperber and Wilson's (1995) account of epizeuxis. Some of the effects of a deliberate repetition might be distinctly unpropositional and are much more like a feeling or physiological sensation. For Sperber and Wilson (1995), however, the recovery of such effects is treated as quite inferential, and depends on reasoning that involves, in the first instance, something propositional. Sperber and Wilson (1995) say that epizeuxis can encourage you to explore the context to recover more and more weakly communicated implicatures (which are propositional), which then leads to or tips over into something more emotional or affective. This could be correct in some more sentimental cases, or cases where the audience has time to consume and chew over a repetition, as in careful and relaxed reading of poetry. However, as we see in this thesis, the expressive cases of X-x or emphatic cases of epizeuxis being prime examples, there are deliberate reproductions of form which produce these very raw, physical, sensation-type responses, but the effects are not recovered in the same considered, reasoned way as the cases just mentioned. In such cases, I want to call these effects non-propositional because they *are* effects and because I am not even sure what kind of representation they involve, if any (it may just be we recognise and meta-cognise about what we physically feel). However, the important point is that these effects are non-propositional effects which do not rely on the necessary accessing of propositional representations for their recovery. The process is instant and reflex-like, as if a button is pushed. As such, Sperber and Wilson's (1995) account requires

some updating to account for the more 'raw' cases of epizeuxis. These repetitions and other communicative behaviours which intentionally provoke such effects in audiences are 'brute force' and 'quick-and-dirty' methods of establishing affective mutuality, whereby it is not only extremely manifest to both communicator and audience that both are experiencing similar sensations and emotions, but their sensations and emotions actually overlap. Perhaps something like the mirror neuron system is involved in the overlap and the recognition of this affective mutuality, which has the power to be of use socially in establishing who is 'in the same boat' from the point of view of how individuals feel. After all, we generally feel closer to those who respond to particular stimuli in similar ways to ourselves, and this, surely, is the power of certain repetitions in terms of their more interactive consequences.

4.3 Interpreting epizeuxis

4.3.1 Defining epizeuxis

Epizeuxis is considered a type of *emphatic* repetition (Pilkington, 2000; Jasinski, 2001; Baldick, 2008). As we saw in chapter one, other 'types' of repetition are also called emphatic. This raises the question of whether these different types of repetition are emphatic in the same sort of way. More fundamentally, it raises the question of how any sense of emphasis emerges. Given that I think emphasis is a property of speaker behaviour, I shall pay attention to the ostensiveness of the speaker behaviour in the production of utterances containing epizeuxis. In particular, I consider the possibility that the effects of epizeuxis are distinguished by the fact that they result from the hearer's perception of an intonation boundary. In such cases, the perception of an intonation boundary can be exploited as part of a speaker's emphatic communicative behaviour.

Without exception, every example of epizeuxis that I have found is minimally of the type X, X. The following are offered as (fairly typical) examples:

> (10a) [From King Lear, Shakespear, cited at literarydevices.net/epizeuxis, accessed 11/2014]
> And my poor fool is hanged! **No,**| **no,** | **no** life!
> Why should a dog, a horse, a rat have life,
> And thou no breath at all? Thou'lt come no more,

Never,|**never,**|**never,**| **never**!

(10b) [From an exchange between John Major and Tony Blair, 1997, taken from americanrhetoric.com/figures/epizeuxis.htm, accessed 04/2014]
Tony Blair: Isn't it extraordinary that the Prime Minister of our country can't even urge his party to support his own position? Yeah. **Weak!**|**Weak!**|**Weak**!

(10c) [The Bible, Samuel 2, Chapter 18, Verse 33, cited in Pilkington (2000, p 124)]
Oh Absalom, **my son,**|**my son**!

(10d) [From Othello, Act 2, Scene 3, cited in Farnsworth (2011, p 4)]
Othello: **Reputation,**|**reputation,**|**reputation**!**|** Oh, I have lost my reputation.

In every written case of epizeuxis I have found, and in every transcribed case of spoken epizeuxis, there is a comma or some other punctuation separating the repeated linguistic material. I have not seen anything called epizeuxis which does not exhibit punctuation symbols, which are generally accepted as representations of intonation boundaries in written texts (Wichmann, 2013). This suggests that an intonation boundary, or its written equivalent, is an essential ingredient of the form of epizeuxis, and would set it apart from other 'types' of repetition for those theorists interested in creating taxonomies of repetition based on form. Yet the presence of this intonation boundary (or its punctuational equivalent) is *never* explicitly mentioned in any description of epizeuxis I have found in the literature. In fact, researchers seem to look past it completely. Pilkington (2000, p 123) simply describes epizeuxis as the immediate repetition of a word or phrase. According to Wales (2001, p 133) epizeuxis is 'a FIGURE of REPETITION, with no words intervening', and Jasinski (2001, p 549) says much the same. Malmkjaer (2003, p 601) defines epizeuxis as the repetition of words or phrases 'without any break', and Garner (2009, p 894) says that epizeuxis is '[t]he immediate, emphatic repetition of a word'. Even in Sperber and Wilson's (1995) cases, the authors make no mention of intonation boundaries even though their examples contain them. Malmkjaer and Garner go as far as to explicitly state that there is *nothing* between repeated tokens in epizeuxis. But there is something there: an intonation boundary. And we need to be able to explain why it is there and what it does. Intentionally produced utterances have particular

forms according to the communicative intentions of the speakers. The decision to choose certain linguistic forms over others is part of a speaker's style. Rational communicators do not produce random, unconsidered forms. The intonation boundaries in epizeuxis are there for a reason, and their presence should not be overlooked in an explanation of how epizeuxis is interpreted.

It is not clear why the intonation boundaries found in all cases of epizeuxis have not been investigated, either from the point of view of literary stylistics, or the point of view of pragmatic stylistics. In terms of the former, it is likely that researchers simply focus on the fact that something has been repeated at all, and, rightly, note that there are particular effects to be had from repeated items that are relatively 'near' to one another. However, as an intonation boundary is not a word, or string of words, and as the study of intonation boundaries falls under the remit of linguistics proper, it is not surprising that intonation boundaries in epizeuxis have flown under the radar somewhat in the stylistics literature. Moreover, some researchers in literary stylistics might not necessarily be aware of the prosodic and pragmatic literature which highlights the role which prosody plays in utterance interpretation (e.g., Wilson & Wharton, 2005; Wichmann, 2013). Nevertheless, it seems that any discussion of the stylistic effects of epizeuxis must feature a discussion of the role that intonation boundaries play in how the repeated material achieves relevance.

As I have suggested, the first question I shall address is whether the intonation boundary characterising epizeuxis plays a role in guiding the hearer towards an interpretation at the level of implicit, rather than explicit, content. In fact, I am going to take this question back even further and ask how the intonation boundaries in (spoken) cases of epizeuxis arise in the first place. What are speakers doing in order that a hearer perceives an intonation boundary, and what are their communicative motivations for doing this?

4.3.2 The perception of intonation boundaries in epizeuxis

According to Cruttenden (1997, p 29), ideally, it should be possible to have a list of definitive external criteria, or phonetic cues, for detecting the placement of an intonation boundary.

Some of these indeed seem to be present where we would want to posit boundaries, but not all are present in all cases or, if such cues are present, they are not all present to the same degree each time. Some phonetic cues are more salient in some contexts than others. The main phonetic cues that Cruttenden (1997, pp. 30-44) says we can draw on in the identification of intonation boundaries are as follows:

> (1) **Pause**: Pauses can be empty, or filled. Empty pauses represent a break in the speech stream. So-called filled pauses, whether intended or otherwise, contain something else such as a drawn-out final syllable.
>
> (2) **Anacrusis**: The presence of extrametrical unstressed syllables after a proposed intonation boundary might serve as evidence of its presence. These syllables can be identified by their rapid tempo compared to the tempo of the tail of the preceding intonation group.
>
> (3) **Phrase-final lengthening**: Often, the last syllable of an intonation group is lengthened. This can be for a number of reasons. It can double-up to be a filled pause as in (1), or it can facilitate:
>
> (4) **Pitch change**: The nucleus in an intonation phrase is the last accented syllable in that phrase. In cases of unmarked nucleus placement, the nucleus may fall on the last syllable in an intonation phrase (e.g. with monosyllabic content words, and content words where the primary stress falls on the final syllable). As a nucleus is accented, it has pitch prominence, and one way that pitch prominence can manifest itself is through some kind of change in pitch.

Cases of spoken epizeuxis *can* exhibit the cues suggested by Cruttenden, just as any stretch of discourse featuring intonation boundaries can. Perhaps, then, if a speaker intends his hearer to perceive an intonation boundary somewhere, all he has to do is ensure that the above phonetic cues are *intentionally* produced in the right place, and to a degree that is salient in the context. We can exert great control over our articulatory apparatus, and so it is, of course, possible for a speaker to do this. What is less clear is that this is what always occurs whenever a hearer perceives an intonation boundary. Must speakers always actively produce these boundary cues in order to trigger the perception of an intonation boundary, or can the production of phonetic boundary cues fall out of something else that a speaker does?

4.3.3 The perception and role of nuclei / nucleic features in the identification and interpretation of intonation boundaries in spoken epizeuxis

There is a boundary-internal criterion which we can use to delimit intonation groups. Each intonation group contains a nucleus, and the nucleus is generally the syllable that has the most prominent pitch accent in an intonation group (Cruttenden, 1997, p 44). Davenport and Hannahs (2013, p 87) say that the nucleus, or tonic syllable, typically exhibits the most noticeable variation in pitch within the intonation group. In unmarked structures, the nucleus is usually the stressed syllable in the final foot in an intonation group (Davenport & Hannahs, 2013, p 87). The presence of a highly stressed syllable exhibiting extremely noticeable pitch variation, or functioning as a starting point for a change in pitch would be treated as the nucleus and can serve as an extra piece of evidence in the identification of intonation groups and their boundaries.

Davenport and Hannahs (2013) have noted that many nuclei are placed according to stress rules governing accent placement in certain *unmarked* structures. In other words, there are linguistic constraints governing where the nucleus of an intonation group is placed. However, in *marked* cases of nucleus placement, there must be other reasons for why the nucleus appears in a different place to the final foot of an intonation group. Or, even in 'normal', unmarked cases of nucleus placement, the existing nucleus could be made even more prominent still. This raises the question of why speakers might place their nuclei in marked positions, or make more prominent nuclei that are already prominent to a degree in an unmarked position.

Asking why a nucleus appears in a marked position has a simple answer within Relevance Theory. We can say that, in a great many grammatically unconditioned cases, the nucleus moves according to which utterance constituent is being highlighted. In other words, the speaker places the nucleus in a certain position to draw attention to the utterance constituent that features it (see Sperber & Wilson's account of focus effects, 1995, pp. 202-217). But what are the consequences of making a nucleus particularly prominent, regardless of where it may

be positioned in the intonation phrase, and how does nucleus prominence contribute to the perception of intonation boundaries, particularly with regards to spoken epizeuxis?

A nucleus typically hosts the most substantial pitch accent in the intonation group, and features noticeable pitch variation that placing that accent might involve. Pitch accent can be thought of as 'some sort of obtrusion of pitch at the point of accent from the pitch of surrounding syllables' (Cruttenden, 1997, p 40). Of course, these elements do not form the only cues for nucleic prominence. Other elements of tone of voice, such as intensity or loudness, can serve to make a nucleus more prominent (Cruttenden, 1997). In some cases, for example in cases of contrastive focus accent placement, a speaker may subconsciously decide to make a particular syllable more prominent in order to draw more attention to it. However, the point is that the presence of a nucleus need not arise from a linguistic rule (such as the sort of rule which might govern contrastive stress placement). Instead, it is possible that the perception of a nucleus arises from something else, namely the fact that a cluster of particular prosodic signals accompanies the speaker's utterance of a given syllable. This proposal in itself is neither new nor especially controversial. Gussenhoven (2002) acknowledges that while nucleic placement is linguistic, it also has other 'natural' or 'universal' motivations. For example, a nucleus may emerge because a speaker wishes to communicate something about his emotional state. As Wilson and Wharton (2005) have shown, in certain contexts, the optimally relevant way for a speaker to communicate his emotional state is to produce clusters of natural prosodic signals. The natural prosodic signals trigger procedures in the domain of emotion reading, but in doing this, the signals the speaker produces *also result in the perception of a nucleus*, or make the perception of an existing nucleus even more prominent. What I would like to do here is introduce a particular example of spoken epizeuxis. It is special because not only does it feature repetition of words in adjacent intonation groups, but it also features repeated nucleic placement with repeated nucleic height. And, in order to understand what role those repeated nuclei play in the perception of intonation boundaries in this particular case of epizeuxis, I would like to address what the speaker intends to communicate through the prosodic features clustering around the nuclei.

The case I examine is an elicited instance of Sperber and Wilson's *there's a fox, a fox in the garden*. The speaker is a thirty one year old female speaker of British English, and was asked to utter *there's a fox, a fox in the garden* in a frightened tone of voice as if she were trying to direct someone's attention to the animal outside. She was not given an example of what the utterance should sound like, or any other clues beyond being told to sound frightened. Below is the prosodic analysis of *a fox, a fox*.

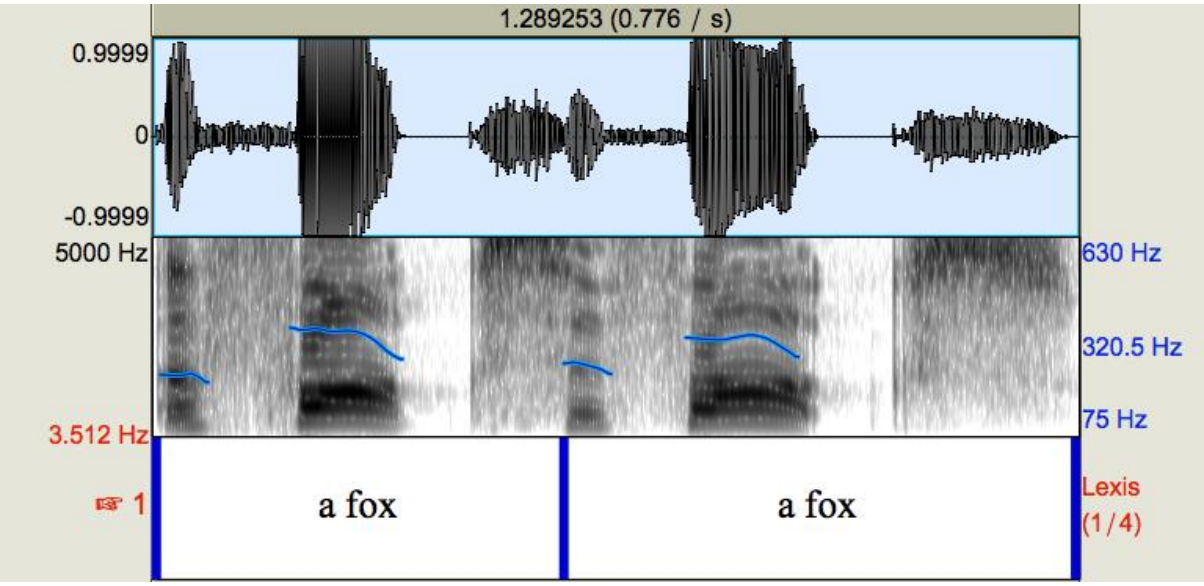

Figure 4 Pitch contour for 'a fox, a fox'

In figure 4, the nuclei of both instances of *fox* are very prominent (and quite high intensity), and occur at a pitch height consistent with excitement or fear. The degree of pitch excursion on the nuclei is relatively salient. Moreover, both nuclei are so perceptually similar in terms of pitch height and contour shape that the hearer would be justified in thinking that they have been intentionally repeated.

As we saw in chapter two, §2.8.1, some of our prosodic behaviours are natural *signals* that 'are genuinely coded and inherently communicative' (Wilson & Wharton, 2005, p 429).
These signals exist to encode information for co-communicators, and are interpreted purely by a decoding process (Wilson & Wharton, 2005, p 430). As I have already said, some prosodic elements are signals that activate emotion-reading procedures. What I would like to do is make sure that the prosodic features that contribute to the perception of the nuclei in this particular case of spoken epizeuxis are similarly natural signals whose interpretation is

governed by a code, and determine what the speaker wished to communicate by their production. What and how do these natural prosodic signals encode?

Carlos Gussenhoven (2002, 2015) has proposed that the interpretation of many prosodic features can be linked to three biological codes: the Frequency Code, the Effort Code, and the Production Code. According to him (2002, pp. 2-3), the Frequency Code concerns what we can decode about power relations in a situation based on computations we can make about the size, age, and gender of co-communicators by exploiting the well-known correlation between larynx size and body size, and the correlation between the size of the larynx and the frequency of the vibrations it produces. The Effort Code concerns a relationship between the effort expended in producing clear articulatory movements and the extent to which an utterance is viewed as important (Gussenhoven, 2002, pp. 4-5). Finally, the Production Code concerns a relationship between articulatory effort and utterance structure (Gussenhoven, 2002, pp. 5-6). For example, sub-glottal air pressure, according to Gussenhoven (2002, p 5), is greater at the beginning of exhalation, and this can be exploited to marry high pitch with the beginnings of utterances, and lower pitch with the ends of utterances. However, is a framework based on three separate codes the right way for us to model the interpretation of natural prosodic features?

There are three problems with Gussenhoven's (2002, 2015) biological codes approach. First, when Gussenhoven talks about effort, he means the speaker's effort. In Relevance Theory, it is acknowledged that both speaker effort and hearer effort affect the comprehension process. However, as Wilson and Wharton (2005, p 441) note, speaker effort and hearer effort do not vary in the same direction. For example, in certain contexts, should a speaker expend more effort on enunciation, this can diminish effort on the hearer's part. For Gussenhoven (2002), the amount of efforts a speaker expends on utterance production *decodes* to the degree he sees his message as important, or the degree to which he can be seen as forthcoming. Because Gussenhoven's (2002) model assumes a code, it means that whenever a speaker uses a certain amount of effort in articulation, then the hearer must recover a particular type of effect under that code. However, speakers sometimes have diminished control over how much articulatory effort they put into utterance interpretation, and sometimes, hearers pay no attention to the amount of articulatory effort the speaker expends. Furthermore, over-

enunciation could hinder communication, slowing down the speed of the interaction. Thus it is not always certain that a clear speaker intends to communicate helpfulness ostensively through Effort, and, in cases where Effort is ostensive, there is no guarantee via decoding that the hearer will interpret the speaker's effort in the way he intends. As Wharton and Wilson (2005, p 441) say, it might be better to treat articulatory effort as a *sign* of speaker helpfulness where it is relevant to do so. This relates to the second problem: the above observation means that Gussenhoven would want to use the codes to interpret prosodic features which might be better categorised as signs, which, although they may be exploited in communication, are not intrinsically communicative and are interpreted completely inferentially (Wharton, 2009). This is problematic because signs, which must be interpreted purely inferentially by definition, are made part of a code. It is more cognitively economical to have less coding, and to explain how natural signs are exploited in communication (see chapter two, §2.8.1).

A third issue is that it is not clear that we need three separate codes for interpreting natural prosodic signals. Clearly, we do draw on information about speaker Effort, or the frequency of the speaker's voice, or levels of articulatory energy when we interpret utterances. However, it seems to me that a particular prosodic feature which is said to be interpreted by one of the codes could equally be interpreted by one of the other codes. Or, that a particular prosodic feature could be interpreted by more than one of the codes at any one time. For example, we could argue that a speaker could put substantial effort into producing a particularly sharp intonation contour that is high in pitch - just like in *a fox, a fox* above. Such a pitch contour is intended to be highly perceptible to a hearer. As such, it might be interpreted using the Effort Code. However, high pitch contours could also be associated with the start of utterances under the Production Code. If such a contour were found utterance-initially, how would we know which code to use? In the absence of a mechanism for deciding which code should kick in during the interpretation process, the model falls down.

Let us consider what Relevance Theory might have to say about the interpretation of pitch height, and marked pitch excursions. As mentioned above, the notion of effort appealed to in Gussenhoven (2002) is one of speaker effort. However, in Relevance Theory, although assumptions about speaker effort may be factored into the context for utterance interpretation, what is most of interest is hearer effort expended in deriving cognitive effects.

Any effort that a hearer is put to should be offset with sufficient cognitive effects. Any utterance of linguistic material could have an indefinite range of interpretations. To optimise their chances of being understood as intended, speakers alter the salience of linguistically possible interpretations (Wilson & Wharton, 2005). Speakers seeking to alter the salience of particular interpretations must alter the form of their utterance in such a way as to encourage a search for the type of cognitive effects that they intend for recovery.

One way that the form of an utterance can be altered to this end is to manipulate its 'prosodic packaging'. Consider (11):

> (11) [Adapted from personal communication, 07/2015. In this context, Kelly had forgotten that Rosie had been reading a very enjoyable book to her, and became excited that they were going to start reading that book again.]
>
> Rosie: What shall we read?
> Kelly: The BOOK!

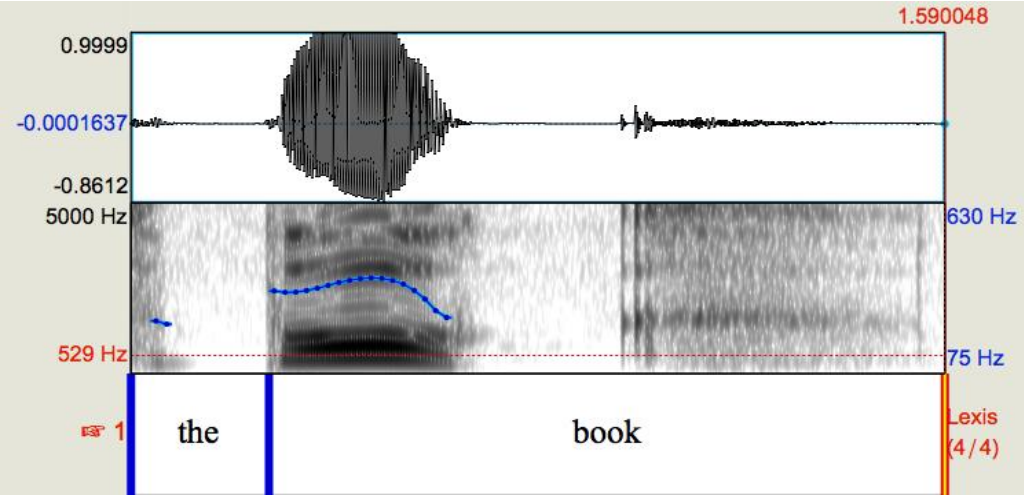

Figure 5 Pitch contour for 'the book'

In figure 5, the nucleus of *the book* falls on *BOOK*. There is a marked pitch excursion, and the nucleus is high in intensity and loudness. Pitch is relatively high. In interpreting *the book*, Rosie has to perform reference assignment. She needs to compute an optimally relevant ad hoc concept BOOK*. There are many books that are manifest to both Rosie and Kelly, so how can Kelly aid Rosie in reference assignment? By making the nucleus *BOOK* very prominent through

high pitch height and a significant pitch excursion, Kelly provides a helping hand to Rosie, prosodically altering the salience of *book*, encouraging her to search her conceptual address for *book* thoroughly and helping to make one of a range of linguistically salient book referents even more salient. This helping hand reduces the likelihood that Rosie will set off down an interpretive path that will not result in the recovery of sufficient cognitive effects to reward her effort. However, this highlighting of *book* does not encode a particular referent; it just *shows* the word *book* to suggest it is worthy of being paid more attention in the identification of explicit content.

Very crucially, the 'prosodic packaging' of *book* in this case is not just responsible for constraining the recovery of a particular explicature. It has a dual function. The prominence of the nucleus (which is in the *least marked* place it can be in this clitic+noun construction) is also increased even further to communicate Kelly's excitement, and it is here where coding can enter the picture. As we saw in chapter two, §2.6, there are different procedures which can be activated in different domains of cognition, including emotion reading, and some of these can be encoded by prosodic signals. It could be considered that the high pitch and pitch excursion on *BOOK* are signals that activate procedures that lead Rosie to compute and calibrate a representation of (an aspect of) Kelly's state of excitement. Moreover, just as you cannot fail to recognise that a true smiler is happy, Rosie cannot fail to recover Kelly's emotional state. The signals encode procedures that force the recovery of certain emotional representations. This is how the 'prosodic packaging' of *book* allows Kelly to both aid Rosie in reference assignment *and* communicate something about her emotional state, and the same sort of explanation is going to apply to *there's a fox, a fox in the garden*. However, I still have to consider what kind of point is served by the speaker repeating the prosodic features of the nuclei on the tokens of *fox*.

The view that prosodic signals encode procedures is not completely unproblematic, however. Wilson and Wharton (2005, p 445) report that Dan Sperber considers that some prosodic behaviours are neither totally natural, nor totally linguistic, but may instead be some kind of culturally stabilised mechanism - such as particular stylised intonation patterns. In the spoken epizeuxis data that I examine, I consider, however, that the natural prosodic signals accompanying the nuclei *are* genuine natural signals which encode procedures in the domain

of emotion reading. This is because the utterances I draw on are typical of spontaneous conversation, and the speakers are not using any stabilised prosodic patterns that one might find associated with the telling of certain formulaic jokes, or with engaging in pantomime audience participation.

In this section, I have shown that prosodic prominence, on one hand, can be employed to make an interpretation more salient in utterance interpretation, but, on the other, it can be (additionally) exploited to activate procedures in the domain of emotion reading. In particular, I showed that a high pitch height might activate a procedure which interacts with the context for the identification of an emotion such as excitement or fear. My suggestion is that a high pitch height (and a significant pitch excursion) that is chiefly intended to communicate an emotional state then raises the prominence of a particular syllable, which is then perceived as a nucleus, or a nucleus is made even more prominent. And, as we have seen, where there is a (very salient) nucleus, it is assumed that there is an intonation phrase. This means that the hearer's *perception* of an intonation phrase follows from her *perception* of a high pitch height and a significant pitch excursion, even if these were chiefly intended by the speaker as a means of activating an emotion-reading procedure. In other words, the perceived nucleus and attendant intonation boundaries can emerge, in some circumstances, from the fact that high pitch and/or a marked pitch excursion have been deployed mainly for the communication of emotion. In such cases, I would say the speaker did not necessarily intend to place specific intonation boundaries in specific places. Boundary placement can be a by-product of where you put the nucleus. What this has to do with the repeated nuclei in *there's a fox, a fox in the garden* is revealed in due course. I now turn to the role of intonation boundaries in utterance interpretation generally.

4.3.4 The general role of intonation boundaries in utterance interpretation

This section of the chapter concerns the interpretation of a kind of repetition – spoken epizeuxis. My interest here is to examine the role that intonation boundaries play in guiding the interpretation of such utterances. To do this, I must first set out what role intonation boundaries are considered to play in the interpretation of any utterance generally – whether or not it contains a repetition of any kind. Intonation boundaries demarcate *chunks* of

linguistic material. Many fields of linguistics are concerned with chunks of particular kinds. Syntax is concerned with the building and parsing of chunks of syntactic structure, and psycholinguistics is concerned with the identification and processing of particular chunks of language, whether they be phonemes, words, or larger phrases. Information Structure looks at the ordering of chunks. Prosodic analysis can be concerned with grouping of chunks. I firstly examine the role that some linguists consider intonation boundaries to play in the interpretation of strings of linguistic material before turning to pertinent findings from the particular field of psycholinguistics.

There exists a continuum of views concerning the role of intonation boundaries in the parsing of strings of linguistic material. At one end of the continuum (see Crystal, 1975; Cruttenden, 1997 for discussion), some consider that intonation boundaries provide solely strict and local cues to scaffold the parsing of syntactic structures. On this view, intonational phrasing should largely exist to overlay syntactic phrasing. Indeed, it is right that some syntactic structures do seem to require the presence of an intonation boundary for their licensing. I think here of some sentence initial adverbials, and of some parentheticals. At the other end of the continuum we find a less strict view: intonation boundaries are fuzzier, both in their perception and in their role, and the chunking of material into intonation groups is considered a gradient phenomenon (Barth-Weingarten, 2011). Bolinger (1972) also proposed that there could be degrees of boundaries which depended on the pitch range in a given intonation group. On this view, the perception and placement of boundaries is fuzzier, and is not limited to overlaying all fixed syntactic units. For example, we can move boundaries to provide a helping hand to a hearer to help them select an intended parse in an ambiguous string. Certainly, the 'strict and local' approach to intonation boundaries has its problems. Boundaries and syntactic chunks do not always overlap (Brown *et al*., 1980; Barth-Weingarten, 2011), and Chafe (1994) has noted that this way of accounting for syntactic structure means that intonation groups and semantic units do not always match up. This strict approach to boundaries would assume that intonation boundaries are always *intentionally* produced, and exist only to scaffold syntactic parsing. This cannot be right. I have made the case that some boundary cues are not deliberately produced, but emerge from a speaker's need to communicate his emotional state through producing highly salient nuclei. What we need is a more general analysis of intonation boundaries which allows us to explain why

speakers can use them to achieve multiple outcomes in communication such as offering a helping hand in syntactic parsing where there is potential ambiguity, or suggesting other processing strategies. A broader, natural account of intonation boundaries can be drawn out of the psycholinguistic literature on the perception and functions of intonation boundaries in processing particular linguistic strings.

There is good evidence in the psycholinguistics literature that intonation boundaries do serve a kind of general chunking or bounding function. For example, Baldwin and Coady (1978, p 373) have shown that intonation boundaries in child reading can trigger a fixation period on the material that is bounded by the breaks. Luo *et al*. (2013) have shown that prosodic boundaries delay the processing of lexical material that comes after the boundary, fitting with Baldwin and Coady's (1978) finding of a fixation period. These studies suggest that intonation boundaries can be used to 'box off' and *highlight* a particular linguistic form (which may or may not have a syntactic motivation). This encourages a hearer or reader to fixate on the linguistic form that has been bounded by the breaks. However, if an intonation boundary is not syntactically motivated, it must be justified in some other way. In the relevance-theoretic framework of this dissertation, it seems that we can say that an intonation boundary serves as a means of delimiting the boundaries of the stimulus which is assessed for relevance. In other words, I consider that intonation boundaries can be used to demarcate *units of relevance* as inputs to particular pragmatic processes. This is very similar to the views taken by Schafer (1997) and Blodgett (2004). Schafer (1997) suggests that the presence of a boundary triggers 'wrap up' of any outstanding semantic and pragmatic processing, and Blodgett (2004) says that a boundary can 'wrap up' any outstanding syntactic parsing. Essentially, in relevance-theoretic terms, it seems that hearers may continue to 'wring out' cognitive effects from what they are processing and their current processing strategy, and then wind this process down in expectation of a new unit of relevance when they perceive boundary cues - including the cue that they have already processed a highly salient nucleus. This provides us with a broader account of intonation boundaries – one in which they can be used for a wide variety of purposes, including, but not limited to assisting the parsing of syntactic structure.

4.3.5 Interpreting *there's a fox, a fox, in the garden*

Figure 6 shows the prosodic analysis for there's a fox, a fox in the garden.

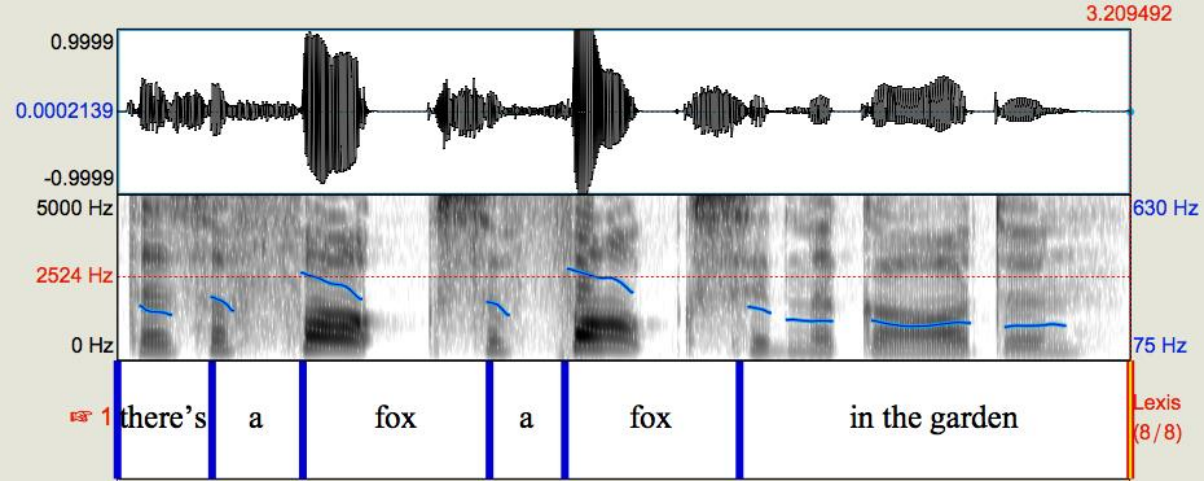

Figure 6 Prosodic analysis for 'there's a fox, a fox in the garden'

I now set out what the hearer would do in interpreting *there's a fox, a fox in the garden*, explaining how different salient aspects of utterance form interact with the context to trigger the recovery of certain effects. In particular, I am really interested in the role the boundaries play, why they might trigger recovery of implicit content, and the point that is served by the speaker ostensively repeating the nuclei.

An important focus of this chapter so far has been on what intonation boundaries are, and how the perception of nuclei can lead to boundary perception. However, in order to consider why the speaker of *a fox, a fox* packaged it prosodically in the way she did, it is first necessary to consider something I've not focussed in this discussion at all: whether or not *a fox, a fox* is a case of *showing*, or at least has an element of *showing* to it. This is quite tricky because, on one hand, it's effectively much of the whole phonological form of an intonation group as well as its prosodic features that have been repeated, so I want to give an account of the repetition as a whole, but, on the other, the unit that has been repeated is a cluster of forms that are interesting in their own right - especially the repeated nuclei. I think that the whole repetition of form of *a fox* can be treated as a case of quite emphatic *showing*. More specifically, I believe it represents a case of *indeterminate showing*, or, number 9 on Sperber and Wilson's (2015) two-dimensional space. A good place to start with the analysis is the repetition of the words

a fox, and the analysis of X-x in chapter three provides insight into how we can approach the adjacent repetition of linguistic material here.

At the end of the last chapter, I explained how onomatopoeic expressions are cases of *showing* and *saying* at the same time, but that the speaker can effectively deliberately and openly *further show* his choice of form by blatantly not choosing a difference utterance form that he could have reasonably been expected to employ in the context. In other words, the speaker exploits and 'plays on' the markedness of the onomatopoeic expression to draw that little bit more attention to it and to suggest that even more non-propositional effects are on offer. The same sort of argument was employed in my analysis of X-x modification. The speaker of an X-x is well aware of the markedness of having a modifier with the same phonological shape as its head, and exploits this to indicate that more processing should be devoted to X in pursuit of more non-propositional effects. In a similar way, it may be that the speaker of *a fox, a fox* is also aware that repeating phonological forms adjacently is somewhat marked. If this is the case, then there is an element of *showing* here. However, if we think about the context in which *a fox, a fox* would be uttered, the markedness of the form is likely not very salient given the fact the attention of the hearer might be focussed on much more salient stimuli in such a context.

My suggestion is that some ostensive repetitions of linguistic form *show* the form of what has already been said before, and lead the hearer to ask why the same form has been repeated, and to reason from this *showing* of form to the first layer of information that the speaker wants to communicate. In contexts which are calm, casual, and without background noise or interruption, or in reading contexts where a) we have stronger expectations that the forms selected by the writer are absolutely intended, and, b) the reader has more time to re-read and reflect on certain features of utterances, such repetitions would likely succeed as cases of *showing* on their own. This is because the speaker or writer could easily exploit the existing markedness of the repeated form without the attention of the hearer being directed towards other salient stimuli. However, this is unlikely to work in the context of *a fox, a fox*. Any *showing* solely involving the repetition of words could be 'drowned out' by other stimuli - it would just be too subtle to be picked up as deliberate display of linguistic form. The hearer here is likely to pay attention to the fox, or where she thinks the fox is, and her processing

resources might be busily computing the implications associated with the presence of a fox outside. She might not really notice the fact that the words *a fox* have been repeated. In this case, the speaker of *a fox, a fox* needs to further *show* her repetition. She does this in a very particular way by forcing a comparison between the 'contents' of the adjacent intonation boundaries, and I explain how shortly.

There is another very particular reason why the speaker needs to *show* her repetition in this case. In chapter one, §1.3.2, I explained that speakers are not always optimally relevant, even if they try to be. I also explained that sometimes we repeat ourselves accidentally, and that not all repetitions are deliberately intended to communicate non-propositional effects. Sometimes speakers are nervous or scared, for example, and might repeat themselves in error, making a mistake in production that leads to an accidental repetition of form. Imagine a context where you see a fox, and you are terrified of them and, perhaps, like I do, you keep chickens, and the presence of a fox in your garden would have serious cognitive effects for you. In such a context, you are likely to be scared. To anyone who is co-present in the context, and to whom it is manifest that you are frightened of foxes and that you keep chickens, the context for interpretation may well include these assumptions. In such a context, it is likely that any repetition of linguistic form you utter would be taken as a production error rather than a deliberate repetition of form designed to engender particular effects. In order to avoid this, the speaker of *a fox, a fox* can manipulate the form of the material in *there's a fox, a fox in the garden* to further draw attention to the repetition, increasing the chance it is recognised as intended. Let us now turn to what the hearer does in interpretation here. Along the way, I show how (perception of) boundary placement leads to the recovery of effects at the implicit level of content.

The speech stream is presented linearly. As such, the hearer encounters | *there's a FOX* | first of all. As soon as she encounters *there's a* she will begin fleshing out a proposition schema in pursuit of an explicature, and will already be narrowing down potential hypotheses about what is being pointed out to her. When she hears *FOX*, what she does is very similar to Rosie in the case of *the book* above. On one hand, the prosodic prominence of *FOX* aids the hearer in formulating an optimally relevant ad hoc concept FOX* during reference assignment. The hearer is encouraged by this highlighting to spend time searching through her conceptual

entry for *fox* more deeply and thoroughly than she would have without this prosodic encouragement. The extra processing effort invested pays off in the form of a finely nuanced ad hoc concept FOX* on one hand, and, given the context, the construction of this concept will give rise to a wide range of weakly communicated implicatures, likely relating to the consequences of the fox being present in the garden for the speaker and for her unfortunate chickens. On the other hand, the speaker can be said to overlay *FOX* with very salient or, as it is an unmarked site for the nucleus, extra prominent prosodic signals which, at the same time, are intended to communicate, via emotion-reading procedures, the scared emotional state of the speaker. These effects fall on the implicit side of content. The prominence of the nucleus, along with the recognition that *there's a fox* could form a syntactically and semantically complete unit, creates the expectation of the perception of an intonation boundary. This (expected) boundary signals that a unit of relevance is coming to a close, and that the hearer should 'wrap up' any interpretive processes she carries out in pursuit of cognitive effects. At this point, the hearer has 'collected' an explicature that satisfies her expectation of relevance in the context, and a series of weakly communicated implicatures about the consequences of a fox in the garden, and a representation of the emotional state of the speaker. These are sufficient effects to offset the processing effort she has invested, and the suggestion of a boundary leads the hearer to subconsciously consider what relevant stimuli may be coming next.

At this point, there are probably only three things that can happen next.

(1) Nothing. The speaker ends her turn and stops speaking.

(2) The speaker continues her turn. There follows an intonation phrase that is different in terms of form and content to the one that just ended e.g. *There's a fox.*| *Look outside!*

(3) The speaker's turn continues. However, the next intonation phrase is an exact repetition of the first, or at least contains a repetition of a noticeable portion of it.

I would say, all other things being equal, that option (3) would, generally, be the least expected. Or, if repetitions were to be expected in this context, there is, as I have said, a possibility that any intentional stylistic repetitions might be misconstrued as production errors in a context where it is manifest that the speaker is nervous and scared. The speaker

needs to do *something* to ensure the repetition of phonological form is recognised as intended, and that any additional non-propositional effects at the level of implicit content are 'squeezed out' of the interpretation. The way the speaker achieves this on this occasion is to aim to produce a second nucleus on the second *FOX* that it is judged to be sufficiently similar to the first *FOX* nucleus as to count as being a repetition (see the discussion of the rings in chapter one, §1.2), and one which has a particular consequence.

When the hearer encounters the second nucleus, what is particularly salient to her is that its prosodic features have been deliberately repeated. This is unusual. It is marked. It is attention-attracting. This leads the hearer to subconsciously ask what point has been served by this. But, first, if the hearer has seen fit to recognise that the *FOX* nuclei are similar, it means that, on some level, she must have compared the nuclei to one another. How else could she judge that the second was an ostensive attempt at a repetition of the first? This comparison of nuclei has a crucial side effect. Recall that Sperber and Wilson (1995, pp. 202-217) have said that nucleus placement suggests which constituent the relevance of an utterance lies in. However, the size of the constituent where the relevance lies has to be pragmatically inferred (see Sperber & Wilson, 1995, pp. 202-217 on focus effects) - nucleus placement could draw attention to the single expression that is the nucleus, or to, say, something like an entire complex subject such as [the MAN whom I saw last week], which contains a nucleus that is itself an entire constituent that could be in focus (i.e, the locus of relevance). When the hearer of *a fox, a fox* sees fit to compare the *FOX* nuclei, she has to pragmatically infer what the placement of the nucleus tells her about 'where the relevance lies'. I suggest that, in this endeavour, it is optimally relevant for the hearer, during this process, to also compare the forms of the adjacent intonation groups that contain the near-identical nuclei. And, when the hearer carries out the comparison of the adjacent intonation groups, what does she notice? The repetition of phonological form that needed to be made more salient in the first place: *a fox, a fox*. In this way, the speaker of *a fox, a fox* further *shows* something she was subtly *showing* already.

I should point out that the repetition of the nucleus has an additional effect - it serves to encourage the hearer to recalibrate her representation of the speaker's emotional state, yielding additional non-propositional effects on the implicit side of communication. However,

this is not so relevant to the discussion in this chapter, as these won't be poetic effects. I address the calibration of emotional states when discussing repeated interjections, expletives and emoji in chapter five, §5.6. What is most of interest is why and how additional weakly communicated implicatures are recovered in this particular case, and what this has to do with intonation boundaries.

As I said above, intonation boundaries can serve to mark the end of a unit of relevance. In other words, the suggestion is that a hearer has fleshed out any relevant explicature, and has recovered sufficient implicit content, and can start 'wrapping up' any processing strategies she undertook in pursuit of optimally relevant cognitive effects in anticipation of processing what might be coming next. In the case of *a fox, a fox*, the second *a fox* is in another intonation group, or unit of relevance. The boundaries delimiting the second intonation group suggest that whatever it contains is also relevant in some way, and that effects are to be sought from processing it. The boundaries force the hearer to concentrate on the second *a fox*; they pen her in to concentrate on processing that form again for relevance. Since *a fox* has been repeated, and as Sperber and Wilson's (1995, pp. 219-222) account suggests, the repetition of *a fox* can only be justified by the assumption that the hearer should 'go back' to the context created when *a fox* was first interpreted, and expand it to collect additional weak implicatures on her own responsibility. What else can she do? The repetition of linguistic form hems the hearer in to the same conceptual region or number of regions that were activated when she processed *a fox* the first time around. To seek relevance elsewhere would not be offset by sufficient cognitive effects in the context. The optimally relevant interpretation is to 'go back' to the first interpretation (and its context) and expand it. The boundaries in this instance of epizeuxis serve to focus further attention on the repetition, suggest it should be re-processed for relevance, and force the hearer to re-examine and build on her initial interpretation. It seems, then, in spoken epizeuxis, that the general chunking function of intonation boundaries interacts with the repeated form and the context to further *show* what has been repeated, and to ensure the hearer fixates on the repeated material and treats it as relevant.

As I have noted before, it appears that there is some sort of connection between intonation boundaries and the recovery of effects at the implicit level of content. I do not actually think there is something special about intonation boundaries that guarantees the level at which

effects are recovered, and there is a simple explanation for what is happening in many cases of epizeuxis. A hearer can always revisit and revise their explicatures. We often do. We re-adjust ad-hoc concepts in light of changing contexts, adding or deleting as new information comes to light. However, as this thesis progresses, it is becoming clearer that we can use various repetitions of form to communicate (aspects of) interpretations that we should struggle to communicate solely via the linguistic semantics of an utterance plus some pragmatic enrichment. Certainly, the hearer of *a fox, a fox* may judge it optimally relevant to 'go back' and adjust her concept FOX* in light of the evolving context, and this may be intended in some contexts. However, in this case, all other things being equal, I think the hearer's first shot at FOX* in this context is probably sufficient to satisfy her expectation of relevance, and so it isn't a main intention of the speaker to have her go back and re-adjust that concept on this occasion. This means that if the point of such repetitions isn't to scaffold the recovery of explicit content, the only thing you have left to do as a hearer is expend more effort on recovering implicit content as part of 'wrap up'. I do think that speakers who ostensively repeat what they have already uttered make manifest the assumption that they chose not to utter something else, or, perhaps, that they could not utter something else. Speakers who make such assumptions manifest may be taken as intending to communicate something that is vague, or nebulous, and possibly vast in terms of the reach of modification to a hearer's cognitive environment, i.e., the type of thing that falls on the implicit side of communication and, though I am loathe to use the metaphor, the type of thing we struggle 'to put into words'. The intonation boundaries in *there's a fox, a fox in the garden* provide a chunking function that indicate how these sort of non-propositional effects should be recovered, and when their recovery should start, stop, and re-start. Thus, the repetition and the boundaries work together in interaction with the context to scaffold the recovery of weakly communicated implicatures, which is in line with Sperber and Wilson's (1995) weak implicature account of repetition, even though they did not explicitly mention the boundaries that are definitely part of epizeuxis. If this analysis is correct, then what we have here is, as I suggested above, a case of *indeterminate showing*, as the import of this particular repetition is very vague and cannot be propositionally paraphrased at all.

Is the repetition in *a fox, a fox* emphatic in any sense? Sperber and Wilson's (1995) *weak implicature account* of repetition predicts there should be some emphatic effects as a result

of processing epizeuxis. I have been suggesting that it is communicative speaker behaviour that should be thought of as emphatic. Earlier, I said that the repetition of phonological form in *a fox, a fox* on its own in a context where the speaker is manifestly scared might not be salient enough to be noticed without additional support from the speaker. What the speaker did to combat this was produce prominent prosodic signals on the *FOX* nuclei, and exploit the placement of boundaries to further *show* the phonological form she repeated. The speaker's *showing* constituted a cluster of highly ostensive behaviours which provided evidence for what the speaker wanted to communicate. As such, my suggestion that emphasis concerns the ostensiveness of speaker behaviour is further supported. Emphasis seems, furthermore, to be related to the level of ostensiveness of behaviours involved in *showing* in ostensive-inferential communication. If the effects of such examples of epizeuxis are themselves not emphatic, are they intensifying in any way? Sperber and Wilson (1995, p 221) suggest that repetitions like *a fox, a fox* yield an increase in effects by getting the hearing to expand the context and add more implicatures to it, and by investing extra processing effort in searching 'deeper' for these effects. This seems to be correct based on the analysis above. However, I am not sure at this stage that it is the effects of epizeuxis which are intensifying, and, moreover, Sperber and Wilson's comments do seem to fit better with my earlier suggestion that intensification, whatever it is, is a processing phenomenon. All we can say for this specific example is that the speaker must *expend* more processing effort in *extending* the context for interpretation, and bear this observation in mind going forward.

4.4 The difference between epizeuxis and apposition

In §4.2.3 of this chapter, I suggested that an advantage of Sperber and Wilson's *weak implicature account* of repetition was that it could potentially be applied to other phenomena involved in the communication of similar effects to repetition. I said that, in particular, it could be applied to appositions, e.g., *he was depressed, flattened*. In personal communication and in her 2008 paper on apposition and affective communication, Diane Blakemore has drawn a link between repetition and apposition, and has said that the way appositions 'communicate an impression of emphasis or intensification' can be compared with the effects communicated by repetitions (Blakemore, 2008, pp. 37-38). She (2008) notes that Sperber and Wilson's (1995) *weak implicature account* of poetic effects can be applied to explain some of the effects of apposition. While the conflation over emphasis and intensification

mentioned in chapter one is present here, Blakemore is correct in this. However, Blakemore's (2008) account of apposition can be enriched further by explicitly adding in a description of the role of intonation boundaries because, in spoken cases of apposition, there will be an intonation boundary between the first part of the apposition, and the second.

Let us return to apposition examples (4a) and (4b), which Blakemore (2008) called cases of 'intensification'.

(4a) Chris was **depressed,** | **flattened**.

(4b) The evening's been **spoilt,** | **ruined**.

I will not provide a full account of apposition here, since, as Blakemore (2008) has shown, there seems to be more than one type from the point of view of the way the effects are recovered. Of the 'intensification' cases of apposition, Blakemore (2008) says that the effects are recovered as follows. In the case of (4b), for example, the interpretation of the second segment of the apposition is some kind of amplification of the interpretation of the first. Her question is then how two segments which are similar in interpretation can differ in 'intensity'. Blakemore (2008) says that the first segment of (4b) encourages the construction of an ad hoc concept SPOIL* which, through its construction using particular assumptions, gives rise to its own implicatures. For example, *spoil* is usually predicated of cases of minor damage, e.g., getting water on a book, or burning the edge of a cupcake. The hearer can recover, as part of constructing this concept, that the evening was not as successful as hoped, and a range of other weak implicatures on her own responsibility. What Blakemore (2008) then says is that the use of *ruin* leads the hearer to access a different range of assumptions - although she does not explicitly talk about ad hoc concept construction at this point. Blakemore (2008) suggests that the hearer is then encouraged to compare the cognitive effects of the latter segment with the former, and says that the relevance lies in the properties that RUIN* has which SPOIL* does not share. I think that this is broadly correct. However, we do not have an account of what the intonation boundary does in such appositions, and, very crucially, if there is a sense of intensification here, I think it emerges from ad hoc concept construction rather than from the recovery of implicit content, even if the *weak implicature account* can correctly be used to suggest how some of the stylistic effects of these appositions are recovered.

Consider the interpretation process of (4b) for the hearer. When she hears *the evening's been spoilt*, she has no idea that an apposite segment is coming - and these are fairly marked, so it's unlikely to be expected, all other things being equal. The syntax and the prosody of the utterance may lead the hearer to a fairly justified conclusion that all she will need to do is compute an ad hoc concept SPOILT*, and collect any optimally relevant weakly communicated implicatures that result from adjusting that concept to apply to evenings rather than to spoilt books and cakes. The nucleus will fall on *spoilt*, the segment is syntactically complete and, so, the hearer has boundary cues that lead her to 'wrap up' ad hoc concept computation, and the derivation of implicatures. Then, however, the hearer is presented with an apposite modifier. There is only an intonation boundary between *spoilt* and *ruined*. There is no conjunction, or any other linguistic evidence for how *ruined* is to be interpreted - there is just the boundary. Just as a hearer interpreting epizeuxis judges it optimally relevant to compare the contents of two adjacent intonation groups; a hearer of an apposition would judge that she must do the same. She has to find a context where the contents of both boundaries can be integrated. Indeed, it is the only thing she can do given there is no encouragement to do otherwise in the form or the linguistic semantics of the utterance. Of course, this will lead to the recovery of more weak implicatures once the hearer is forced by the second boundary to consider how the apposition can be integrated into the context. This is the sense in which the *weak implicature account* can be applied to apposition as well. However, it is at this point where my account of the apposition in (4b) differs slightly from Blakemore's.

Ruined sits in its own intonation group. Blakemore's account seems to suggest (though she does not explicitly say it) that the hearer would compute a concept RUINED*, which would then lead to a different range of implicatures being recovered. I analyse what the hearer does as follows: the hearer is already carrying out pragmatic work in the conceptual region unlocked by *spoilt*, and recovers SPOILT* which satisfies her expectation of relevance in the context. She then encounters *ruined*. I don't think the hearer constructs a separate concept RUINED* at this point. It is likely that the optimally relevant and more economical strategy is to *expend* effort on *further modulating* the SPOILT* concept in a particular way. The conceptual regions for *spoil* and *ruin* likely overlap in that they probably share a high number of connections. The fact that *ruined* is in its own intonation boundary forces the hearer to

fixate on it and, in the absence of other cues to do otherwise, modulate her concept SPOILT* based on information that would normally be readily accessed through the conceptual address for *ruin*, e.g., assumptions about complete destruction rather than minor damage. We might be tempted to think of this as a case of broadening (see chapter two, §2.5.4). However, we could also think of it as a case of narrowing. It seems to me, and Blakemore (2008) said it of some of the first segments in her appositions, that the concept communicated by the first segment is quite a general one. What I think is that the concept expressed by the first segment in some appositions starts off quite general and is then narrowed to become quite a particular occasion-specific concept that brings with it a set of particular weakly communicated implicatures that the hearer is responsible for recovering and which would not be recovered by computing the concept that only the first segment would communicate. Since it is considered that boundaries can be effortful, and can slow down processing (Pilkington, 2000), and since the concept communicated by *spoilt, ruined* requires the activation of more than one conceptual address, this puts the hearer to more effort than if there were no apposition. Of course, this effort is offset by the recovery of a finely-tuned concept, and a range of particular poetic effects, so the use of apposition has a justification.

If the slightly amended analysis I have just provided is along the right lines, then the speaker did recover extra and different effects on the implicit side of communication. However, what is of interest is that I am not sure this is the source of any sense of intensification for (4b). If we look at what the hearer did when processing the apposition, she was encouraged to *expend* more processing effort on f*urther narrowing* a concept that she had already adjusted. This was in pursuit of the identification of explicit utterance content. Perhaps the sense of intensification in apposition arises from the hearer expending more effort during narrowing in ad hoc concept computation. If this analysis is correct, this means that intonation boundaries are unlikely guarantee that effects are recovered at either the explicit or implicit level, as my initial hypothesis suggested earlier in the chapter. The data discussed in the remainder of chapter four, and in chapter five allow me to comment more concretely on the role of intonation boundaries in communication in my concluding chapter. In any case, the way intensification obtains in the identification of explicit content in certain appositions is very similar to how it obtains in some cases of the interpretation of material repeated in the

same intonation group, as discussed in chapter five - through narrowing in pursuit of explicit content.

4.5 'Long distance' repetition: intensification in recovery of implicit content

4.5.1 The problem of 'distance'

The appositions and cases of epizeuxis addressed above concerned material in immediately adjacent intonation groups. That is to say, there was no linguistic material, or 'distance' in terms of space-on-the-page or time elapsed between an original and its repeated counterpart, or the first segment of an apposition, and the second. However, there are repetitions which have stylistic effects, but which are not quite as adjacent to their originals as in epizeuxis. Consider a slightly amended version of *there's a fox, a fox in the garden*.

(12) There's **a FOX** | in the GARden |, **a FOX**.

(12) also has stylistic effects. However, the repetition occurs non-adjacently to the original, even if it is relatively 'near' in some sense. The key questions raised by the amended utterance in (12) are as follows: What is the difference in interpretation between *a fox, a fox* and the repetition in (12)? And what aspect of utterance form can we pin any difference in interpretation down to? Relevance Theory is able to suggest an answer to these questions and these answers offer a way into analysing the nature of other 'types' of repetition that occur over greater stretches of 'distance'. The discussion in this part of the chapter allows me to define a cognitively motivated notion of 'distance' which should be useful to anyone attempting to describe the processing and interpretation of repetition within pragmatic stylistics.

It is actually very hard for us to notice, let alone articulate, any difference in interpretation between *there's a fox, a fox in the garden* and *there's a fox in the garden, a fox*. In terms of explicit content, any differences need not be particularly salient or of great consequence. Both utterances do have non-propositional effects, however. On one hand, there are poetic non-propositional effects arising from the weak communication of implicatures. On the other hand, there is the matter of the communication of an emotional state through prosodic signals of nuclei. The difference in form between the examples is not great. We would struggle to talk about their individual effects anyway owing to ineffability, and the nebulousness of

what is communicated, so it would be doubly hard to try and describe exactly what differs in terms of the effects of *a fox, a fox* and (12). The only clue that we have that there ought to be a difference in effects is the slight difference of form - the second *a fox* is in a different place, and so there is, in some sense, more 'distance' between the original and its repetition. However, I am not sure that we have a good notion of 'distance' to work with in pragmatic stylistics. While we might have a pre-theoretic, intuitive notion of distance, this will not be adequate for an explanatory account of non-adjacent repetitions.

In research on coherence, there is discussion of distance between segments involved in particular coherence relations, and this relates to the notions of local coherence and global coherence that were mentioned in chapter one of this thesis. However, the discussion is cashed out using the terms adjacency and non-adjacency (or less adjacency) (Speyer & Fetzer, 2014). Adjacency can be thought of as a structural relationship, but also as a semantic and pragmatic connection between segments that are close to one another (Spyer & Fetzer, 2014, p 98). The more adjacent two segments considered to stand in a coherence relation, the more local the coherence. The less adjacent two segments are, the more they would be considered to contribute to global coherence (Lenk, 1998). It has been suggested that a theory of text coherence should be able to account for a notion of distance, as it appears, for example, in the case of anaphora resolution, that anaphoric noun phrases can 'find' an antecedent 'quite far away' in a discourse (Stede, 2012, p 123). It has been suggested that the notion of distance from the beginning of a text or discourse to its end might come into play (Stede, 2012, p 123), and 'distance' is thought to contribute to the hierarchical nature of Mann and Thompson-like structural relationships in a text (Liebert *et al*., 1997; Speyer & Fetzer, 2014). It is not clear that 'distance' should mean just structural distance. I am not sure that humans would actively compute the mathematical length of a text and keep track of the length of text expired between multiple segments and use the results of computations to diagnose different types of coherence relation while trying to make sense of a discourse or text on-the-fly. Speyer and Fetzer (2014), mentioning Sperber and Wilson (1995), do suggest, however, that the notion of accessibility may be important in explaining elements of adjacency. Perhaps, then, for an understanding of 'distance' and any role it plays in the recognition and processing of particular repetitions, we need to move away from any purely mathematical calculations and representations of structural 'distance' in terms of 'length' in the mind of a hearer and reader,

and look for an explanation of 'distance' that is couched in terms of what the audience is doing in terms of interpretive processes and processing effort. However, rather than thinking about accessibility, I think it might be more beneficial to think about how a notion of 'distance' might relate to activation levels of interpretations and the mental addresses or spaces involved in their recovery. (That activation is important in the understanding of repetition is not a new idea. Bakker (1963) has also proposed this, and I turn to his account in a short while.)

Relevance theory can offer a cognitive basis for a theoretical notion of 'distance'. We can treat 'distance' in terms of processing effort, where this is understood as the effort entailed by processing the linguistic forms in question plus any inferential work involved in accessing contextual assumptions during the derivation of explicit and implicit content. The notion of activation is important here. Essentially, with repetitions over more or less 'distance', what happens is as follows: the original is processed, and its interpretation is recovered. At that time, the information used in that interpretation is highly activated, and the interpretation recovered is highly activated. If, for some reason, it is optimally relevant in the immediate context to keep working on that interpretation, the speaker can do so. However, as most of our utterances do differ in form from one another, and because contexts evolve and need updating, the chances are that your attention is drawn from the interpretation you just recovered, and the form you just processed, and you start to attend to new stimuli and new contextual assumptions. As such, the activation levels of the interpretation you just recovered and the activation levels of any conceptual regions used in that interpretation will drop. Your attentional and processing resources are diverted elsewhere. Suppose, several minutes later, you hear an exact repetition of form that you consider to be ostensive, and so you decide to attend to it to recover specific effects that you have reason to believe are on offer. I have already suggested that, with some repetitions, the repetition forces an audience to 'go back' and extend and re-work an interpretation. However, this is extremely effortful. You have expended a lot of effort in the intermediate time processing other stimuli and attending to their interpretations. As a result, you have to re-activate mental regions that have become de-activated, and likely to a degree that is greater than the first time around (as you are searching for different or extra effects), and you have to re-activate your original interpretation in order to add to it or amend it in some useful way. The more that your

attention has been diverted elsewhere to process other stimuli, the more difficult it will be to 'get back' to your original interpretation and increase those activation levels. If we think of 'distance' in terms of processing effort expended in this type of interpretive endeavour, it might help us to understand why there is a feeling that repetitions that are physically right next to each other seem much more noticeable than ones that do occur with relatively more linguistic material in between. It also allows us to explain why the effects of *a fox, a fox* and (12) should be very subtly different, even if it is hard for us to conceptualise this.

In (12), there is material in an intonation group intervening between the original *a fox* and its repetition. We can predict that whatever the hearer of (12) recovers in terms of non-propositional poetic-type effects and an emotional representation of the hearer, these representations (and the mental regions involved in their recovery) will decline in activation as the hearer attends to the new material in the next intonation group. The hearer then has to 'go back' and re-activate everything again, and to a greater degree, and with an expectation that she should look for even more effects, and all of this is quite effortful. In fact, in (12), compared to *a fox, a fox*, we would have to say that the communication is slightly weaker, and the hearer has more responsibility for recognising the repetition, and for recovering extra weak implicatures from processing it. In *a fox, a fox*, although there is an intonation boundary present between repeated segments, there is much less 'distance' than in (12) in terms of diverting processing power elsewhere and then diverting it back to reactivate regions and interpretations. Due to the recency of the first *a fox*, there's much less effort expended on 'going back' to the first interpretation than in (12).

If a person's processing resources are taken up with processing a lot of intervening material or stimuli between an original and its repetition(s), the prediction is that where this occurs, stylistic repetitions should be relatively less ostensive, and thus harder to recognise than cases like *a fox, a fox*, and even (12). I think that this is the case. However, as should be suggested by chapter one of this thesis, just because a piece of text has been repeated 'at distance', it may not have a place in this thesis if it is not also intended to be recognised as ostensively communicating non-propositional effects. It is now time to turn to one of the main questions that this thesis addresses: how is it we know that a repetition has been ostensively produced for the communication of stylistic effects? To answer this question, I first examine 'long distance' repetitions that fall outside the scope of this thesis.

4.5.2 'Long distance' repetitions which fall outside the scope of the thesis

Whether written or spoken, it is clear that repetitions exist that have a quite a lot of intervening material to process, and where the context has evolved quite significantly in the time since the original was first produced. Consider (13) and (14):

> (13) [From 'Miranda', BBC1, 01/2014]
>
> Miranda at 3 minutes 34 seconds: **Plunge** the loo before I go…expunge the bath…any activities ending in –unge should be banned…I am on an –unge protest. Although **PLUNGE** is a lovely word…**plunge**…(to camera) **PLUNGE**!
>
> Miranda at 10 minutes 21 seconds: Ooh! Crack. Cheeky! Crack! **PLUNGE**!...**Plunge** m' crack.
>
> (14) [From '300: Rise of an Empire' (2014) (142 minutes' running length)]
>
> Queen Gorgo **[at the start of the film]**: The Oracle's words stand as a warning. A prophecy. Sparta will fall. All of Greece will fall. And Persian fire will reduce to cinder. For Athens is a pile of stone and wood and cloth and dust. And, as dust, will vanish into the wind. Only the Athenians themselves exist, and the fate of the world hangs on their every syllable. Only the Athenians exist and only stout wooden ships can save them. Wooden ships, and a tidal wave of heroes' blood.
>
> Queen Gorgo **[in the closing scene of the film]**: The same monologue is repeated verbatim.

Both (13) and (14) are clearly associated with the ostensive communication of non-propositional effects, and their repetitions occur after the processing of substantial linguistic material (and other contextual elements) following the production of the original material. *Miranda* is a scripted television sitcom. As such, it is part of the context for interpretation that any elements of form found in the sitcom are meant to be there, as the script will have been revised and edited by the writers. The audience is therefore entitled to expect, perhaps to a slightly higher degree than in spontaneous discourse, that any repetitions they notice are there intentionally. The repetition of *plunge* in (13) is associated with humorous effects concerning the onomatopoeic word, and amusing sexual or toilet-humour connotations this low frequency word might have. The low level of frequency of *plunge* and its onomatopoeic

nature arguably contribute to its ostensiveness. Moreover, Miranda prosodically highlights *plunge* more than once, which means that the audience might recover a representation of her emotional state (which is recalibrated each time), as well as recovering a wide range of weak implicatures. In (14), the repetition of Queen Gorgo's monologue involves the exact reproduction of a fairly long stretch of discourse, and after more than two hours have elapsed. The effect of this repetition is to provoke the reactivation of feelings and contextual assumptions recovered by the audience at the start of the film, and to have those feelings and assumptions re-processed and revised in light of the evolved context of the whole film. What makes the repetition in (14) ostensive is the amount of linguistic material that is repeated verbatim. It is not just one word, or one line, but several lines consecutively. The repeated material in (13) and (14) are reasonably ostensive in that they are both marked - the former because *plunge* is low frequency and onomatopoeic, and the latter because a relatively large chunk of linguistic material is repeated. This surely aids the noticeability of these repetitions. It is not clear that the behaviour of *both* the speakers here is emphatic, however. The repetitions are certainly ostensive enough to be picked up by an audience, but I am not sure we would call Queen Gorgo emphatic. Miranda we might, however, and I address this when this repetition is analysed in more detail later. However, I do want to say that these repetitions lead to some intensification during processing involved in the recovery of implicit content, as I shall show. This suggests that there may be a disconnect between emphasis and intensification. While I think emphatic speaker behaviour during *showing* would lead to some kind of intensification in processing in the recovery of certain effects, it is not clear that every instance of intensification is a result of emphatic speaker behaviour during *showing*. The interpretation of some repetitions 'at distance' may be a case of this disconnect.

The study of the repetitions found in (13) and (14) obviously falls under the remit of this thesis. However, there are repetitions which also occur 'at some distance' but which we would not want to include in this study. Consider (15) and (16):

(15) [From Tannen (1987, p. 591)]

D: Do you read?

P: Do I **read**?

D: Do you read things just for fun?

P: Yeah. Right now I'm **reading** *Norma Jean the Termite Queen*.

(16) [Taken from *English Vocabulary in Use* (Redman, 1997, pp. 26-27).]

Page 26, line 1: A **compound noun** is formed from two nouns…

Page 26, line 3: **Compound nouns** are usually written as two words…

Page 27, line 1: Find **compound nouns** on the opposite page…

Page 27, line 2: Complete these sentences with suitable **compound nouns**…

Page 27, line 7: Try creating your own **compound nouns**…

In the case of (15), we might want to say that P repeats *read(ing)* at some distance. However, this cannot be analysed as repetition that is intended to produce a particular stylistic effect. It is more like a case of necessary re-use. *Read* is a high-frequency lexeme for which we have few synonyms. P is simply re-using a word to express a concept that he needs to continue the conversation, as we saw in §1.3.3 of chapter one. In (16), the repetition of *compound noun(s)* is definitely deliberate, as the data is taken from an EFL textbook, and these are carefully produced and edited with specific purposes in mind. However, these purposes are purely pedagogic and not stylistic. Some of the repetitions of *compound noun* are simply there to help the reader work out what she should do in the task. Some of them are simply there to help the reader assimilate new information into her conceptual entry for *compound noun* as she learns facts about new grammatical terms.

The data in (15) and (16) raise an interesting problem. From the point of view of (presentation of) linguistic form, there is absolutely no difference in the way that these non-stylistic repetitions look when compared to cases that seem quite clearly intended to communicate non-propositional effects, such as (13) and (14) above. Chunks of varying sizes of linguistic material can be repeated over more or less distance in both instances. Why is it that (15) and (16) are judged not to be intended to trigger the recovery of stylistic effects while (13) and (14) are? In fact, let's take this question back a step further, and ask, first of all, how it might be that someone notices that something has been deliberately repeated in the first place before they even ask themselves whether they should recover any stylistic effects.

4.5.3 Repetition and redundancy

In 1948, Claude Elwood Shannon published an article: *A Mathematical Theory of Communication* while he was working for a telecommunications company. He sought to model how communication worked. As part of his model, he noted that there are environmental elements which can interfere with the clarity of transmission of a given message. These elements are known as 'noise'. Shannon (1948) proposed that noise could be combatted if redundancy is introduced into the system. Redundancy can be thought of as elements in a system which perform the same function as one another, increasing the chance that the message is understood despite the presence of any noise in the system. If one element is not successfully understood, that aspect of the message will be taken care of by another redundant element. Natural language grammars clearly also have the potential for redundancy. An example of redundancy in the grammar of English concerns the third person singular ending *–s*, and the third person pronouns. The third person pronouns communicate that the verb should be understood as expressing the third person, and the present tense indicative. Even if the *–s* were not phonologically realised, the hearer could still access what its semantics encodes, as it can be recovered by interpreting the pronoun. This type of redundancy is known as grammatical redundancy (Wit & Gillette, 1999). As this type of redundancy seems concerned with coded elements being 'doubled up' to ensure that decoding is successful, I will not pursue grammatical redundancy here. However, the illustration is useful for elucidating the concept of redundancy. Part of the task of explaining why deliberate and stylistic repetitions are recognised and processed as intended involves explaining why we don't treat them as redundant.

Let us turn to what Wit and Gillette (1999) have to say about what they call *contextual redundancy*. They (1999, p 12) note that linguistic redundancy does a lot more than simply improve comprehensibility. Looking at their account, we find a starting point for explaining why deliberate stylistic repetitions aren't considered redundant, whether this is understood from a theoretical perspective, or a layman's understanding of the term (i.e., 'useless', or 'superfluous'). According to Wit and Gillette (1999, p 9):

> 'Contextual redundancy is the repetition of information that is, in a grammatical sense, nonobligatory. This repetition consists of the reproduction of identical elements of information or of elements that are only apparently identical. Contextual

redundancy is not systematically generated by grammatical rules, although nongrammatical circumstances may suggest or require its use. Such circumstances include socio linguistic and psycho linguistic factors.'

Wit and Gillette (1999, pp. 9-11) propose four types of contextual redundancy, though they note that these categories 'do not seem to be mutually exclusive':

(1) **Identical or Synonymous Repetition** – these share the form of the epizeuxis cases which I have been discussing (but also seem to include cases of apposition), e.g., *the green, green grass of home* (Wit & Gillette, 1999, p 10).

(2) **Isolating, Salient Repetition** – These include cases where the semantics of one lexical item includes or guarantees part of the semantics of another lexical item, e.g., *the salty sea* > the sea is salty anyway. (Wit & Gillette, 1999, p 10)

(3) **Contrasting Repetition** – Occurs when words which are already contrasted in a sense (e.g., antonyms) are additionally contrasted through some kind of emphasis, e.g., Spanish ***A mi*** *me gusta el cafe y* ***a ti*** *no te gusta* (literally: *to me to me enjoy the coffee and to you not to you enjoy*) (Wit & Gillette, 1999, p 10).

(4) **Distinguishing, Differentiating Repetition** – '[t]he disambiguation of a word or expression in an unambiguous context with another word or expression that is considered non-repetitive…' (Wit & Gillette, 1999, p. 11), e.g., Birmingham, England (where there are multiple possible referents for Birmingham, e.g., the Birmingham in Alabama, USA).

As well as proposing a taxonomy of four 'types' of contextual redundancy, Wit and Gillette (1999, p 12) propose a taxonomy of six major functions for contextual redundancy. These are as follows:

(1) **Enhancing comprehensibility** – repetitions produced to combat problems with the 'sender' or the 'receiver', or any environmental factors hindering communication.

(2) **Resolving ambiguity** - to increase precision of expression.

(3) **Isolating a feature** – to focus the audience on a salient feature, e.g., 'salty sea' focuses attention on the feature 'salty'.

(4) **Contrasting elements** – where, for example, additional stress is employed to achieve (additional) juxtaposition of elements which already are contrasted in some sense. Noted as *emphasising* or *intensifying* single elements at times.

(5) **Emphasising or intensifying** – the authors talk of pronoun redundancy from various languages even though they have said they do not want to talk about grammatical redundancy. They do, however, go on to give epizeuxis-like examples and note that these may engender poetic effects even though they propose the separate category below.

(6) **Creating 'poetic' effect** – Wit and Gillette (1999, p 15) say '[t]he final function of creating poetic effect is, in some sense, a rest category. It encapsulates all uses of redundancy with no clear semantic purpose, but with an intention to shock, to please, to horrify, to move…'. They go on to say these repetitions often have an intensifying or emphatic effect, and note that these repetitions are permitted to occur at the paragraph level, not just at sentence level (Wit & Gillette, 1999, p 15).

This account of non-obligatory repetitions offers some useful insights. For one, it links some repetitions, correctly, with senses of intensification and emphasis. For another, the authors recognise that non-obligatory repetitions can have a variety of functions. Moreover, the authors correctly note that some repetitions are present for reasons other than grammatical necessity. One of the reasons given is 'psycho linguistic'. Here, I take the authors to perhaps mean that some repetitions exist to have a particular effect on the interpretation process, i.e., there is a cognitive motivation, which I think is correct. Finally, we have an account of repetitions of form which explicitly says that repetition can occur at above the sentence level, and takes into account the fact that we must somehow be able to identify repetitions over more or less 'distance'. However, there are also a number of serious issues in Wit and Gillette's (1999) account of contextual redundancy which must be resolved in order to provide a useful account of how non-adjacent repetitions are understood – especially those which occur over some considerable 'distance'.

The authors acknowledge that there is overlap between some of the categories in their taxonomy. If a clear cut distinction cannot be made between categories in a taxonomy, this brings into question its usefulness as a theoretical tool, and it should be amended or

abandoned accordingly. Furthermore, it seems the authors seek to link particular forms with certain effects. That is to say, if a so-called 'isolating repetition' is present, the audience will *always* focus on the feature that it is guaranteed to isolate, and that because of the linguistic semantics - this would be the feature of saltiness in *the salty sea*, for example. However, clearly, what aspects of encyclopaedic knowledge are accessed in ad hoc concept construction will depend on the context, which includes assumptions about the communicator's intentions - this kind of repetition cannot guarantee what would be focussed on in interpretation or to what degree. It also seems that Wit and Gillette (1999) fall prey to some of the conflation regarding emphasis and intensification that I discussed in chapter one, §1.6. On one hand, they say that 'poetic' linguistic redundancy can result in emphasis or intensification and associate these repetitions with a desire to move people or stir up some emotional response, which suggests they consider emphasis and intensification to be effects. On the other hand, however, they seem to suggest that forms are emphatic (i.e., some repeated pronouns) at the same time.

It is also not clear what role is afforded to pragmatics in this account. The authors appeal to a notion of context, and talk of repetitions which have non-semantic functions, which suggests that they must have some kind of pragmatic wastebasket in mind which takes care of the interpretation of certain repetitions. Moreover, it seems that when they talk about a 'socio linguistic' motivation for a given case of redundancy, they might mean 'pragmatic' in the sense of changing what you utter to take into account who you are talking to (see §4.5.5 of this chapter). However, Wit and Gillette (1999) offer no cognitive account of how hearers interpret the repetitions. They seem to seek to link particular forms with particular effects. It is not really explained how a given repetition is identified, or how its function is decided upon. Furthermore, and quite worryingly for the data I have presented in this chapter, their model is not actually able to account for repetitions which occur above the 'level' of the paragraph. The authors rightly note that repetition can occur within paragraphs, and not just within sentences, but they do not offer any account of repetitions over large(r) stretches of discourse, text, or time. Finally, although Wit and Gillette (1999) do seek to produce a taxonomy of repetitions and effects, they fail to make an important distinction. They do not distinguish between deliberate stylistic repetitions, and repetitions which are simply cases of necessary re-use or cases of the speaker making sure that he has been heard properly (see

'enhancing comprehensibility' and 'resolving ambiguity' above). It is especially important to make this distinction when explaining how 'long distance' stylistic repetitions are recognised in the first place, as I have suggested above.

Wit and Gillette (1999) don't call what they examine *repetition*, and stick to using the term *redundancy*, even if they are trying to suggest that some repetitions of form aren't actually redundant because they have other consequences than 'shoring up' a message or improving comprehensibility. As such, it is important to recognise they are not trying to give a full account of stylistic repetitions, particularly those that occur after a good deal of linguistic material has been presented between an original and its repetition(s). However, a critical examination of their account suggests that any explanation of 'long distance' repetition must take care to not conflate form and effects, and requires clear notions of emphasis and intensification. Moreover, we need a pragmatic account of how all stylistic repetitions are recognised and identified as intended to communicate stylistic effects, but without appealing to taxonomies that seek to link forms with the guaranteed recovery of particular effects. Finally, I think it is important to stop thinking about redundancy, and call repetition *repetition*. If a repeated form does not exist to prevent communication from breaking down, or is not a case of incidental or necessary use, then it should be called *repetition*. This is not to say that there is no redundancy in ostensive-inferential behaviours in communication, however.

As I explained above, language systems can exhibit redundancy and, so, ostensive-inferential communication that involves language could inherit some redundancy in this way. Moreover, communicators can utter things again in order to increase their chance of being understood as intended - perhaps in a noisy environment where they might expect communication to have a good chance of failing due to environmental interference, for example. Ostensive inferential communication involving nonverbal behaviours can also exhibit redundancy. We might point to something and nod towards it at the same time in order to draw an audience's attention to it. If you ask me what time it is, I might either point or nod to the clock, or both. In the case where I do both, you may see me point but miss my subtle nod - in this case, it could be said that the nod is, from an information theory perspective, redundant. However, just because some repetitions of ostensive-inferential behaviours in communication are redundant, it does not mean that all repetitions are redundant. If you consider the examples of stylistic repetition and non-stylistic repetition put forward so far for discussion, it does not

seem that the stylistic cases are intended to 'shore up' communication and prevent breakdowns. Stylistic repetitions don't seem to be produced for the same ends as, say, someone repeating an utterance to recreate his original attempt at optimal relevance when he wasn't heard the first time. This will be problematic for anyone claiming that all repetitions exemplify some type of redundancy in communication. Of course, to demonstrate that 'long distance' repetitions are not redundant, I have to explain how they achieve relevance. During the course of the analyses, I address why these interpretations result in a particular type of intensification, and whether or not there is any emphatic communicative behaviour present.

4.5.4 Expectations, ostensiveness and 'distance' in recognition of 'long distance' repetitions

Below are reproduced three examples that I introduced earlier in this chapter.

(9) [From 'Anthem', by Ayn Rand (2010) (Total pages: 102)]

Page 5: The sleeping halls are white and clean and bare of all things, save one hundred beds.

Page 6: The sleeping halls are white and clean and bare of all things, save one hundred beds.

Page 14: The sleeping halls are white and clean and bare of all things, save one hundred beds.

(13) [From 'Miranda', BBC1, 01/2014]

Miranda at 3 minutes 34 seconds: **Plunge** the loo before I go…expunge the bath…any activities ending in –unge should be banned…I am on an –unge protest. Although **PLUNGE** is a lovely word…**plunge**…(to camera) **PLUNGE**!

Miranda at 10 minutes 21 seconds: Ooh! Crack. Cheeky! Crack! **PLUNGE**!...**Plunge** m' crack.

(14) [From '300: Rise of an Empire' (2014) (142 minutes' running length)]

Queen Gorgo **[at the start of the film]**: The Oracle's words stand as a warning. A prophecy. Sparta will fall. All of Greece will fall. And Persian fire will reduce to cinder.

> For Athens is a pile of stone and wood and cloth and dust. And, as dust, will vanish into the wind. Only the Athenians themselves exist, and the fate of the world hangs on their every syllable. Only the Athenians exist and only stout wooden ships can save them. Wooden ships, and a tidal wave of heroes' blood. **[The whole monologue is repeated right at the end of the film.]**

Earlier, I established that each of these repetitions is associated with the recovery of stylistic effects, and I believe that this is how the authors of the cited works intended them to be interpreted. I also treat the repetitions in (9), (13) and (14) as cases of *indeterminate showing*. The communicators 'play on' and draw additional attention to the forms that they repeat. The import of these repetitions is very vague and open-ended, and cannot be paraphrased propositionally at all. As I have suggested elsewhere, a communicator deploying such repetitions can be taken to be exploiting the fact that he did not utter a novel form in a context where he might have been expected to do so. This is marked, because, all other things being equal, we generally expect to encounter new and different forms as communication proceeds. Of course, the chance of having this identified by an audience is related to the salience of the repetition of form. In more subtle cases, or cases where the repetition might be 'drowned out' by other stimuli competing for attentional resources, the *showing* may not appear very ostensive - particularly over 'distance'. However, there are ways for the linguistic forms *shown* in 'long distance' repetition to be made more salient, and/or more expected. For one, the authors of the texts that yielded (9), (13) and (14) all exploit particular expectations and assumptions that one might reasonably expect their audiences to bring to bear on the interpretation process when reading a novel, or watching a TV comedy or film.

In Relevance Theory, as I set out in chapter two, §2.2, the context (for utterance interpretation) can be defined as follows:

> 'The set of premises used in interpreting an utterance…constitutes what is generally known as the *context*. A context is a psychological construct, a subset of the hearer's assumptions about the world…A context in this sense is not limited to information about the immediate physical environment or the immediately preceding utterances: **expectations about the future**, scientific hypotheses or religious beliefs, anecdotal memories, **general cultural assumptions**, beliefs about the mental state of the

speaker, may all play a role in interpretation.' (Sperber & Wilson, 1995, pp. 15-16, my emphasis).

Humans constantly make predictions about what is coming next. When we walk, we predict where our foot is relative to where the floor is expected to be. When we communicate with each other, we make predictions about what will be uttered next, and what the speaker might intend by that utterance. When we read novels, or watch television programmes and films, we make predictions about the type of language, devices and content which we might encounter. Predictions about the forms that we might expect to encounter in certain types of text genre or media format explain why we recognise some stylistic repetitions as deliberate. Sperber and Wilson (1995, pp. 15-16) consider that expectations about the future, and about cultural assumptions and, presumably, predictions regarding genre expectations will be brought to bear on the context for interpretation.

(9) is a repetition from a novella. The type of individual who would be likely to read an Ayn Rand title is likely to be familiar with lots of books, and is aware that repetition is a common device in fiction. As a result, the reader of (9) would have some level of expectation based on general cultural assumptions that she would be likely to encounter repetitions in this text. If you expect to encounter something, it is more likely that you will recognise it. You may even already be devoting small amounts of cognitive resources to that which you are looking for, especially for avid, considered readers. In (13), let's suppose that the viewer of *Miranda* watches the show regularly, and is aware that Miranda commonly repeats expressions she finds amusing for comedic effect. The *Miranda* viewer has an expectation based on general cultural assumptions, including assumptions about the typical format of her favourite show, that she is likely to encounter repetition. The expectation that this repetition is deliberate, as in the other cases discussed here, is justified in part by the scriptedness of the show. The audience of the film example in (14) will also bring expectations about genre and other general cultural assumptions to bear on the interpretation process. Viewers are more or less aware that scriptwriters and directors use particular devices in film, and one of those is presenting (and, thus, from the point of view of the writers and directors, *showing*) elements of a film that you have already heard or seen, e.g., in flashbacks or montages. Moreover, *300: Rise of An Empire*, I think, features certain aspects of style, and certain rhetorical devices that would linked by the ordinary person to classical texts of ancient times - the time period in

which the film was set. Consider the below dialogue from an official trailer of the film, which some viewers will have seen in advance of viewing the film, and which may have lead some viewers to want to see the film on the basis of what they think it will be:

> (17) [Official Trailer 1 for *300: Rise of an Empire*, Warner Brothers, viewed on YouTube, December, 2014.]
>
> Queen Gorgo: It begins as a whisper, a promise. The lightest of breezes dances above the death cries of three hundred men. That breeze became a wind, a wind my brothers of sacrifice. A wind of freedom. A wind of justice. A wind of vengeance.
>
> (This dialogue appears in the actual film, and is quite representative of the way in which many characters speak.)

The trailer dialogue itself contains repetition. It exhibits metaphoric use, apposition, and a measured, well-crafted, theatrical and dramatic style. These are the type of things that the layman might well associate with the classical works that the film is intended to evoke and emulate. Particularly if someone has seen the trailer in advance of viewing the film, it is clear why viewers would be justified in expecting to encounter some repetition for some kind of effect, and might already be somewhat primed to look for it. Thus, (9), (13) and (14) demonstrate that communicators, particularly authors or scriptwriters that seek to evoke certain effects in their audience, can exploit cultural expectations and genre expectations to slightly increase the chance that repetitions will be identified by audiences as intended.

Another way that a repetition at some 'distance' can be made more salient is to make sure that you repeat something that is already quite marked anyway. Essentially, in these cases, loosely put, a communicator can exploit the fact that the original material that was repeated was already quite ostensive and served to 'trip up' an audience and encourage them to attend to it in pursuit of extra or different effects the first time around - this is the sense that I mean *marked* in here. This is the case for (9), (13) *and* (14). (9) is marked in terms of form in three ways. First, it hosts *save*, which is much lower frequency than the synonymous expression 'except for'. Second, multiple conjunctions are present where listing with commas would be possible. Finally, an entire sentence has been repeated, and one which is neither particularly short nor structurally simple. What is repeated in (13) is already the type of form that you

would pay more attention to in the anticipation of stylistic effects (onomatopoeia). In (14), what is repeated actually contains repetition - fairly adjacent lexical and syntactic repetition which cannot be explained (solely) in terms of necessary re-use. My suggestion is that such fairly adjacent repetitions are, all other things being equal, easier and less-effortful to recognise than those which require the processing of much more material in between the original and its repetition(s). Thus, the repetition-within-the-repetition in (14) was already quite ostensive. In this way, these examples show that communicators can increase the chance of their repetitions being identified by repeating something that is already noticeable and marked. Very loosely put, seeing or hearing something 'odd' once is of some interest, but should something marked occur more than once in a book or film, for example, it is not likely to be a coincidence, and one might be more inclined to attend to it as deliberate, stylistic repetition.

Finally, I would like to emphasise a point I just made above. It should be clear from the examples and discussion presented in this chapter that one way that a repetition can be made more, or less, salient (depending on how strong you want your communication to be) is to manipulate the 'distance' between an original and its repetition(s). The higher the proportions of repetitions in a text, the more salient they will be, all other things being equal. This is because the activation levels of the initial interpretation will still be relatively high, and the form that was repeated will be relatively at the front of the audience's mind. Of course, aforementioned genre expectations can interact with this, including assumptions about the length of a book or a film, for example. This general observation allows us to come up with a working notion of 'distance' which can be meaningfully used in the discussion of stylistic repetition data. I propose that 'distance' be understood as a relative notion. We will not be able to say that particular repetitions are 'long distance' if they occur every ten pages, or every twenty minutes, for example. Repetition over distance is to be potentially understood as any deliberate repetition of linguistic form which is not directly adjacent. Whether a repetition is considered to be a 'long distance' is something that emerges post hoc, and is identified considering how much, if any, repetition an audience expects to encounter, and may involve assumptions about text or discourse length. It cannot be specified in advance, and is unlikely to be something we would treat in a formulaic, predetermined manner. Thus, in a ten-minute conversation, 'long distance' repetition might be assumed when the same

sentence is uttered in four times. In a text message or tweet, however, repetition at distance might be understood as occurring within the scope of a few hundred characters. 'Distance' is always going to be relative. Our understanding of the notion will also need to take into account the activation levels and processing effort required in recognising and interpreting repetitions. This is linked with why repetitions of the kind seen in (9), (13) and (14) should not be considered redundant in any sense of the word, and with the type of processing they engender and the effects they result in. In any case, it seems that the ostensiveness of repetitions plays a role in why they are recognised. Their salience creates an expectation that there will be effects on offer. The effort the hearer is put to will only be offset if a repetition makes a useful update to her cognitive environment, i.e, it is not redundant.

Redundancy is a failsafe built into a system to ensure that one single message is conveyed. Stylistic repetitions could only be redundant if it can be shown that they exist to guarantee the transmission or exact re-transmission of some single message. In the final section of this chapter, I show that stylistic repetitions are not redundant because they don't exist to guarantee the transmission of some existing, or *old* message should there be noise in the system. Such repetitions exist to scaffold the recovery of some *new* or *adjusted* effect, and can only be taken that way - to treat them as simply giving access to what might be termed 'old' information would not justify the effort that audiences are put to in processing them.

4.5.5 Why 'long distance' repetitions don't yield just *old* information

The field of Information Structure concerns the way that information is 'packaged' within sentences taking into account, among other things, the mind of the addressee (Chafe, 1976; Krifka, 2006). One important distinction made in Information Structure (IS) is the distinction between information that is *old* (or given, familiar), and information that is *new* (unfamiliar) (see, for example, Chafe, 1976; Ward & Birner, 2001; Krifka, 2006). Chafe (1976, p 30) treats given or *old* information as knowledge which the speaker can already assume to be in the mind of the addressee when an utterance is produced. He (*ibid*.) goes on to say that *new* information is that information that a speaker believes himself to be introducing into the mind of his addressee by means of an utterance. IS is not designed to account for the specific case of stylistic repetition, whether over any distance or not. The concern is simply with how information in sentences is packaged given consideration of the speaker, hearer and

discourse in communication in general. Moreover, if some researchers of IS treat it as the study of the structure of sentences or utterances, they may not have any interest in any recurrence of linguistic form above the sentence/utterance level, i.e., over more significant expanses of discourse or text. I was not able to find a single publication that addressed *stylistic* repetition from the point of view of Information Structure. However, we might imagine how repetitions of form could be viewed by some working on IS.

For Krifka (2006, pp. 2-4), one of the key tasks in communication is management of the Common Ground (CG). The CG is defined as a model of information exchange which accounts for how information that is mutually known is shared and continuously modified or updated in communication (Krifka, 2006, pp. 3-4). On one hand, the CG consists of propositions which are mutually accepted, as well as any entities that have been introduced into the discourse before (Krifka, 2006, p 4). This is called CG content (*ibid*.). However, on the other hand, Krifka notes that anything like CG would have to include information about the goals and interests of the participants (*ibid*.). It seems that CG content falls squarely on the side of what contributes to the truth conditions of utterances, while any aspects of structure that do not contribute to truth conditions would be treated as 'pragmatic use' and would be subsumed under the heading of content management, which is a task that is undertaken by the speaker and hearer together, and concerns the way that CG content is to develop (*ibid*.). I understand Krifka to mean that communicators modify the 'packaging' of their utterances in line with their interests and abilities, and the interests and abilities of the speaker. There is, of course, no doubt that communicators do this (see discussion of optimal relevance in chapter two, §2.4.2).

With this in mind, I would like to think about what some IS researchers might say about the communicative stylistic repetitions presented so far in this chapter. It appears that if a repetition were to serve a communicative goal and did not contribute to truth-conditional content, it would be called a 'pragmatic' repetition. If a repetition related to factual information and/or truth-conditional content, it would be treated as part of the CG content to which it would presumably introduce *old* or *new* information, and would not be treated as 'pragmatic'. Here, pragmatics is usage-based, and semantics is the domain of the truth-conditional. As per the discussion in chapter two, this is not the right way to cut the semantics-

pragmatics cake. As Carston (2002) showed, pragmatics shoulders more of the responsibility for recovering the proposition expressed than thought previously.

There are issues with whether or not, and how, repetitions might be categorised as part of content management or as contributing to CG content. 'Pragmatic goals' seem to be things like turn-taking or organising discourse in a particular way based on the anticipated needs of your co-communicator (see Krifka, 2006). As such, this is a usage-based conception of pragmatics, and not one that allows a significant role for pragmatic inference, not in the recovery of what Relevance Theorists call implicit content, and certainly not concerning the recovery of what is thought of as co-extensive with semantics for many outside Relevance Theory, i.e., truth-conditional content. This is where we have our first problem. Faced with a repetition, on this approach, I don't know how a hearer should decide if it should signal some kind of move or organisation of discourse on the 'pragmatic' side of things, or whether it should indicate something about CG content. In fact, this matter is even more complicated - if a repetition would not be judged to contribute to CG content, it must be taken as part of content management, and that would mean that stylistic repetitions would never be picked up on for the communication of stylistic effects - repetitions would just be 'moves' or 'turns' or 'strategies'. If repetitions 'at distance' were even noticed, and judged not to contribute to CG content, they would probably be taken as cases of repetition which represent an original attempt at optimal relevance, for example, cases where you repeat yourself because someone did not hear you the first time around.

Repetitions that serve some sort of usage-based, or social, or conversational function are not produced to lead to the recovery of intended stylistic effects and so I am not concerned with what researchers in IS might say about 'pragmatic' repetition and CG management. What I want to do is get to a picture of 'long distance' repetition where I can show that these are not redundant because they *do* scaffold the recovery of *new* or *amended* interpretations. For this picture to emerge, it is helpful to show that stylistic repetitions cannot be taken as contributing *old* information to the CG - which is how I suspect they might be taken by some working on IS.

'*New* information is defined as entities and propositions whose semantic content is not already present in the common ground of the discourse; *old* information, predictably that is

present in the same way' (Delin, 1987, pp. 115-116). On this view, information in the CG is going to be propositional; the common ground is made up of propositions. Moreover, 'old' and 'new' information concerns *semantic content*. This will either be co-extensive with truth-conditional content, or some view of semantics that relies heavily on the linguistic semantics to provide some representation that requires minimal work from pragmatic inference, e.g., just reference assignment and disambiguation. Let's put aside the issue that any communicative repetitions which are borderline or non-linguistic would not be able to contribute any information to a CG on this view. Let us also put aside the fact that any repetitions that communicated non-propositional representations would also be unable to contribute to the CG. On this view, presumably, you would have to say that repetitions of linguistic material would have to be seen as yielding *old* information. There is a sense where it could be argued that what you have already manifestly uttered before has already resulted in an interpretation, and so already contributed to the CG. To repeat something is simply to re-introduce or re-activate the information that is already there either in the discourse or the mind of the speaker, and this could be construed as *old*. However, I suggest that repetitions, if anything, would have to be classed as providing some new or improved information to a CG, if one is taking such an approach. After all, if there were no new or improved information on offer for being put to the effort of processing sometimes very ostensive and laborious repetitions, they would not be an effective and fairly common device employed in the communication of non-propositional effects.

Perhaps, then, we need to think about stepping away from the notion of *old* and *new* information in Information Structure in order to understand how we might account for the effects of repetition, and, also, to prevent stylistic repetition from being considered informationally redundant. In a paper by Egbert J. Bakker (1993), a similar suggestion has been made. A key premise of this paper is to suggest that the old / new distinction should be abandoned, and that repetitions, rather, should be treated in terms of *activation*. Bakker (1993) approaches repetition from the perspective of the oral traditions, more specifically, from the perspective of re-enactment and performance in oral traditions. Essentially, he (1993) notes that oral cultures rely on a lot of performance and re-enactment of texts – not just for their transmission and preservation, but for the effects on offer. He makes the point that, in such performances, what we might consider to be 'old' because it has been recently

encountered cannot really be considered 'old'. Performances are interesting and enjoyable, not even though they feature repetition, but *because* they have repetitions which are associated with particular effects. As he puts it (1993, pp. 7-8), '[w]hat is "known" can be highly salient in terms of perception', and he goes on to say that information need not be totally new in order for it to have some effect. If a protagonist in a play is 'mentioned' again and again, effects can be had by adjusting the concept of the protagonist in light of changing contextual assumptions as the performance unfolds. He also notes that the speaker who is only informative, and who only offers 'new' information can actually be a complete bore. 'What matters in speech' says Bakker (1993, p. 9), 'is not whether something is new or old information (knowledge) but the dynamic cognitive process of *activation*, the appearance in the speaker's and listener's consciousness of an idea out of inactivity' (my italics).

I now apply a consideration of activation to the 'long distance' Ayn Rand repetition in (9) to demonstrate that the repetition has to be understood in resulting in extra cognitive effects through the collection of very weakly communicated weak implicatures, and so cannot be understood as redundant in any sense of the word.

4.5.6 'Long distance' repetition and intensification in pursuit of implicit content

Consider (9) again, reproduced below:

> (9) [From 'Anthem', by Ayn Rand (2010) (Total pages: 102)]
>
> Page 5: The sleeping halls are white and clean and bare of all things, save one hundred beds.
>
> Page 6: The sleeping halls are white and clean and bare of all things, save one hundred beds.
>
> Page 14: The sleeping halls are white and clean and bare of all things, save one hundred beds.

In (9), the original and its repetitions are part of a description of different living quarters that citizens of Ayn Rand's imaginary dystopic world are forced to live in. When the reader encounters the words *the sleeping halls are white and clean and bare of all things, save one hundred beds* for the first time, she will do two things. She will construct an explicature that

expresses a proposition that makes a claim about the way the room looks in the first set of sleeping quarters. She will also recover a set of weakly communicated implicatures on the basis of the context for utterance interpretation (which may include contextual assumptions concerning the preceding text, and genre expectations), and her expectations of relevance in that context. However, as this is an original, and not a repetition, she will recover a set of weakly communicated implicatures that are more strongly communicated than any she will recover from a later repetition, as she is not put to any effort of having to recognise that anything is repeated, and of having to infer what processing she should undertake as a result of any repetition. Below are some of the implicatures that the speaker could recover on her own responsibility. The list is illustrative, and is not fully determinate - it's merely an example.

(18) Set of weakly communicated implicatures that hearer of (9) could recover.

[One hundred people sleep in this room]
[The humans in this story are raised like battery hens]
[Parts of this story are very similar to elements of Aldous Huxley's 'Brave New World']
[These rooms sound like hospital wards]
[It must be awful to share a room with 100 others]
[I cannot imagine having no possessions in my room]
[This novella is clearly a dystopic novella]

As the reader is recovering these weakly communicated implicatures in pursuit of relevant cognitive effects, clearly, they become much more manifest to her. The readiness with which they could be employed in any inferences, or with which they could be accessed and adopted as part of the evolving context for interpretation of the text, increases. In other words, these particular assumptions are much more *activated* as they are attended to and perhaps incorporated into the context. At the end of the original in (9), the reader is presented with a full stop. As I explained earlier in the chapter, punctuation can 'stand in' for intonational cues (Wichmann, 2013), and so the reader here will begin to 'wrap up' the construction of any explicatures and the collection of the range of weakly communicated implicatures, and prepare herself to encounter new units of relevance, as one would normally expect in a novella. The reader is then presented with a page of text which does not contain repetition. As such, the reader's attention turns to processing the new linguistic material, and to

attending to and incorporating new or different contextual assumptions into the ever-shifting context for utterance interpretation. As her attentional resources are diverted elsewhere, the activation levels of any contextual assumptions that the reader used to interpret the original utterance in (9) dip. Moreover, the activation levels of the interpretation she recovered from (9) also dip, and are less at the forefront of her cognitive environment as she concentrates on processing the previously unseen linguistic material of the coming text.

A page later the reader is presented with exactly the same utterance again. At this point, we could say that the repetition represents a subtle-ish case of *showing*. It could be argued that Ayn Rand takes advantage of genre expectations, and the existing markedness of what she repeats to openly display the repeated linguistic form. At this point, if we assume that this makes the first repeated form in (9) salient enough to attract attention to it, the reader will need to reason from being presented with this form, to the layer of information that she is intended to recover from processing it in a particular way. The reader certainly can't take the repetition as simply to be processed as if it were being encountered for the first time. What would that yield in terms of extra or different cognitive effects? If that were all the reader should do, the repetition would simply deliver what the linguistic semantics delivered the first time around, and a few weak implicatures. There would be absolutely no useful improvement to the reader's cognitive environment if the repetition were processed in the same way as the original. As we saw in my discussion of Sperber and Wilson's (1995) account of epizeuxis, the way that its effects are recovered seem to rely on inferring that one has to 'go back' to an original interpretation and re-work it for additional cognitive effects. In a similar way, what 'long distance' repetitions do is send the audience back to *expend* effort on *re-activating* and *extending* the context for interpretation. This necessarily involves dredging up an existing interpretation, and so we can see why some might associate stylistic repetitions with 'old' information. However, as we see below, the resulting interpretation from the first repetition in (9) is not 'old', because it is reworked and quite different from the one that was originally recovered.

Having been presented with the same utterance a page after the original was presented, let's think about what is happening in terms of processing and interpretation from the reader's perspective. As above, let's assume that Ayn Rand's deliberate display of what I suggested

was an already marked utterance is salient enough to attract attention, and the hearer decides that the only optimally relevant way to process the utterance is to 'go back' to and manipulate that existing interpretation in some way. The activation levels of the initial interpretation, including the set of weak implicatures the reader recovered, have dipped as the reader interpreted the intervening text. However, the levels of activation are still reasonably high, and any assumptions that were manifest to her the first time around are certainly more manifest and accessible than those which are not currently manifest to her at all. Accessing only those assumptions, and only the initial interpretation would not be very relevant to the reader. Encouraged to attend to the repetition by means of *showing*, the reader reasons that she needs to attend to the original form and its interpretation, and should *expend* time and effort 'digging deeper' in pursuit of effects. As such, the original interpretation will become more activated than it was the first time around, as the reader is going to expend effort extending the context and adding in more weakly communicated implicatures than she did the first time. Much more at her own responsibility (because she is presented with the same linguistic form rather than a new one), the hearer judges that all she can do is *expend* effort on seeking *more* weak implicatures that satisfy her expectation of relevance, and she may add more to her list, such as those in (19):

(19)
[These sleeping quarters sound uncomfortable]
[Communal sleeping is definitely not something I would enjoy]
[The future sounds really bleak]
[Oppressive states can interfere in any area of your life, including sleeping arrangements]

The repetitions of form in (9) are bounded by a full stop, which is interpreted as standing in for an intonation boundary, as I set out above. The reader, when she gets to the end of the first repetition, determines she should 'wrap up' processing of the repeated utterance and, again, prepares herself to process new units of relevance. This time, almost seven pages of linguistic material intervene between the first repetition and the second and final one. As such, the activation levels of the interpretation that the reader has just reworked and extended at some effort falls, and falls much more than in the case of processing the original

and the first repetition. However, when the line repeated in (9) is encountered for the third time, this is judged to be quite ostensive indeed. The reader cannot miss it. It was definitely intended. One repetition could be incidental. More than one is likely not, given the context.

Since the second repetition in (9) is much more ostensive than the first, the hearer is encouraged to *expend* even more effort re-activating her on-going interpretation, *extending* the context for interpretation further still, and adding in even more weakly communicated implicatures. But what we have is an interpretive situation where, on one hand, the activation levels for the reworked interpretation are fairly low owing to the processing of seven pages intervening text, but, on the other, the speaker is suggesting to the reader by means of a very ostensive repetition that she should expend a lot of effort here on re-activating and improving the interpretation. A lot of activation and effort is needed to jump from the lower level of activation to the increased level - a level of activation that is higher than when the utterance was processed initially and, also, when first repeated. Furthermore, more effort is expended in adding further weak implicatures to the reworked interpretation, which means that more assumptions are more activated - the reach of activation is greater in terms of the cognitive environment. The communication now is very weak indeed. It may even be that the reader does not really add definite assumptions to the activated set of weak implicatures that she entertains; the utterance may simply serve to mildly increase the manifestness of quite a wide range of assumptions, and any range of these could satisfy the reader's expectation of relevance. The effect is increasingly cumulative and impressionistic.

If this is correct, it seems that stylistic repetitions cannot be taken as solely giving access to old information even if they encourage a hearer to attend to an old interpretation while pursuing updated effects. Moreover, the effort of recognising and processing these repetitions appears to be offset by improved and updated interpretations that contain more weak implicatures recovered at the hearer's responsibility. In (9), (13), and (14), the repetitions would also serve to create greater cognitive and affective mutuality with the characters uttering them, which, in terms of consumption of fictional works, might engender greater empathy with those characters depicted as using repetitions.

4.6 Epizeuxis, 'long distance' repetition, and insights into boundaries, emphasis and intensification

This chapter began with an examination of Sperber and Wilson's (1995) account of epizeuxis. There were several things that I wished to consider in this chapter, including whether or not the *weak implicature account of repetition* could be extended to other phenomena, including repetition 'at distance', and apposition. I also wanted to think about the role that 'distance' could play in how repetitions are interpreted, and consider what the repetitions in this chapter might tell us about intensification and emphasis. I also had a tentative hypothesis that intonation boundaries in repetition might restrict hearers or readers to recovering effects at the level of implicit content, as this is where the effects of many repetitions appear to be recovered. Certainly, it seems that the *weak implicature account of repetition* can be extended to repetitions that occur at distance, as the effects of cases like (9), (13) and (14) do lie in 'going back' and adding in further weakly communicated implicatures to an existing interpretation. The account also sheds light on how appositions might be interpreted, as Blakemore (2008) claimed. Thus, Sperber and Wilson's (1995) *weak implicature account* of epizeuxis is vindicated to some extent. However, it is not clear that their prediction that repetitions have 'emphatic effects' is borne out. Not all of the repetitions addressed in this chapter have a sense of emphasis about them. In particular, the repetitions 'at (a great deal of) distance' don't feel emphatic at all. The only 'long distance' case that might be considered to have a sense of emphasis would be the *Miranda* case in (13). If we consider what makes this repetition emphatic, it is not just the repetition itself, but it is Miranda's *showing* behaviour involved in further displaying the repetition - she looks to camera, produces a highly salient nucleus on *plunge*, and surrounds *plunge* by clear intonation boundaries. Some tokens of *plunge* are also placed close together, decreasing the degree to which attention is focused elsewhere between tokens, and further increasing the salience of any repeated material. As such, emphasis certainly seems to be something the speaker does, and something to do with quite ostensive *showing* behaviour. 'Distance' seems to be a way of manipulating the salience of a stylistic repetition, contributing to the overall ostensiveness of its display.

In terms of intensification, I have been suggesting that this is a processing phenomenon. If we look at the interpretations in terms of weak implicatures for both adjacent and less-and-less

adjacent repetitions, in all cases, the speaker is to 'go back' and *expend* more effort on *extending* the context for interpretation, as Sperber and Wilson (1995) set out for epizeuxis. However, in doing this, existing interpretations are re-activated and *further extended* and *activated* as new weakly communicated implicatures are added in, and as the context for interpretation evolves and is updated. 'Distance' can also play a role in the amount of effort required in the activation involved in processing repetitions, and greater 'distance' can lead to the investment of relatively more effort and, thus, result in an expectation that more or different effects should be sought. In any case, I believe that the discussion of more-adjacent and less-adjacent repetitions in this chapter suggests that it is possible to have intensification in pursuit of *implicit* content, and that, at this stage, I suggest it is identified in processing when a hearer is encouraged to *expend more effort in extending an interpretation*, and in a way that involves *re-activation* and *extra-activation* during utterance interpretation. My findings here link with the comments made by authors writing about repetition whereby it was considered connected with a sense of continuation and prolongation (see chapter one, §1.6). Re-activation of interpretations would certainly explain this, along with any sense that some repetitions are cumulative, or 'build' in some way.

One of the key aims of this thesis was to set out how it is that particular repetitions are recognised, and are recognised as intended to communicate non-propositional effects. The discussion above makes the case that, sometimes, this is the only way that stylistic repetitions can be taken - it is not optimally relevant to treat them in any other way. Many stylistic repetitions of the kind addressed in this chapter are quite ostensive, and so they make demands on our processing resources and create an expectation of relevance, and this is why we recognise them in the first place. These repetitions don't just contribute old information, or serve some kind of communicative functions in terms of 'moves' or 'strategies' in the organisation of discourse. What I suggest is that speakers, in more or less subtle cases, either display or full on *show* the repeated form they have already uttered to encourage audiences to reason from what is *shown* to what the speaker/writer wants to communicate - which is likely a fairly indeterminate range of weak implicatures here. As such, all the stylistic repetitions in this chapter are cases of *indeterminate showing*.

Finally, it should emerge from this chapter that one thing a communicator can do to improve the chance that his displayed repetition is definitely picked up on is to further highlight or

delimit it in some way. A speaker like Miranda could prosodically highlight an aspect of a repetition. Or, a speaker could put intonation boundaries around a repetition to further delimit what is being repeated. Thus, just as I can draw a circle around an amount I want you to attend to on a bank statement, I can put intonation boundaries around what I want to *show* in terms of linguistic form - like Miranda did in (13). What else do intonation boundaries do when they interact with repetition? Although I have demonstrated that they play a role in suggesting when an audience is to start or stop collecting particular effects, this appears to be a by-product of a general 'wrap up' function that intonation boundaries serve, and, so, it does not seem that boundaries guarantee whether effects should be collected at a particular level of content. To shed more light on this, the next chapter examines how repetitions of material within the same (intonation) groups achieve relevance. This chapter allows us to discuss some of the cases I mentioned earlier that Sperber and Wilson (1995) cannot explain, including some repetitions of nonverbal material.

Chapter Five: The Repetition of Material Within (Intonation) Groups

5.1 The existence of repeated material within intonation groups

In the last chapter, it was established that utterances can contain repetitions of complete intonation groups. The repeated intonation groups can be adjacent, or they can have some intervening linguistic material between them. Such repetitions were generally shown to contribute to the recovery of effects at the *implicit* level of content. However, there are also particular repetitions of linguistic material which occur within a *single* intonation group and which do not seem to primarily constrain the recovery of implicit content. These repetitions have not been addressed in this thesis so far, and they are not explicitly mentioned in any text that I reviewed during my research. Consider the following examples:

(1a) He's a big big guy.
(1b) We went for a long long long walk.
(1c) That's one dead dead corpse.
(1d) I've got a full full schedule today.
(1e) It's really really exciting what we're doing for New Year.
(1f) Our kitten was abandoned. Cleo is very very very small.
(1g) Rosie: Do you want to see the new James Bond film tonight?
Kelly: Yes yes yes. / No no no.

These repetitions are *not* found in a contrastive context, and they *do not* feature the distinctive intonation pattern found in the X-x cases in chapter three.

The examples in (1a-1g) feature a variety of expressions: the so-called relative adjectives *big* and *long*, and the so-called absolute adjectives *dead* and *full*, the so-called intensifiers *very* and *really*, and the negative polarity particles *yes* and *no*. At first glance, it seems that the repetitions in (1a-1g) interact with the recovery of explicit content. All other things being equal, a 'big big' guy is certainly physically bigger than a guy who is just 'big'; a 'long long long' walk is surely greater in length than a walk that is just 'long'. While we might argue that the patient in (1c) is either dead or not, it does seem that the concept DEAD* expressed here is a very particular one. We might imagine two mortician colleagues discussing a single corpse

that is exhibiting full rigor mortis, signs of blood pooling, and bloating, all of which would not be present in someone newly deceased. The repetition here seems to influence the nature and extent of the deadness of the patient that the hearer recovers via the ad hoc concept DEAD* in that context. The same sort of scenario would apply to the 'full full' diary schedule in (1d). A plan that is 'really really' exciting is more exciting than one merely described as just exciting; Cleo the cat is clearly much smaller than one might imagine if she was only described as 'very' small. These repetitions very obviously interact with the truth conditions of the utterances that host them, and, from a lexical-pragmatic perspective, these repetitions (apart from (1g)) seem to influence the computation of very finely adjusted ad hoc concepts. As such, it looks as if these repetitions of linguistic material within a single intonation group are responsible for constraining the identification of explicit content. (The polarity particles *yes* and *no* could be treated in Relevance Theory as leading to the recovery of higher level explicatures. These involve the embedding of a propositional form under some kind of comment or attitude description (Carston, 2002 (and see chapter two, §2.5.2)) and, as such, their repetitions are principally treated as interacting with the recovery of explicit content as well.) The main point of this chapter is to provide a detailed explanation of how these repetitions achieve relevance.

As I set out in chapter two, §2.5 and §2.6, in Relevance Theory, expressions are usually treated as either conceptual or procedural. Conceptual expressions can contribute to the truth conditions of an utterance by yielding a constituent of a propositional representation, e.g., the word *dog*; conceptual expressions can, however, be conceptual and NOT be truth-conditional such as the sentence adverbials *seriously* and *sadly*. Procedural expressions do not encode concepts themselves and many do not contribute to truth conditions, e.g., *but* and *however*. Nevertheless, some procedural expressions do affect truth-conditional utterance content - this is considered the case for pronouns and possibly some possessives (Aitken, 2009; Wilson, 2011). Most of the expressions in (1a-1g) above are clearly conceptual. We can introspect on their meanings, and visualise these in the mind's eye. However, some of the expressions are quite semantically bleached, and it's not immediately clear what sort of concept they might encode (*really*, *very*), while others (*yes*/*no*) are not compositional, and not conceptual, even if they do interact with the recovery of explicit content. For this reason, I think it's worth looking at the nature of the semantics of these expressions in much more

detail to understand what happens when these expressions are repeated, especially those modifiers called intensifiers - *really* and *very*. Robyn Carston has acknowledged that there is not currently much work being done on modification within lexical pragmatics (personal communication, 2015). In fact, within relevance-theoretic lexical pragmatics there appears to be little work which deals exclusively with the topic of modification at all. This is surprising given that modifiers could be considered the linguistic means *de rigeur* for scaffolding the computation of a particular ad hoc concept.

This list of expressions that features in (1a-1g) is not intended to be exhaustive. The point is simply to show that such repetitions within single intonation groups actually exist. In the literature I reviewed, these types of repetitions are not mentioned at all - including in Sperber and Wilson's (1995) account of repetition. The only thing I could find was a single example in Wilson and Matsui (1998), where examples of repetition are used to make the point that linguistic material can be found repeated below the level of the clause, which it is argued that some coherence theorists are unable to account for (Wilson & Matsui, 1998). The example is as follows:

(2) [Adapted from Wilson & Matsui, 1998, p 25]
John's a very very clever man.

When spoken, we can see that the repetition in (2) occurs within a single intonation group.

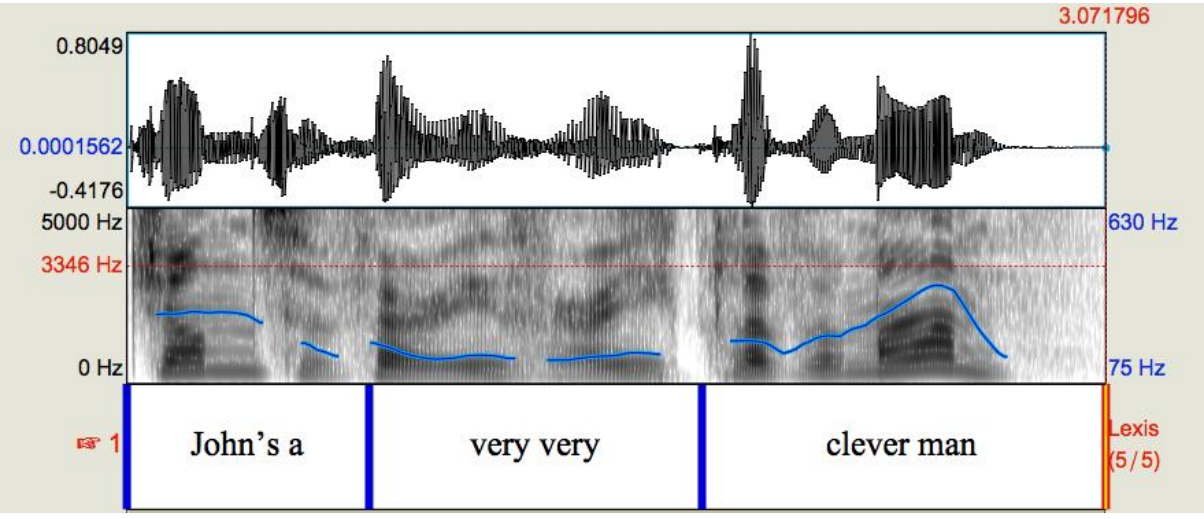

Figure 7 Prosodic analysis for 'John's a very very clever man'

It is unlikely that Wilson and Matsui accidentally omitted to write a comma between the tokens of *very* in (2). However, the authors draw no attention to the fact that there is no comma where prescriptive writing rules require one. From the point of view of form, then, repetitions within a single intonation group are not explicitly recognised. As their form is not addressed in the literature, there is naturally no account I am aware of that addresses how these repetitions interact with the context to constrain the identification of explicit content. This chapter provides such an account. In the last chapter, we saw that adjacent and non-adjacent repetitions of intonation groups contributed to implicit content. Above, the repetitions within intonation groups seem to contribute to the interpretation of explicit content. Perhaps it is the case, after all, that the repetition of material within an intonation group always constrains the recovery of explicit content. Continuing to examine this hypothesis is a key aim of this chapter, and the discussion here extends our understanding of the roles that intonation boundaries play in communication, including our understanding of their role in delimiting units of relevance.

This chapter may also be able to shed light on why certain expressions are considered unacceptable when repeated in particular contexts. It is thought by some (Kennedy, 1999; Kennedy & McNally, 2005) that so-called absolute adjectives have meanings that are not contextually determined. On this view, the meanings of *full*, *dead*, *spotty*, or *closed* are not modulated in context. If these accounts are correct, it would predict that we should not be able to find acceptable repetitions of certain adjectives within the same intonation boundary. Indeed, this seems as if it may be the case:

(3a) ??? He chose a **spotty spotty** shirt.
(3b) ??? I've never met someone with such a **closed closed closed** heart.

Any account of the repetitions in (1a-1g) above must include an explanation of why (3a) and (3b) seem unacceptable. Any examples of adjectival modification presented thus far in this chapter have been cases where the repeated adjectival modifiers occur *before* the argument it modifies. These will be the focus of the first section of this chapter. However, there are also repetitions which meet the formal criteria of repeated linguistic material within an intonation

boundary, but which seem not only to be unacceptable, but quite likely judged as ungrammatical as well.

(4a) * The walk was long long.

(4b) * I am tired tired.

The insertion of an intonation boundary between repeated adjectives renders (4a) and (4b) perfectly acceptable. It turns them into epizeuxis, in fact. However, without the presence of an intonation boundary, this type of repetition is unacceptable. It seems that repeated gradable adjectives used as attributive modifiers are fine whether or not an intonation boundary separates tokens, but they are not able to appear as predicative complements when not separated by an intonation boundary. During the course of this chapter, a pragmatic explanation emerges for this distributional restriction.

An account of how the repeated expressions in (1a-1g) achieve relevance will need to include the usual considerations of how the semantics of the phenomena in question interact with the context and pragmatic processes to yield their interpretations. To properly understand how much of these interpretations can be attributed to the linguistic semantics of the repeated expressions themselves, and how much of the interpretations can be attributed to the fact an expression has been repeated, it is necessary to offer an account of the semantics (and pragmatics) of each of the expressions in (1a-1g) as the analyses progress. I must also comment on whether there is any element of *showing* to the repetitions of linguistic material within intonation groups.

In the first section of this chapter, the majority of the data that are analysed feature instances of modifying expressions. Some of these cases are adjectival, while others involve degree modification. According to McNally (2016, pp. 26-27), modification involves a word or phrase combining with another to produce a phrase of the same semantic type. Modification can be broadly thought of as a process whereby 'a modifier adds additional, non-essential descriptive content to that contributed by the expression that it combines with' (McNally, 2016, p 1). Degree modification involves the combining of an expression yielding a gradable property or quality which also yields information about the degree to which that property or quality holds

of the argument being modified (McNally, 2016, p 21). According to McNally (2016), research into modification has been problematic, and research into degree modification is considered particularly complex. She (2016, p 27) notes that '...a fully adequate theory of how best to explain the division of labour between the semantics of the modifier and the modifiee, the semantic composition rules, and context remains to be developed'. The analyses required for (1a-1f) will require consideration of these issues, and provide insight into the degree of responsibility for the interpretation carried by the context and pragmatic inference in the interpretation of the expressions in question. From a relevance-theoretic perspective, these analyses will go some way to closing the current gap in research on modification in lexical pragmatics.

5.2 Semantic accounts of gradable adjectives

The adjectives in (1a-1d) above are gradable adjectives. Gradable adjectives are sometimes treated as expressions which map their arguments onto measurements or degrees (Kennedy & McNally, 2005, p 10). Their core meanings are said to involve relations to scalar concepts which objects can be ordered against (Kennedy, 1999; Syrett *et al*., 2009). It is generally assumed that any relations between objects and degrees on a scale is the result of a function encoded by the gradable adjective (Bartsch & Venneman, 1973; Kennedy, 1999). The function takes an object as an input and returns as an output a degree to which that object possesses the property in question (Klein, 1991). On this view, when processing an adjective such as *red* or *empty*, a function mandates that some kind of scale is accessed and the degree to which the property in question is predicated of its argument is pinpointed on the scale. It is considered that the structures of these scales are part of the conventional meaning of gradable adjectives (Sassoon & Toledo, 2011, p.3). Gradable adjectives should be thought of as being involved in the communication of representations into which is built some kind of degree of a particular property such as length, height, expense, bumpiness, spottiness, and so on. Importantly, gradable adjectives are argued to be adjectives whose semantics mandates that they are interpreted with respect to some kind of comparison class or standard of comparison (Kennedy, 2007; van Rooij, 2011). For example, to understand how 'tall' someone is, a comparison would need to be made to a set of other entities of which the property of tallness holds to greater or lesser degrees. Finally, Bolinger (1972, p 15) associates

adjectives, along with many adverbs, with 'manifestations of degree and intensity'. However, he (*ibid*.) notes that not all adjectives seem to express degree.

Kennedy and McNally (2005, p 13) say that there are three parameters which have to be specified in the lexical entries for gradable adjectives. These are a set of degrees or measurement values, a dimension specific to the type of measurement, e.g., cost or height, and an ordering relation which is useful in distinguishing between pairs of antonyms (*ibid*.). Gradable adjectives are often sub-divided into two categories - relative adjectives and gradable adjectives (Unger, 1975; Kennedy & McNally, 2005; Kennedy, 2007). The justification for this distinction is based on largely on semantic and distributional factors, and the main difference that is posited to exist between the two categories is the nature of the comparison classes or standards of comparison said to be involved in the computation of their meaning (see Sassoon & Toledo, 2011).

Relative adjectives are said to mandate the computation of a contextually-determined comparison class as part of their interpretation (Syrett *et al*., 2009, p.2). There is a standard of comparison that is some kind of degree or norm for the comparison class (*ibid*.), on this view. If we say today is *hot*, we have to determine some kind of norm level of hotness, and the degree of hotness predicated of today must exceed that contextually-determined norm or average for the comparison class. It is in this respect that formal semanticists would acknowledge that the meanings of such adjectives are context-dependent. On the other hand, the interpretations of so-called absolute adjectives involve reference to comparison classes that are generally considered fixed and independent of the context (Kennedy & McNally, 2005). Exemplar adjectives for this category are *full*, *open*, *dead*, *invisible*, and *spotty*. It is easy to see why theorists would consider that the meanings of such adjectives are not contextually determined. There does not appear to be much scope for modulation of meaning. After all, when Harry Potter puts on his invisibility cloak, he's either invisible or he's not.

There are a number of arguments and pieces of empirical evidence employed to argue for a split in gradable adjectives based on whether they are context-sensitive or not. An important argument has its foundations in distributional evidence from degree modifiers, i.e., which degree modifiers a gradable adjective can legitimately combine with. Kennedy and McNally

(2005) and Kennedy (2007) use the degree modifiers *perfectly* and *slightly* to motivate a typology of four scales that can be associated with various gradable adjectives. The scales are set out as follows with a visual representation of each scale and distributional evidence for its existence (examples and visuals due to Kennedy (2007), and Sassoon & Toledo (2011)):

OPEN SCALE STRUCTURE (*tall*, *short*) - Members of this category mandate the computation of a context-dependent midpoint on a scale.
Scale visual: 0----------0
Distributional evidence #slightly tall, #perfectly tall

LOWER CLOSED SCALE STRUCTURE (*dirty*, *wet*) - Members of this category have a minimum endpoint on a scale.
Scale visual: x----------0
Distributional evidence: slightly dirty, #perfectly dirty

UPPER CLOSED SCALE STRUCTURE (*clean*, *dry*) - Members of this category have a maximum endpoint on a scale.
Scale visual: 0----------x
Distributional evidence: #slightly clean, perfectly clean

TOTALLY CLOSED SCALE STRUCTURE (*full*, *empty*) - Members of this category have minimum and maximum endpoints on a scale.
Scale visual: x----------x
Distributional evidence: slightly full, perfectly full

On this view, if the standard employed in arriving at the comparison class is somewhere on a midpoint of a scale, and the ends of the scale are completely open, then an adjective is *relative*. If the standard for comparison is found at one of the endpoints of a scale, then the fixed, encoded standard for comparison can be used, and there is no need to resort to the context during interpretation of the adjective. Absolute adjectives with a totally closed scale structure are known as total absolute adjectives, while those with a partially open scale structure are called partial absolute adjectives (Burnett, 2012).

Following Sassoon and Toledo's (2011, p 3) account of the argument, and Siegel (1979), 'for-phrases' also seem to support the typology proposed by Kennedy and McNally. With relative adjectives, it is possible to say 'x is y for a z', e.g., *Cleo is small for a cat*. However, with *dirty*, *clean*, and *full*, which have at least one closed endpoint, these 'for-phrases' appear odd.

(5a) ? This cloth is dirty for a tablecloth.

(5b) ? This needle is clean for a sewing needle.

(5c) ? This glass is full for an espresso glass.

Another piece of argumentation that is put forward in favour of the context-insensitivity of so-called absolute adjectives is that certain adjectives seem to give rise to particular inference patterns, and the patterns of inference that emerge are said to be tied to the standard of comparison referred to in the interpretation of the adjective (see Sassoon & Toledo, 2011, p 4). For example, Sassoon and Toledo (*ibid*.) explain that a case such as *Florentine is bigger than Egglentine* entails that Florentine is indeed bigger than Egglentine, but it does not entail that either Florentine or Egglentine are bigger than any contextually-determined midpoint standard on a scale, i.e., that either of them are larger than some sort of norm. However, with so-called absolute adjectives, there appear to be entailments involving another entity in a comparison exceeding a minimum standard or failing to meet a maximum standard. For example, if you say *Your house is dirtier than my house*, there is thought to be an entailment that my house is dirty because it exceeds the minimum standard for dirtiness. If you say *Your bottle is emptier than my bottle*, there is said to be an entailment that your bottle is not empty. The presence of these inference patterns seems to suggest that the interpretation of absolute adjectives does not require any reference to a context.

Finally, it has been proposed that relative adjectives and absolute adjectives are distinguishable by means of a prosodic test which involves inserting them into a contrastive setting. It is said that this sounds acceptable for relative adjectives, but not for absolute ones (Unger, 1975; Kennedy, 2007):

(6a) That house is BIG, but it could be bigger.

(6b) That dress is SPOTTy, ? but it could be spottier.

This 'accentuation test' forces a 'precise interpretation' for absolute adjectives, which then leads to a contradiction, thought to render the utterance unacceptable (Burnett, 2012, p 8).

As Sassoon and Toledo (2011) explain, the proposed differences between relative and absolute adjectives with respect to context-sensitivity have been developed by Kennedy (2007) into a theory of vagueness. For Kennedy (2007), relative adjectives are vague, while

any adjective drawing on a fixed endpoint requires no reference to the context for interpretation and is, therefore, not vague. Kennedy's (2007) approach relies on an 'interpretive economy' that requires truth conditions to be worked out only on the basis of the conventional meaning of an expression, where possible, and he assumes that the context should only support an interpretation as a 'last resort'. As often is the case with approaches to phenomena that attribute a large portion of the responsibility for the interpretation to the linguistic semantics, the point at which pragmatics and the context are required to step in is not described. This may indicate that pragmatics has more of a role to play in the interpretation of adjectives than is typically thought. Traditionally, an expression is deemed vague if it has a fuzzy or imprecise extension, and the category of entities to which that expression can apply has borderline cases (Barker, 2002). Colour adjectives are the paradigm case of vague adjectives (Kortmann & Loebner, 2002). The term vagueness has long been the subject of discussion in linguistics and the philosophy of language (see Keefe, 2006) and I will not attempt an account of it here. However, it's worth noting that Kennedy (2007) considers vague expressions to be those that are contextually variable, and admit borderline cases, and he says that only relative adjectives can exhibit this vagueness. In Relevance Theory, vagueness of adjectives would be treated as cases of approximation (Wilson, 2014), and approximation involves the modulation of ad hoc concepts. If we recast what Kennedy says in relevance-theoretic terms, then the prediction would be that so-called relative adjectives are amenable to conceptual adjustment, while so-called absolute adjectives do not readily submit to modulation. Pre-theoretically, it seems like this picture might be correct. A door is either closed or not. An individual is either dead or alive. A dress is either spotty, or it isn't. Any difference between these two types of adjectives seems to hinge on their perceived inherent 'adjustability'.

5.2.1 The possible context-sensitivity of absolute adjectives

I have so far outlined a situation where gradable adjectives encode functions which lead to the mapping of a relation between an object and a degree on a scale for a dimension such as height, cost, bumpiness, closedness, etc. Reference to comparison classes and standards of comparison are required during the interpretation of gradable adjectives on such approaches. It has been proposed that some scales are part of the lexically encoded meaning of adjectives,

and a split has been proposed between relative adjectives, which are said to be interpreted with reference to the context, and absolute adjectives, which are said to be interpreted independent of the context. There is, however, some disagreement about whether gradable adjectives can actually be split on the basis of whether or not they are interpreted with respect to the context (see Sassoon & Toledo, 2011). Kennedy and McNally's (2005) approach would say that *dead* and *full* are absolute adjectives which are not interpreted with respect to the context. However, repetition examples (1c) and (1d) reproduced below suggest that this viewpoint may be too rigid.

(1c) That's one dead dead corpse.

(1d) I've got a full full schedule today.

In (1c) and (1d), it does seem that the repeated 'absolute' adjectives play at least some role in the communication of quite nuanced concepts DEAD* and FULL*, and there is a sense that the speakers here wish to communicate that the corpse is particularly or extremely dead, or that the schedule is very highly packed. If this is right, there is a suggestion that conceptual adjustment is at work in the interpretation of every kind of adjective.

Rotstein and Winter's (2004) and Sassoon and Toledo's (2011) work considers the suggestion that absolute adjectives *might be* context sensitive in the same way that relative adjectives clearly are. There are a number of reasons why they suggest this. One key argument against categorisations of adjectives on the basis of the scales they employ is that of so-called 'standard shifts'. A standard shift occurs when the standard of comparison that is supposedly employed in the interpretation of an adjective changes (Syrett *et al*., 2005). What counts as tall is going to differ in various contexts depending on what an entity is tall *in relation to*. I am 5"2. If I visited the planet Endor where the Ewok inhabitants are about a meter in height, I should be considered tall in comparison to them. Back home on earth, however, few would consider me tall with respect to my immediate con-specifics. It has been argued that there are no standard shifts for absolute adjectives because their interpretations are supposedly not context-specific. However, cases involving absolute adjectives and standard shifts do seem to exist. As Sassoon and Toledo (2011, p 5) note, following Cruse (1980), the standards of comparison for *dirty* and *clean* depend on what you are talking about. They (*ibid*.) explain

that the knife that you eat your dinner off would be considered clean under very different circumstances to those that would need to apply to felicitously call a surgical knife clean. Likewise, a tuxedo would be considered dirty if it had a couple of specks of dirt on it, whilst a child's t-shirt would be considered dirty if it had big splodges of mud on it. It seems like standard shifts might be the norm rather than a rarity (*ibid*.). Sassoon and Toledo (*ibid*.) claim that it is actually very difficult to find an instance where the absolute adjective *full* doesn't exhibit a shift, for example. When you return a hire car, its tank can be considered full if you could still fit in at least a few more drops of petrol, and a wine glass is generally considered full when the alcohol fills only half the glass (this example is due to McNally, 2011). In addition to this, if it were correct that absolute adjectives are not interpreted with respect to context, it would mean, as Sassoon and Toledo (2011, p 6) note, that saying *the tank is full* is the same as saying *the tank is completely full*. It would also mean that you could not acceptably say *the tank is full, but you can still top it off* (*ibid*.). However, it is perfectly fine to say this in some contexts. As such, perhaps so-called absolute adjectives are context-sensitive after all.

5.3 A relevance-theoretic account of repeated gradable adjectives

5.3.1 The lexical pragmatics and processing of gradable adjectives

Gradable adjectives are used as examples within lexical pragmatic discussions about word meaning, online concept construction, lexical entries, linguistic underdeterminacy, approximation and free enrichment, for example (Carston, 2002). *Green*, *happy*, *hot*, *closed* and *raw* have all appeared as examples, and they are considered to have a conceptual semantics in that they yield or 'point to' conceptual material which comes to feature as a constituent (or in a constituent) of a mental representation (Carston, 2002; Wilson, 2011). *Green*, *happ*y, and *hot* would be considered relative adjectives based on the discussion in the last section, while *closed* and *raw* would be analysed as absolute adjectives.

These expressions likely do not encode atomic concepts given the roles of modulation and the context in interpreting word meaning outlined in chapter two. We might instead want to say gradable adjectives encode some kind of pointer to addresses in memory where this information is stored (Carston, 2002; Wilson, 2011). The lexical information found there would include morphosyntactic information about what kind of expressions the adjectives

can modify, and any distributional constraints. For example, we might expect to find it stated that gradable adjectives can modify other adjectives, or information about which degree modifiers they can combine with. The encyclopaedic information would include world knowledge, memories, experiences, deduced knowledge, and so on. In the case of *hot*, for example, I might access scientific knowledge about temperatures (e.g., the temperature values of body heat or boiling water), and I might access memories about uncomfortable days travelling in Spain at the height of summer. It is also possible that information may be accessible concerning the type of occasion that the word *hot* has been used on before, and other metalinguistic knowledge. *Raw* would be given a similar treatment on this approach - when an individual processes the word *raw*, we might expect to find similar distributional and syntactic information, and it seems reasonable to expect that similar sorts of encyclopaedic information would be accessed.

If this is so far correct, then there is potentially no difference between relative and absolute adjectives in terms of the type of lexical and encyclopaedic information made accessible during the interpretation process. However, there is the matter of any semantically encoded functions or scales to deal with. According to Rayo (2011, p 3), semantic rules or things akin to semantic rules can be found associated with various 'grab bags' of information for particular expressions and which help to determine the entities to which a term can be applied. In Relevance Theory (Sperber & Wilson, 1995; Carston, 2002), it is considered that logical rules can be part of the information that is accessed when computing an ad hoc concept. As such, it is possible that a function may be stored for a particular word along the lines of the ones proposed in the semantic accounts. It may just be that each type of adjective indeed behaves differently in terms of the functions and/or the scales employed in interpretation, as the authors of many semantic accounts posit. However, another view that we could take is that 1) adjectives don't lexically encode a specific function for generating or computing scales, and/or 2) open/closed configurations of scales are not lexically encoded at all, and emerge in context, if at all. I consider that any (relevant) scales and the placement of a particular degree of a property on that scale is contextually determined, happening if and only if it is optimally relevant in the context. Comparison classes are also contextually determined, and only if it is relevant for the interpretation to compute a comparison.

In chapter three, §3.4, it was explained that contrastive interpretations were likely not lexically encoded by particular adjectives, at least for English. Sedivy *et al*. (1999) showed that intersective adjectives can be processed incrementally, and it was the possible or actual presence of a competitor referent that led to the emergence of a contrastive interpretation. Sedivy *et al*.'s (1999) study also addressed the interpretation of scalar adjectives, and found that these can also be interpreted incrementally, and that contextual information can be integrated into interpretations involving scalar adjectives even before the head is encountered. It should be noted that Sedivy *et al*. (1999, p 142) report it is thought that incremental processing of adjectives likely proceeds when 'the adjective fails to have a stable core meaning'. It is also useful to mention that they (*ibid*.) say that incrementality can proceed by establishing contrasts, and I presume also comparisons, not just by using information from the immediate visual context, but also by considering relationships between objects and their representations in memory. The fact that Sedivy *et al*. (1999) speak of 'core meanings' means, as we saw in the discussion of prototypes in chapter three, that the notion of accessibility could play a role in the interpretation of gradable adjectives. Nevertheless, the point is, since comparisons and contrasts can be made by accessing representations in memory, we do not have to rely solely on the linguistic context or the immediate physical environment to interpret gradable adjectives.

Sedivy *et al*. (1999, p 117) note that the typical way to interpret scalar adjectives such as *tall* is to refer to a comparison class, and they say that, often, the comparison class is determined by making use of information that is yielded by the head. However, they (*ibid*.) also say that it should be possible fix a value for a scalar adjective as soon as the comparison class becomes available - one does not need to necessarily process the head that supposedly yields a comparison class. It is possible to determine a comparison class with the early integration of contextually available information instead - and that contextual information, as I argued above, need not come from the linguistic context. That information could be any set of relevant manifest assumptions. Based on the findings and discussion in Sedivy *et al*.'s (1999) paper, it seems that not only can contrastive interpretations emerge early with contextual support, but it also seems reasonable to propose that any comparison classes, where it is relevant to make a comparison during interpretation, might also emerge early, incrementally, and with support from such contextual information.

If this is right, then I am able to argue that any comparisons, functions or scales that might be made use of when interpreting gradable adjectives can be emergent rather than linguistically encoded. Let us think again about what the semantic accounts say about what gradable adjectives encode, and how they are interpreted. Kennedy and McNally (2005, p 13) said there were three parameters which have to be *lexically encoded*: a set of degrees or measurement values, a dimension specific to the type of measurement, e.g., cost or height, and an ordering relation which is useful in distinguishing between pairs of antonyms (*ibid*.). Reference must also be made to some kind of comparison class or standard. Closed end points of scales must be lexically encoded (Kennedy, 2007). Kennedy and McNally (2005, p 8) state that the scalar properties of gradable adjectives are largely predictable from the entities to which they apply. We can only decide, for example, if something is tall if we know what is being called *tall*, and what it is considered tall in relation to. For relative adjectives, reference needs to be made to a standard of comparison that is contextually determined. For absolute adjectives, this will not be contextually determined. It is important to note that the examples supplied in Kennedy and McNally (2005) are mostly cases where a noun phrase appears before the adjective predicated of it, i.e., cases of postpositive modification. In other words, the way that the examples are structured ensures that the modifiee/head is presented before the adjective predicated of it. As such, it is easy to argue that the semantics of the head necessarily makes some contribution to the interpretation of the postpositive adjective. Nevertheless, the assumption on this approach is that gradable adjectives map their *arguments* onto degrees, and these ordered degrees form a scale, and gradable adjectives are often analysed *as relations between individuals and degrees*. I do think Kennedy and McNally (2005) are correct to say that scale-like properties can be predicted, in a sense, from the entities to which they apply. However, I am less convinced that the linguistic semantics of utterances hosting gradable adjectives must encode this information.

On this approach, when interpreting gradable adjectives, it seems as if a hearer first needs access to information yielded by the modifiee in order to compute any scalar properties of an adjectives. This type of semantic approach tends to assign quite a lot of the interpretive burden to the linguistic semantics, and does not generally allow pragmatic inference to intervene unless it is absolutely necessary (e.g, Kennedy, 2007). As such, we have to assume

that the interpretation of the adjective is dependent on information yielded by the head, and the suggestion is that this information is lexically decoded rather than pragmatically determined in context. On this view, there are two predictions that would follow concerning the interpretation of gradable adjectives:

(1) a hearer must encounter and process the modifiee in order to complete the processing of the adjective, and;
(2) there should be no contextual adjustment of adjectival meaning for absolute adjectives.

First, I deal with prediction 1 which seems to hold fine for cases such as (7a) and (7b) below.

(7a) My supervisor is tall.
(7b) That parrot is definitely deceased.

In (7a) and (7b), the hearer is able to process the entity of which the gradable adjective is predicated before he or she encounters the adjective, and needs to derive or access a scale, and compute a relation between an entity and a property - like in Kennedy and McNally's (2005) cases. However, we have to ask what would happen in cases such as (8a) and (8b) below, where the adjective is encountered before the noun:

(8a) She is a tall supervisor.
(8b) That is a deceased parrot.

You might simply say that the semantics of the adjectives have certain values that need to be filled in order for them to be interpreted, and you just have to wait until you encounter the head noun to obtain any information you need to predict scales and compute any relations. Only, when we take into account Sedivy *et al*.'s (1999) account of incremental processing and the role of context in the interpretation of scalar adjectives, it seems like you don't need to actually wait to encounter the modifiee to be able to assign reference using adjectives - the context seems able to supply plenty of information that allows reference assignment to take place very early on indeed. For example, consider this dialogue between a brother and sister talking about warm places to go on holiday:

(9) Chris: What sort of destinations are warm at this time of year?

Laura: Malta is a warm place this time of year.

Chris does not need to encounter the word *place* in Laura's utterance to interpret her *warm*. From his own question, the word *destination* is already salient, and the concept it comes to express and likely other related concepts are going to be activated to a greater degree than if that word had not been uttered. Chris is entitled to expect, all other things being equal, that Laura's utterance is an answer to his question, and is likely going to yield information about some kind of place. As soon as he encounters the adjective *warm* in a pre-modifying syntactic context, he need only hypothesise what is likely to be modified, and, given the context, and, in particular, the linguistic context, he will likely hypothesise that the concept that will be modified will correspond to some kind of destination or place. As such, Chris already begins computing a concept of WARM PLACE* in line with his expectations of optimal relevance, and when the head *place* is encountered, he does not then compute a concept, but simply confirms that his initial interpretive hypothesis was correct. In fact, it is possible to argue that as soon as Chris interprets *Malta is*, and hypothesises that Laura's utterance is an answer to his question, he may be making forward inferences about reference assignment, and may correctly assign reference without processing *any* of the adjective phrase in question. Thus, we can see that we don't need to process the modifiee first. In fact, we may not even need to process an adjective before embarking on reference assignment - we would simply make a hypothesis, and confirm this, or tweak what we end up with if our hypotheses are incorrect. If this is right, reference assignment involving gradable adjectives can proceed without even processing the adjective, never mind the head. This means that it is possible to argue that any scales, or functions for computing degrees or scales need not be lexically encoded, but can be emergent in context if relevant, as can any comparisons that enter the picture, and this assumption should hold for both relative and absolute adjectives. I return to this point shortly.

Let us now turn to the second prediction made above - there should be little or no contextual adjustment possible for so-called absolute adjectives. Or, more specifically, any standards of comparison involved in the interpretation of absolute adjectives are not derived in context, but are lexically encoded. First of all, it is not clear that the interpretation of every gradable adjective requires reference to a comparison class. I think there is a sense in example (9) that it's mainly only relevant how warm Malta is at this time of year - no other destinations are

salient. We might, of course, argue that Chris' utterance forces the recovery of a comparison class. However, that would mean that the comparison class is already salient in the linguistic context, and is not yielded from Laura's uttering of *warm*, which would be problematic for the formal accounts. Since I do not think that any gradable adjective encodes anything like a comparison class, or a function for deriving a comparison class, I cannot argue that absolute adjectives have the same semantics as relative ones from the point of view of the formal accounts above, and claim that absolute adjectives also encode functions for the recovery of contextually-determined comparison classes. However, I can argue that (1) absolute adjectives are also subject to early, contextually-supported interpretation, just like *warm* in (9) above, and that (2) they are, in many cases, also extremely amenable to contextual adjustment.

Let us now turn to how a gradable adjective like *warm* is interpreted. It is now necessary to discuss how repeated relative gradable adjectives contribute to ad hoc concept construction, as this discussion is vital for explaining why absolute adjectives are also interpreted with respect to the context. Where adjectives are involved, the hearer undertakes a process of modification (in this case, narrowing) in pursuit of reference assignment. On processing *warm*, a hearer is directed to a mental address where lexical, encyclopaedic and logical information is stored that enables her to construct an ad hoc concept in context and in line with her expectation of optimal relevance. The question is if and how comparison classes and degrees become involved in the interpretation. Clearly, reference assignment involving adjectives is compositional (see the discussion of *green salad* in chapter three, §3.5.4). The concept WARM PLACE* will contain conceptual material found under the conceptual addresses for *warm* and under *place*. However, I don't think the semantics of *warm* contributes a pre-defined degree or a scale to utterance interpretation - it simply provides evidence to scaffold the construction of a concept of a referent WARM PLACE*. The ontological status of degrees is not clear. It is something that it is difficult to introspect on. Certainly, the ad hoc concept WARM PLACE* may have a degree built into it or, when called to mind, seem to be something we can order with respect to temperature. Perhaps any sense of a degree is just perceptible post hoc. *Warm* simply gives access to a mental address whereby optimally relevant information about warm things can be taken and built into a concept of WARM PLACE*, and this is driven by the context and our expectations of relevance.

We know that *warm* can be applied to the Arctic at 8 degrees, but it's not information about this type of warm that we would apply to Malta in the context of (9), as it would not be relevant to do this. We would select information about previous holidays to Malta, memories of TV adverts about Malta, perhaps images of the sun, or impressions of temperature - such things can be made accessible during the interpretation of conceptual items, after all. What we can do, however, afterwards, is judge that the entity referred to by *warm place* is warmer than, say, the Arctic in February, and it is clear that different concepts of WARM PLACE* could be ordered post hoc, and, as such, give rise to different logical implications. Any degrees built into an ad hoc concept are not simply recovered from the semantics and necessarily built into the compositionally more complex concept WARM PLACE*; they would emerge if and when it is relevant for us to compute a degree.

5.3.2 Interpreting repeated gradable adjectives

Single gradable adjectives constrain narrowing during ad hoc concept construction in reference assignment. Repeated adjectives can also constrain narrowing during ad hoc concept construction. I think they can, in some cases, iconically lead to a degree or an extent being built into a concept if relevant, or, considered another way, the number of repeated tokens provides evidence for how much conceptual adjustment should take place, whether that involves a degree or not.

Let's think about this with respect to (10a), (10b) and (10c) below.

(10a) We went|on a long long walk.|
(10b) We went|on a long long long walk.|
(10c) We went|on a long long long long walk.|

It's clear that the difference in interpretation between (10a), (10b) and (10c) has something to do with how long the hearer determines the walk to actually be. What is relevant is the length of the walk that the speaker intends to communicate. In context, the optimally relevant processing strategy is to assume that the repeated adjectives encourage the hearer to compute the degree to which the walk is long, and we might argue there is an iconic relationship between the number of tokens produced, and the extent to which the walk is

determined to be long. In this case, we might say that a degree of length is computed which is built into a concept LONG LONG LONG LONG WALK*, for example, and the repetition of the same adjective within the same intonation group 'traps' the hearer at a particular mental address, and the prosodic grouping of the tokens within the same intonation boundary suggests that they should be taken together as evidence about the extent to which that concept is to be modulated, i.e., the repeated tokens are to be interpreted as part of a single unit of relevance. The hearer is encouraged to *expend effort* on *further narrowing* the concept, which has the effect of increasing the extent to which the walk is considered long. I come back to this later to determine what, if any, element of *showing* is at work in such cases.

At this point, I focus the discussion on total absolute adjectives. Burnett (2012) has reviewed the literature on gradable adjectives, and has subjected utterances containing partial absolute adjectives such as *wet* and *sick* to various linguistic tests. She concluded that they behave the same way as so-called relative adjectives. The reader is advised to consult this paper for a fuller treatment of the issue. However, as we saw, there is a consensus in the semantics literature that total absolute adjectives are not interpreted with respect to the context, and it is this perspective that I challenge. It seems clear to me that there are many common, everyday cases of so-called total absolute adjectives where the resulting ad hoc concept is very fine-grained indeed, and conceptual adjustment seems to be able to proceed 'beyond', in a sense, what we might expect to be the closed end of a scale.

For one, we can see that there are some linguistic tests that are passed by absolute adjectives with suitable contextual support. We've already seen, for example, that *full* is more adjustable than some semantic accounts predict, e.g.:

> (11) The gas tank is full, but you can still top it off. It's not completely full yet. (Example due to Sassoon & Toledo, 2011, p 6).
> (1d) I've got a full full schedule today.

Dead seems to be equally adjustable, too, even though we think of *dead* as being a prototypical absolute adjective.

> (12) [Consultant to trainee doctor during a teaching session.]

Your patient is dead, but she could be more dead. There's still a little brain activity registering.

(1c) (reproduced from above) That's one dead dead corpse.

The absolute adjectives *full* and *dead* are repeatable in the same way as *long* is, and they can be given a similar treatment to relative adjectives in terms of their contributions to interpretations. In order to communicate a very highly adjusted concept of *full* and *dead*, multiple repetitions of the adjectives are possible within the same intonation boundary. If the concepts expressed by these words were not adjustable in context, their repetition would likely not occur as it would put the hearer to effort for little or no cognitive reward.

Let's think more about *dead*. Normally, dead is dead. We usually think of it as an either/or - you're either dead or you're not. On the semantic accounts, if you're dead, the scale involved in interpretation is closed, and you don't need to refer to the context to work out how dead a person is - you're just dead. It seems to me that there's a lot more flexibility in these either/or concepts than one might first think. The adjustment may be effortful for the hearer, but the payoff is a very nuanced concept, which is going to result in its own very particular implicatures to boot. Consider (13):

(13)

Mortician Jones: The new guy's in a bad way.

Mortician Toll: Yep - That's one dead dead dead dead corpse.

In a similar scenario to (13) above, we can imagine that what Mortician Toll wishes to communicate is a concept DEAD* that is highly adjusted to go further than simply picking out entities whose hearts have stopped beating - this corpse is bloated, decomposing, rigid - he's a very particular type of dead, and a type of dead that is more dead than any kind of core, predetermined, absolute sense. There is every reason, in the interests of economy, to suppose that interpretation proceeds the same way for *dead* as it does for, say, *long*. During processing, the intonational grouping and the repetition of the same adjective over and over restricts Mortician Jones to the same conceptual address, and the only optimally relevant processing strategy for her to adopt is to assume that the repeated tokens are clues as to how

and how much the concept DEAD should be narrowed in that context. In other words, the repetitions provide evidence for how the concept is to be modulated.

If absolute adjectives are also adjustable, and there are no guaranteed 'endpoints' encoded in their semantics it does still remain to be explained why:

(1) Particular entailment patterns put forward by various researchers sometimes emerge;

(2) Certain absolute adjectives appear totally unacceptable in certain contexts, regardless of whether or not repetition is involved.

First of all, it's important to note that I'm not claiming that it's possible to repeat every kind of gradable adjective more than once within the same intonation group. Acceptability depends on multiple factors. Acceptability can be a result of phonotactic constraints, and certain words will be too long or too phonologically complex to be easily repeated within the same intonation boundary. I will not be considering such cases here. The focus here is on certain entailments or acceptability judgments associated with particular absolute adjectives in the literature. The issue is not that certain entailments or acceptability judgments never arise. It's clear that there are many circumstances when *dead* entails *not alive*, for example, or where something which is described as *full* is indeed completely full, and cannot be filled up more. However, in linguistics research, particularly in formal semantics, the aim is to abstract away from context to work out what the semantics of an expression contributes to an interpretation. As such, the kind of contexts found in linguistics papers tend to be impoverished - the focus is primarily on the immediate linguistic context, and little reference is made to the circumstances of uttering, or the (imagined) communicative intentions of any speakers.

In such impoverished contexts, we have to think about what is accessed when the reader tries to interpret the adjective in question. Without any guidance from a rich context, what the reader will access is whatever information is most accessible to him/her from the linguistic context, and that is likely to be what some semanticists would call a 'core' meaning, or what some would call a prototype - think back to the discussion in chapter three, §3.4.3. It could be argued that the scale endpoints supposedly encoded by absolute adjectives have a

prototypical flavour in the Roschian sense. However, in Relevance Theory, we might say that, in the absence of contextual support, what might be most accessible is a concept of, say, *dead* or *full* that has already been deployed on other occasions of use, and where the encyclopaedic information is organised such that the state of being alive or not, or full or not is particularly foregrounded because it is the most *accessible* information. As such, all other things being equal, a hearer would access a kind of concept that would seem quite rigid, and prototypical, and would not be adjusted much (as there is no strong context to encourage this). This is exactly the type of concept that would yield fairly robust entailment patterns, and seem to guarantee the unacceptability of certain expressions in certain (linguistic) contexts. Couch these utterances in the right context, and one including salient speaker intentions, and these patterns can disappear, as we have seen above. As such, the patterns, the judgments and the entailments that are put forward as evidence for the context-dependence of absolute adjectives need not necessarily be present. This, coupled with the repetition data, suggests that all gradable adjectives can be interpreted in the same way, and are both adjustable in context with neither encoding any kind of scale, endpoint, comparisons, or special functions.

An objection that could be raised here concerns whether these repeated adjectives are cases of reduplication. This cannot be ruled out on the basis of prosody and morphophonology at this point. There is a sense in which these repetitions can feel like one long word. It has been suggested that repeated adjectives or repeated intensifiers could be cases of reduplication (Pullum, 2006), and that, perhaps, what is yielded by that reduplication is intensification (Pullum, 2006). Based on my discussion of reduplication in chapter three, §3.2, it was determined that reduplication is a grammatically mandated process where its contribution to the interpretation is delivered by the linguistic semantics - i.e., reduplication must linguistically encode its meaning or function. I will not go into great detail here, but the issue of reduplication will come up again in the discussion of degree modifiers later in this chapter, and I will return to the matter of reduplication and repeated modifiers there.

5.3.3 Repeated gradable adjectives as cases of *showing*

One gradable adjective does not a case of *showing* make. What a single gradable adjective does make is a case of *saying* - production of indirect linguistically coded evidence for an ad

hoc concept that is intended for communication. However, if you add in repeated instances of the same adjective within the same intonation group, this represents a case of *showing*. In the same way as was suggested for other repetition phenomena, the speaker displays the form he chooses by exploiting its markedness, and the realisation that he manifestly did not utter something else that he might have been expected to in the context. In the case of repeated gradable adjectives, I think it is reasonable to say the speaker might have been expected to utter a degree modifier such as *really* or *very*, for example, or other distinct adjectives. In this way, the speaker uttering repeated gradable adjectives 'plays on' or displays the repeated forms, attracting attention to them, and leading the hearer to reason from the form he *shows* to the first layer of information to be recovered. There is a further way that this form can be *shown*, which has a two-fold consequence - place the repeated forms within the same intonation group. On one hand, this suggests that whatever is within the intonation group should be treated as a single unit of relevance to be processed together in pursuit of effects, and, on the other, the boundaries of the intonation group serve to further highlight and *show* the repeated forms, just like when I said in the last chapter that you could draw a circle around an amount on a bank statement to *show* and draw attention to that number. What do the repeated forms do, and what is the interaction with the context in cases like (13) above? The first adjective unlocks a conceptual address and the speaker is 'hemmed in' there to fixate on the repetitions by the intonation boundary; subsequent repetitions draw attention to themselves by virtue of their marked form, and the only optimally relevant way to treat them in context is to have the number of tokens iconically suggest the amount of narrowing required in pursuit of ad hoc concept construction.

Is this kind of repetition stylistic, and does this kind of repetition represent a case of *indeterminate showing*? This kind of repetition has stylistic effects in the way that metaphor does, as we saw in chapter two, §2.7. They arise from the mutual adjustment of explicit and implicit content. In (13), DEAD* results in the communication of particular weak implicatures recovered on Mortician Jones' own responsibility and which emerge as a result of her adjusting the so-called absolute adjective *dead* in the way suggested by the context, just as the audience of *Juliet is the sun* is encouraged to recover weak implicatures from computing a very specific SUN* concept in context. However, the Shakespeare case would be treated as a case of *indeterminate saying*, since all the evidence provided for the vague interpretation is

linguistic and fairly indirect. Although the gradable adjectives which are repeated are linguistic material, they together do not represent a case of *saying*. The original adjective does, but the subsequent repeated tokens have attention drawn to them by the speaker, and the only optimally relevant way to interpret them (as you are led to fixate on them by the intonation boundaries) is to treat them as iconically suggesting how much adjustment should be invested in ad hoc concept construction. The repetition *shows*. The concept computed in such cases is also weakly communicated because there is repetition rather than the uttering of additional distinct adjectives or intensifiers. Moreover, the implicatures that emerge from the mutual adjustment of explicit and implicit content are also weakly communicated, and I don't think the import of gradable adjectives repeated within the same intonation boundary can be properly paraphrased without loss of effects. For this reason, I would say cases like (13) are cases of *indeterminate showing*.

5.3.4 The additional effects of repeated adjectives within intonation groups

Repeated adjectives within a single intonation group can result in more than encouraging a hearer to *expend* more effort on *further adjusting* a concept. Put another way, there may be a 'tipping point' where the markedness and effort of processing multiple adjectives encourages the hearer to adopt a different or additional processing strategy. At this contextually determined 'tipping point', the repetitions serve less to suggest how a particular concept and attendant poetic effects might be recovered, and, instead, are judged to be indicative of an attitude, or the emotional state of the speaker. In other words, such repetitions increasingly draw attention to themselves, and scaffold the recovery of an attitude, or suggest how a representation of an emotional state is to be calibrated. Consider (14a) and (14b):

> (14a) [Laura is complaining about a walk that she was taken on by her brother on her only day off.]
> Laura: We went for a long long long long LONG walk.
>
> (14b) [Dave's football team has just been beaten 4-0. He is devastated, as this means his beloved team will be relegated.]

Dave: It's an awful awful AWful result.

In both (14a) and (14b), the repetitions contribute to the recovery of very particular ad hoc concepts, e.g., a walk that is of a length that one could be annoyed about it, or a result that is so terrible that it will have a significant impact on a team's standing. However, in (14a), we also recover an attitude, and in (14b), we also recover an emotional state. We can imagine that Laura, who is clearly fed up in (14a), sighs before she speaks, speaks quite slowly, and produces alongside her utterance iterated circular-motion hand gestures indicative of long duration. Faced with this paralinguistic evidence, the hearer may find it optimally relevant to process the very marked multiple repetition tokens as indicative of the extent and nature of Laura's negative attitude to the walk. In the case of (14b), Dave might turn the corners of his mouth down, and if the nucleus of his utterance is produced with marked pitch and volume, this may provide evidence about the extent of his emotional state. What we have, then, is repeated adjectives that still contribute to the explicit content through ad hoc concept construction, but which can also be exploited and coupled with other behaviour for extra effects, just as we saw with X-x in chapter three. In the case of (14a), we end up with an attitude, which would also be part of the explicit content. However, in the case of (14b), the extra effects fall on the implicit side of communication.

What makes the hearers of (14a) and (14b) decide to adopt the additional processing strategy of computing an attitude or an emotional state? High token number (all other things being equal) within the context, and a context where there is paralinguistic support for recovering an attitude or emotional state. I would suggest here that these repeated tokens are cases of quite ostensive *showing*, and ones that also have indeterminate effects. It will be hard to paraphrase what Laura thinks and feels in (14a), and effects would be lost in the attempt. The same applies to (14b) - we cannot paraphrase Dave's emotional state at all. Thus, (14a) and (14b) are also cases of *indeterminate showing*. These cases are 9s on Sperber and Wilson's (2015) two-dimensional space, and they are interesting because they provide information about the interaction between repetition and intonation boundaries. Essentially, we have repetitions of gradable adjectives that can contribute to explicit and implicit content, or both, but which look the same from the point of view of intonational grouping. It seems as if I may

need to abandon any thoughts that intonation boundaries themselves guarantee the recovery of content at a specific level.

5.3.5 Intensification, emphasis and repeated gradable adjectives

Bolinger (1972) said that there is a sense in which any modifying expression can have an intensificational flavour to it, i.e., any adjective, for example, might be associated with intensification. What I propose is that all adjectives are necessarily associated with *adjustment* rather than necessarily associated with intensification. Ad hoc concepts can be adjusted *a little*, or they can be adjusted *a lot* when there is sufficient evidence in context that this is required to satisfy your expectation of relevance, e.g., adjectives repeated in the same intonation group. What if intensification, at least in ad hoc concept construction, is simply (the perception of) a lot of modulation? Intensification at the word level could be seen as the extreme end of a continuum of processing between adjustment and intensification, and would be recognised, post hoc, as a process of 'a lot more' narrowing. If this is correct, this type of intensification would seem qualitatively different to the intensification undertaken in pursuit of implicit content in the last chapter. The intensification there depended on re-activation and greater activation of sets of weakly communicated implicatures. Here, we just have 'more narrowing' in ad hoc concept construction in the cases where the point of the repetition is mainly to constrain the identification of an ad hoc concept. It may be that intensification always involves 'more' of something in processing, but that there are different kinds of it. Are speakers who employ multiple repeated gradable adjectives within intonation groups in any sense emphatic? I think in many cases, yes. Consider (14b) again:

> (14b) [Dave's football team has just been beaten 4-0. He is devastated, as this means his beloved team will be relegated.]
> Dave: It's an awful awful AWful result.

Consider the prosodic analysis of (14b) below in figure 8:

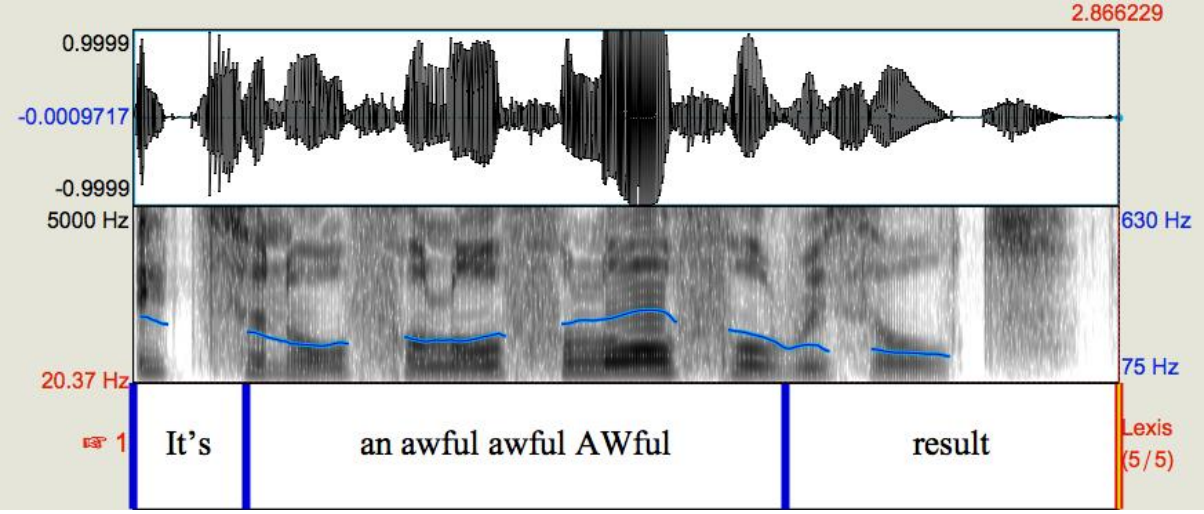

Figure 8 Prosodic analysis for 'It's an awful awful awful result'

Dave openly and deliberately frowns, makes his nucleus on the third *awful* prosodically salient, and repeats several tokens of *awful* within the same intonation group. He may be somewhat loud and the nucleus exhibits high intensity. All of these communicate behaviours here *show*. Dave's *showing* behaviour around and involving the repetition represents a cluster of *very ostensive* communicative behaviours. It would seem appropriate to call his *showing* behaviour emphatic.

5.3.6 Solving a puzzle

Finally, before moving on, I am now able to explain why you can utter (15a), and (15b), but not (15c).

(15a) We went for a long long walk.

(15b) We went for a long, long walk.

(15c) * The walk we went on was long long.

If adjectives repeated within the same intonation group encourage the recovery of a highly adjusted concept, then (15c) is ruled out on the grounds of processing effort. Let's imagine that (15c) is uttered in a context where it is manifest that the speaker went on holiday to the Peak District, where walking is popular. Assumptions about long walks are already marginally more salient than they ordinarily would be. As soon as the hearer processes *the walk* in this context, she makes forward predictions about ad hoc concept construction, and this will be guided by some of her assumptions about walking in the Peak District. When she encounters

long, it is likely that she has already computed a concept LONG* that satisfies her expectation of relevance in that context and, so, it would put the hearer to extra effort for no extra effects to have her adjust that concept even further, which is what the repetition in (15c) might be thought to do. The repetition in (15a) is acceptable, however, as is (15b), because, in that particular case, the hearer is entitled to assume that she can revisit her existing interpretation to derive further effects on the basis of the intonational grouping.

5.4 The repetition of degree modifiers *really* and *very*

5.4.1 Introducing (repeated) degree modifiers

In addition to gradable adjectives, other modifying expressions can also be repeated within the same intonation group. For example:

(16a) I **really really** love you.
(16b) I'm **very very** cross.

Sometimes, the number of tokens of the repeated expression is substantial, as in (17a) and (17b):

(17a) Would I recommend postgraduate study? It's **very very very very** HARD.
(17b) I'm **really really really really really** BORED.

It has been suggested in a blog post by Geoffrey Pullum (2006) that the number of repeated tokens should be a signal of the degree of intensification required. However, Pullum was writing about repeated degree modifiers separated by intonation boundaries, and does not supply a full account of how this works in that particular blog post. Nevertheless, Pullum (2006) is quite clear that the process at work here is one of intensification, not emphasis, but he still appears to interchange the terms, and he also does not provide a definition of what he means by them. Nevertheless, I will be adapting Pullum's insight to repeated degree modifiers *within* intonation groups, and the aim here is to more fully explain how they achieve relevance.

The repeated expressions in (16a-16b) and (17a-17b) are generally accepted to be degree modifiers playing a role in *boosting*, *maximising*, or somehow *intensifying* a (degree of a) property. The expressions *really* and *very* in (16a-16b) and (17a-17b) above are referred to as intensifiers (Bolinger, 1972), but the range of expressions called intensifiers are not usually thought to constitute a lexical category of their own. These particular degree modifiers are generally treated as adverbs (in English), sometimes modifying verbs, but modifying other adjectives and adverbs in particular (Huddleston & Pullum, 2010, p 563). Beltrama and Bochnak (2015) explain that intensifiers are expressions which boost a degree of a property. Quirk *et al*. (1985) thought that intensifiers were expressions that scaled a property upwards from some kind of assumed norm. Kennedy and McNally (2005) consider that the semantic effect of intensifiers such as *very* is to 'adjust' any contextually determined standard of comparison when paired with relative adjectives. Huddleston and Pullum (2010, p 1165) consider that intensifiers are any degree adverbs which mean something akin to 'highly, very, extremely'. However, they (*ibid*.) have also commented that it may not be possible to posit a unitary class of expressions called intensifiers.

It is sometimes thought that intensifiers are quite semantically bleached in some respects (Bolinger, 1972). Nevertheless, it is clear that intensifiers do play a role in the recovery of propositional content, and there is a good deal of work in semantics and syntax approaching the topic from a formal perspective. In this section, the focus is on how single and repeated degree modifiers are interpreted, concentrating on the semantics and pragmatics of certain so-called intensifiers, and focussing mainly on cases where repeated intensifiers contribute to the identification of explicit content. However, in some contexts, there is an undeniable emotional or expressive element to the interpretation of some multiply repeated intensifiers, and this is addressed in due course.

5.4.2 The (possible) syntax of repeated intensifiers

Huddleston and Pullum (2010, pp. 542-572) explain that multiple modifiers of this kind can be analysed in two ways syntactically, exhibiting either stacked modification or sub-modification, e.g.:

(18a) He [really [quite [rarely [submits his work on time]]]]. (stacked modification)

(18b) He [[really quite rarely] submits his work on time]. (sub-modification)

Stacked modification occurs when modifiers successively modify the head, while sub-modification involves simultaneous modification of the head (Payne & Huddleston, 2002). It is not clear whether *very very very* or *really really* repeated within the same intonation group would be treated as cases of stacked modification or sub-modification, e.g.:

(19a) It was [very [very [very [cold]]]].

(19b) It was [[very very very] cold].

The discussion of how prosody, semantics, pragmatics and context interact in the interpretation of repeated intensifiers should shed light on the syntactic structure of sentences that contain them.

5.4.3 The semantics and functions of intensifying degree modifiers

Many expressions have been called degree modifiers, e.g., *hardly*, *truly*, *fairly* (Huddleston & Pullum, 2010), and *perfect*, *little* and *a bit* (Bolinger, 1972), to name a few. For the discussion here, I restrict the analysis to two expressions - *really* and *very*. It is generally agreed that *really* and *very* are degree modifiers. Specifically, they are generally called intensifiers. Bolinger (1972, p 17) considers that intensification is a grammatical process that is associated with the semantic loading of adjectives, nouns and verbs. He (1972, p 18) says that intensification involves functional morphemes 'closer to the heart of grammar than...adjectives'. Bolinger (1972, p 17) notes 'I use the term *intensifier* for any device that scales a quality, whether up or down or somewhere between the two', and he (1972, p 20) also comments that 'intensification is the linguistic expression of exaggeration and depreciation'. Bolinger (1972, p 20) then goes on to say that there are several classes of lexical intensifiers which are categorised on the basis of how they interact with scales.

BOOSTERS - sit at the upper part of a scale and 'look up[wards]', e.g., *perfect*, *very, really*

COMPROMISERS - sit in the middle of a scale and 'try to look both ways at once', e.g., *fairly, quite, rather*

DIMINISHERS - sit at the lower end of a scale and 'look down[wards]', e.g., *little*

MINIMIZERS - sit at the lower end of scales, e.g., *a bit*

Very and *really* are 'boosters' considered to be involved in 'upward scaling' of degrees, as seen in (20a) and (20b) below:

(20a) The film was **really** good.

(20b) The film was **very** good.

Like Huddleston and Pullum, Bolinger (1972) also does not wish to 'exhaust' any category of intensifiers - he (1972, p 21) says it would be impossible to do so. Nevertheless, he (1972, p 23) says '[i]nvestigation will probably reveal that virtually any adverb modifying an adjective tends to have or develop an intensifying meaning' and, very interestingly, there is a footnote on the same page which suggests that *any modifier tends to be taken in an intensifying sense*, which I mentioned in the discussion of gradable adjectives above. In terms of the semantic contribution of any so-called intensifiers (in some contexts), Bolinger (1972, p 153-154) feels that they 'add little or no lexical meaning of their own', and, in some contexts, including instances where particular modifiers are repeated, multiple intensifiers can lead to subtle 'semantic repetition' which can sometimes seem redundant. Some of these comments are indicative that Bolinger sees many degree modifiers, including many of the so-called intensifiers, as perhaps somewhat semantically more bleached or less rich than, say, adjectives. Indeed, he does explicitly say that he considers many degree modifiers to be fairly 'grammaticized', and the reason for this is made clear in my analysis of the *conceptual* semantics of *really* and *very*.

Let us now turn to what Kennedy and McNally have to say about degree modification, and intensification. McNally (2016, p 21) says that degree modifiers combine with expressions that 'describe' gradable properties, and 'provide information about the degree to which that property holds of its argument'. Here, McNally (*ibid*.) notes that intensifiers are classed as degree modifiers alongside other expressions, including measure phrases, some manner verbs, comparatives, and adjectives which indicate 'some sort of extreme size'. Degree

morphemes, including degree modifiers, 'denote functions from (gradable) adjective meanings to properties of individuals' and their job is to 'saturate the degree argument of the adjective' (Kennedy & McNally, 2005, p 44). The syntactic and semantic distributions of degree modifiers are considered to be explained by the nature of scales associated with the expressions they modify (open/closed), and by whether standards of comparison for a modified gradable adjective are absolute or relative along the lines set out earlier in the chapter. The semantics of degree modifiers are said to be 'sensitive' to these features (Kennedy & McNally, 2005, p 8). This suggests that the semantics of certain degree modifiers, including the intensifiers I analyse, must encode some means of 'looking for' gradable predicates which they can legally combine with. Moreover, if they encode any functions for predicating degrees of a property of an individual, the semantics of particular degree modifiers will specify the kind of inputs they can compute over, and will also restrict the kind of outputs that they can produce. For example, Kennedy and McNally (2005, pp. 38-46) say that the degree modifier *completely* 'restricts the degree argument of a gradable adjective to being a maximum on the adjective's scale', while *much* is considered to 'look for' an adjective with a lower closed scale, e.g., *admired* or *regretted*.

For McNally (2016, p 23), what defines an expression as an intensifier concerns the constraints it imposes 'on the standard for the adjective [it combines] with and on the standard for the resulting modified expression'. She (2016, p 25) says that some intensifiers are iterable, recognising that you can have utterances containing 'very, very' (although, again, the intonation boundary is present in the written form). McNally (*ibid*.) goes on to say that intensifiers can be reduplicated, but this is not addressed in detail. Moreover, just as I set out above, she also identifies that it is not easy to see what syntactic structure is at work when intensifiers are 'iterated' - does the 'outermost' intensifier modify the rest of the phrase as a whole, or does it modify the next intensifier? Kennedy and McNally (2005, pp. 184-186) note that intensifiers are 'special' because they cannot be modified by anything apart from other intensifiers. As far as they are concerned, the semantic function of intensifiers is to 'manipulate' the standard for comparison which is considered, on these accounts, to play a crucial role in the interpretation of gradable adjectives (*ibid*.), and the semantic effect is the actual adjusting of these standards. Finally, since certain intensifiers, when combined with particular gradable adjectives, supposedly yield a predicate whose interpretation is derived

from an adjusted standard of comparison, further intensification is permitted (*ibid*.) - presumably, each time another intensifier is applied, the standard of comparison would simply be adjusted further upward. On Kennedy and McNally's (2005) account, repeated intensifiers would have to contribute to an incremental, recursive process where a new standard is computed each time an intensifier is repeated. However, they appear to mention only those cases where there is an intonation boundary between 'iterated' intensifiers.

The research I reviewed regarding the interpretation and functions of degree modifiers indicates that their distribution is restricted by combination with particular types of adjectives. Recall that Kennedy and McNally (2005, p 8) said that the semantics of such expressions are sensitive to the features of particular adjectives. For example, it is said that *very* can only combine with particular types of modifiee. It is not considered possible to utter (21b) and (21c), but (21a) is fine:

(21a) Michelle is very happy now she is marrying Dave.
(21b) ? That coffee is very fantastic.
(21c) ? When I study at home, I often leave my back door very open.

Very, for example, is considered to combine with a modifiee where the relative standard of comparison can be raised by a particular amount. *Happy* is considered a relative adjective, and so *very* can 'scale up' the degree of happiness predicated of Michelle in (21a) for example. However, something that is fantastic might be considered to already be at the pinnacle of that quality. Likewise, doors are generally either open or they're not. As such, the prediction is that *very* cannot modify these concepts. As noted above, Kennedy and McNally would consider that certain intensifiers can only combine with certain types of adjectives with certain types of open/closed scale configurations. However, if, as I suggested in §5.3.1, there are no scales lexically encoded by certain adjectives, and no guaranteed standards of comparison are lexically encoded by them either, we cannot argue that the distribution of certain intensifiers is dependent on what is lexically encoded by the semantics of the adjectives they modify. That is not to say that certain patterns of distribution and certain acceptability judgments do not exist. I simply suggest that, perhaps, other non-semantic factors are necessarily at work.

Following Bolinger (1972), I consider that intensifying degree modifiers are somewhat semantically bleached, and are on the way, at least, to having a more grammatical flavour to them than 'content words' such as adjectives and verbs. In relevance-theoretic terms, I show they have a particular type of conceptual semantics before setting out how they contribute to an interpretation when produced in isolation before investigating the contribution they make when repeated within the same intonation group.

5.4.4 A relevance-theoretic account of the semantics and pragmatics of *really* and *very*

I noted above that I wanted to investigate the nature of the semantics of the expressions addressed in this chapter. In particular, I wanted to look at the very particular conceptual semantics of *really* and *very*, and establish that it is quite correct to treat them conceptually, even if, as Bolinger (1972) thought, they seem to have a bleached, grammatical flavour.

There are a number of diagnostics for distinguishing between conceptual and procedural expressions, as I explained in chapter two. One such test involves ease-of-translation - conceptual expressions are readily transformed into equivalents in other tongues, and this applies to *really* and *very*.

(22a) I am **really** tired.
Ich bin **wirklich** muede.
Je suis **vraiment** fatiguee.
(22b) Cleo is **very** beautiful.
Cleo ist **sehr** schoen.
Cleo est **tres** belle.

The data in (22a) and (22b) suggest that such expressions have a conceptual semantics (and they *do*). However, it is very difficult for speakers to pin down the meanings of *really* and *very* in isolation. This semantic ineffability *can* be indicative of procedural meaning, as we saw in chapter two, §2.6. Moreover, if an expression is procedural, speakers are often unable to explain its meaning using a conceptual paraphrase, and instead offer explanations of how the expression is used. In an (admittedly limited) survey of non-linguist acquaintances, native

speakers of English offered up the type of response that might indicate procedural meaning for *really* and *very*:

Me: What does 'really' mean?

Native speaker 1: It's a positive response to confirm something.

Native speaker 2: 'evidently'.

Native speaker 3: It means an emphasis to your point.

Native speaker 4: Emphatically 'something'.

Native speaker 5: It IS.

Me: What does 'very' mean?

Native speaker 1: It's something to reinforce something.

Native speaker 2: 'Definitely' and lots of emphasis on 'definitely'. It's an emphasis word.

Native speaker 3: It means a lot of.

Native speaker 4: More intensively 'something', whatever that something is.

Native speaker 5: Nothing - it is used to suggest something is 'more'.

On its own, the conceptual paraphrase test is not a reliable test for procedural meaning. It only targets encyclopaedic information associated with an expression. However, not all conceptual expressions have encyclopaedic information associated with them. *And* is a case in point - it is conceptual, because it enters into compositional representations but, beyond perhaps metalinguistic knowledge about particular occasions of use, it is not associated with any encyclopaedic information. *And* activates lexical and logical information associated with it, but has no encyclopaedic entry. What *and* contributes is a function for deriving a representation that contains a constituent that corresponds to two entities that have been combined together by means of *and*. Native speakers of English can have difficulty pinning down what *and* means in conceptual terms, and give similar usage-based answers to my informants above if asked what *and* means. If *and* can be conceptual, so can *really* or *very* in this way, and I argue these words *also* encode a function, but one that is constrained by considerations of relevance instead of very specific semantic constraints on inputs and outputs. Moreover, If *really* and *very* do not have encyclopaedic entries, this would explain

Bolinger's (1972) suggestion that they are semantically bleached, and have a flavour of the grammatical to them.

Let's suppose that (23a) and (23b) give rise to the recovery of the ad hoc concepts VERY HOT* and REALLY SWEET*, respectively.

> (23a) [Uttered in a garden in Manchester, in March, and the temperature is an unseasonal 14 degrees.]
> Luke: It's very hot.

> (23b) [Uttered after receiving a short lift from friend who was travelling in the same direction anyway.]
> Julia: You're really good.

As we saw in chapter two, §2.5.3, many linguistic expressions radically underdetermine what they communicate, and each expression is capable of communicating one of an indefinite range of concepts on each occasion of use. *Hot* can express concepts of hotness to do with physical attractiveness, spiciness, desirability or the illegality of stolen goods, as well as a potentially indefinite range of concepts concerning temperature - just as it communicates a particular temperature-related concept in (23a). *Good*, meanwhile, results in the same underdeterminacy, and can communicate concepts related to being kind (as in (23b)), being virtuous, being of decent quality, and so on. Clearly, from these examples, you cannot determine how hot *very hot* is until one has ascertained what kind of hotness we are dealing with. Likewise, one cannot work out how good *really good* is unless one first has a hypothesis about what type of goodness we are talking about. In other words, a hearer has to first start to construct a concept of HOT* or GOOD* before *really* and *very* can make their contribution to the interpretation. In the case of (23a), we must determine that *hot* communicates a concept of a level of warmth that corresponds to the currently, unseasonably warm March day. Luke is not communicating the type of hot one experiences in the Sahara Desert, for example. In the case of (23b), the hearer computes an ad hoc concept that can be predicated of a person who is generous enough to give a lift, but only when they are not going out of their way, and who is not being overly kind. In other words, before *very* and *really* come into play, it seems the hearer must undertake an initial step of narrowing in terms of adjectival

meaning so she has a hypothesis about a concept that can then enter into further modification.

There is a paradox here. Intensifiers appear before the adjectives that they modify. However, it appears that their interpretation depends on hypotheses about the interpretation of an expression that is anticipated to occur later in the speech stream. However, this need not be problematic. It is accepted in Relevance Theory that utterance interpretation proceeds online and in parallel, and that we can make backwards and forwards inferences, adjusting aspects of the interpretation that have gone before, whilst making predictions about what is to come, which, if needed, can be confirmed or tweaked in line with our expectations of relevance. Furthermore, as explained in chapter three, and in the discussion above, it is possible for an optimally relevant ad hoc concept to be computed before an adjective is even presented for processing when the context is rich enough. Consider the below scenario.

> (24) [Luke and David are sitting in their garden in Manchester. It is March, but the temperature is 14 degrees, and it is sunny. As it is March, they are both wearing heavy coats. Luke catches David's eye, looks up at the sky, wipes his brow, and sighs.]
> Luke: It's very hot.

When David processes Luke's utterance, it is already salient to him that it is warm, and likely particularly so for a March day in northern England. Moreover, it is mutually manifest to Luke and David that Luke has deliberately and communicatively looked at the sky, and it is mutually manifest to Luke and David that Luke is likely hot. As such, hypotheses about the coming utterance can be constructed as soon as *its* is interpreted. Given the context, this is not going to be taken as a possessive pronoun - David will begin to make hypotheses about an utterance based on a sentence containing a copula and an adjective of some kind, and an interpretation where *it* corresponds to the day in question will satisfy David's expectation of relevance in this context. Thus, it is plausible that David is already narrowing a concept of HOT* before he even encounters the intensifier *very*. As such, there will be a concept activated in David's mind that is capable of being an input to a function encoded by *very*, allowing *very* to provide its contribution to the interpretation before the adjective *hot* is actually encountered.

I consider that *very* and *really* encode a function which takes a narrowed ad hoc concept as its argument, and returns a further narrowed concept as its value. In the case of (24), for example, David would have started to narrow a concept HOT* that relates to weather, and can be predicated of days out in the garden. *Very* can compute over this and mandate additional narrowing. How is the path of that narrowing constrained? It can be taken care of by the assumption that hearers follow the Relevance Theoretic Comprehension Procedure, and stop processing as soon as they hit upon an interpretation that satisfies their expectation of relevance in context. All that is required is that the hearer is not put to gratuitous effort by having to further unnecessarily narrow the concept that is modified by the intensifier. Essentially, it is fine for a particular intensifier to combine with a particular adjective if and only if the further modifying that the function of the intensifier mandates would yield an adjusted concept that provides sufficient cognitive effects for the processing effort that this further adjustment costs. If this is correct, we can explain why certain modifier+adjective combinations are unacceptable or ungrammatical using a pragmatic explanation rather than a semantic one - the effort entailed by the additional narrowing is not offset by sufficient effects. Let's think again about (21b) and (21c):

(21b) ? That coffee is very fantastic.
(21c) ? When I study at home, I leave my windows very open.

Fantastic and *open*, all other things being equal, may not always be amenable to much further narrowing on top of the initial narrowing process. If something is fantastic, it's already extremely good, so, in many contexts, it is not appropriate to put the hearer to the effort of trying to narrow this concept further. Likewise, our experience of things that are open tells us that things are generally open or they are not. However, we *can* adjust this concept, although it would have to be in a very strong context, e.g.:

(25) [Two colleagues are discussing how stuffy it is in their university offices. They are annoyed because they have safety catches on their windows so they cannot open them very far.]
Professor Toll: I hate writing papers at work. It's so stuffy. These bloody catches!
Professor Jones: I know. When I study at home, I leave my windows VEry open.

Uttered in this context, *very open* is perfectly acceptable. In fact, the intensifier might even be prosodically highlighted to suggest that further narrowing is worthwhile, or to lend expressive effects to the interpretation. The concept of OPEN* here that is subject to further narrowing is one where the window is unrestricted by a safety catch, and the pane that opens outwards is fully opened, almost touching the wall on the side of the house. Thus, additional narrowing via an intensifier does not put Professor Toll to undue processing effort for no cognitive effects. The payoff will be a finely nuanced concept, which will be associated with its own particular implicatures. Let us now think about how *repeated* intensifiers interact with the recovery of explicit content.

5.4.5 Repeated intensifiers and the identification of explicit content

It was established above that *really* and *very* can be 'iterated' and modify each other successively. Possible syntactic structures for the examples in (1e) and (1f) are:

(1e) It's [really [really [exciting]] what we're doing for New Year.
(1f) Our kitten was abandoned. Cleo is [very [very [very [small]]]].

The syntactic structures in (1e) and (1f) here are stacked. The head is modified by each intensifier successively. I think this is what happens when degree modifiers are separated off from one another by intonation boundaries, but not when they occur within the same intonation boundary. It seems to me that intensifiers repeated within the same intonation boundary might behave in a similar way as repeated gradable adjectives do. The first adjective takes the hearer to a conceptual space for ad hoc concept construction, and the additional tokens, via *showing*, iconically suggest how much adjustment should be undertaken. For words like *really* and *very*, the linguistically encoded function mandates that an already narrowed ad hoc concept be further narrowed in the way I suggested above, and the repeated intensifiers then serve as evidence for how much narrowing it is optimally relevant to undertake in pursuit of modulating a concept. In other words, in terms of syntax, these would be cases of sub-modification, whereby all modifiers interact with the head at once - and this makes sense because they are all grouped together prosodically. Repeated intensifiers separated by intonation boundaries would instead be analysed as stacked. This highlights the

relationship between syntactic structure and prosodic grouping. Repeated intensifiers will also be a case of *indeterminate showing*, or number 9 on Sperber and Wilson's (2015) grid space. In the same way as for repeated gradable adjectives, the speaker displays the already ostensive repeated modifiers, further *showing* them by deliberately putting them inside a single intonation group. The more tokens you have, the more the repetition draws attention to itself. In terms of explicit content, with each repetition, the weakly communicated concept that the intensifiers modify becomes more weakly communicated still, as the hearer must adjust more and more on her own responsibility. Any attendant weak implicatures will become more weakly communicated as the hearer is 'penned in' by the intonation boundary to fixate on the repetition and 'wring out' more effects on her own responsibility. The stylistic effects here arise, as for repeated gradable adjectives, from the mutual adjustment of explicit and implicit content, just like with metaphor. The interpretation cannot easily be paraphrased, and so this is *indeterminate showing*.

5.4.6 A tipping point: repeated intensifiers and the communication of attitudes or emotional representations

Recall that, for the repeated gradable adjectives, in certain contexts, there seems to be a 'tipping point' whereby the number of tokens is judged to be particularly ostensive and particularly marked in context, and the hearer decides it is optimally relevant to adopt an additional or alternative processing strategy to merely computing an ad hoc concept. This also appears to be the case for some repeated *really*s and *very*s:

(26a) I'm really really really really REALLy pissed off.
(26b) I'm very very very VERy bored.

In (26a) and (26b), there are two possibilities. The first possibility is the speaker does not intend for this expressive or emotional element to be communicated, and merely intends for his deliberately repeated tokens to show how much adjustment should take place in the identification of explicit content. Here, the recovery of an emotional and possibly expressive or attitudinal aspect to the interpretation is accidental - a bad mood or extreme boredom can be inferred, but not intended for communication by the speaker. The second possibility is that

the speaker has deliberately and markedly produced so many tokens of the intensifier that the optimally relevant way to take these is as *showing* something about how to calibrate the speaker's emotional state. This could also proceed iconically. The more tokens, the more the degree of commitment to any higher level explicatures recovered, or the greater the extent of the emotional state calibrated.

5.4.7 Repeated gradable adjectives and intensifiers within single intonation groups are not reduplicative (in English)

In a similar way to with the gradable adjectives, here, I have made the case that the semantics of certain degree modifiers are neither as rich nor as rigid as previously supposed. However, this viewpoint leaves us with the same quandary as I raised at the end of §5.3.5 - are we not dealing with cases of reduplication instead of repetition? If I am right that these repeated adjectives and repeated intensifying degree modifiers *show*, then they cannot be cases of reduplication. If they were, they would necessarily be cases of *saying*. Reduplication is defined as a grammatically mandated process which requires that the meaning or function of the reduplication to be delivered by the linguistic semantics. In the cases reviewed so far in this chapter, this is not happening. It seems like the semantics of the first adjective or the first intensifier make a defined contribution to the interpretation, but no semantic material, including further functions, are contributed by the subsequent repetitions - they merely serve as non-coded evidence for suggesting what processing path is to be followed. Repeated within the same intonation group, these do not constitute cases of reduplication.

In this section, I focussed on the repetition of expressions called intensifiers. However, I said very little about the very thing that gives them their name - intensification. Nevertheless, in the section on gradable adjectives, taking inspiration from Bolinger, I did suggest that all intensification might be is a very perceptible amount of adjustment. Now, adjectives will always lead to at least a little narrowing. However, all other things being equal, intensifiers will generally lead to more narrowing than naked adjectives because they mandate narrowing on top of narrowing - and this is why they probably are more readily associated with intensification, or, as I see it, (noticeable) additional narrowing in pursuit of ad hoc concept construction. Intensification involving (repeated) modifiers therefore appears to be part of

processing rather than an effect of processing. It was noted above that researchers have been reluctant to posit the existence of a class of expressions called intensifiers. Based on my discussion, it may actually be possible to do this, at least for a small set of expressions. For any degree modifier that takes a narrowed ad hoc concept as its argument and returns a further narrowed concept as its value, we can say we have an intensifier.

Finally, several times I have drawn attention to comments of Bolinger's. I think that his comments about intensifiers 'scaling up' or 'scaling down' and about intensification being concerned with exaggeration and depreciation are potentially very important. If I'm right that intensification of the kind addressed in this chapter is just the perception of more narrowing, then this fits with what Bolinger says about intensifiers being able to get involved with interpretations that have a different quality than up, or bigger, or better (augmentation) or lead to movement along the upper part of a scale. They could also be involved in interpretations that 'scale down', or involve depreciation, or even diminuation. Whilst it is outside the scope of this PhD, it would be interesting to look at the morphology, semantics and pragmatics of morphemes that have been called diminutives and augmentatives to see if some of the insights gained here could shed light on the interpretation of these expressions as well.

5.5 The repetition of *yes* and *no* within the same intonation group

The polarity discourse particles *yes* and *no* can also be repeated within the same intonation group.

(1g) Rosie: Do you want to see the new James Bond film tonight?
Kelly: Yes yes yes. / No no no.

When *yes* or *no* are uttered in isolation as an answer to a yes-no question, they can be treated as discourse particles contributing to a higher-level explicature. They might encode a procedure for the recovery of higher-level explicatures. A higher-level explicature occurs when a proposition (or sub-propositional constituent) expressed by an utterance is embedded under a propositional attitude description, a speech act-type description or some

kind of comment on a lower-level, embedded proposition (see chapter two, §2.5.2). Let us consider how this might work for a modified version of example (1g) below:

(27) Rosie: Do you want to see the new James Bond film tonight?
Kelly: Yes, I do want to see the new James Bond Film tonight.

Rosie's choice of a yes/no question suggests to Kelly what type of information she would find relevant - a yes/no answer. She want to know whether *p* or not-*p* is the case, *p* being [KELLY WOULD LIKE TO SEE THE NEW JAMES BOND FILM TONIGHT]. In this context, it is reasonable for Rosie to assume that Kelly's utterance is an answer to her question. Certainly, one of the explicatures that Kelly communicated by way of her answer is that *p* is the case. She does want to see the film that evening. However, Kelly can also be understood to be asserting that she wants to see the film that evening, and assertions can be analysed as a kind of higher-level explicature associated with declarative utterances. *Yes* is involved, then, in the recovery of the higher-level explicature that Kelly is asserting the proposition *p* - she would like to see the film that night.

What we might say with regards to repeated instances of *yes* or *no* within intonation boundaries is that the analysis might proceed along the same lines as it did for the repetition of gradable adjectives and intensifiers within intonation groups. It seems as if additional instances of the same adjective or the same intensifier serve as evidence for *expending effort on further adjusting* a particular conceptual representation. In some cases of repeated *yes* and *no*, I suggest that the first instance of the expression triggers a procedure for the recovery of a higher-level explicature, but subsequent repetitions within the same intonation boundary encourage the hearer to adjust, or, rather, calibrate the degree of the strength of the assertion being made.

(1g) Rosie: Do you want to see the new James Bond film tonight?
Kelly: Yes yes yes. / No no no. (The former is uttered with high pitch, while the latter exhibits a low falling tone.)

In (1g), we can say that Rosie is encouraged by the repetitions of *yes* or *no* to expend processing effort on building into her interpretation the degree to which Kelly is making her

assertion. Again, as we saw for repeated adjectives, and for repeated intensifying degree modifiers, there seems to be this 'tipping point' whereby the number of tokens are subconsciously judged within a given context as being indicative of another processing strategy - one where a more emotional (aspect of an) interpretation may be highly relevant. We can imagine that repeated tokens of *yes* or *no*, particularly if accompanied by pitch in the upper portions of someone's pitch range, and gestures and facial expressions associated with excitement or anger, could induce a hearer to recover an emotional representation of the speaker. This might be recovered but not communicated by the speaker, but, clearly, it is possible for a speaker to deliberately construct their utterance so as to feature a (likely rather marked) set of repeated particles such as *yes* or *no*. In this case, the additional tokens would *show* in the same way as repeated gradable adjectives and intensifiers, and suggest that a hearer should invest an optimally relevant amount of additional effort in calibrating the extent or severity of a speaker's emotional state. Here, the hearer is encouraged to *expend* additional processing effort on *calibration*. This would be weak communication, and the interpretation would be indeterminate. It is not possible to paraphrase the import of repeated *yes* and *no* particles here, in both higher-explicature and emotional cases. Effects would be 'lost in translation'. These would also be cases of *indeterminate showing*. Are speakers repeating *yes* and *no* creating any sense of emphasis? If Kelly in (1g) were to gesture ostensively alongside her repeated *yes* particles, and speak in an excited voice, this, along with the repetition and the fact the repetition is further *shown* by the intonation boundaries would suggest that her behaviour is emphatic. We once again see a cluster of highly ostensive *showing* behaviours here. In terms of intensification, there is a suggestion that it can occur in pursuit of the construction of higher-level explicatures, *and* in pursuit of the calibration of emotional representations, which would fall on the implicit side of content. Intensification appears to be undertaking *more* of an optimally relevant processing strategy.

If we reflect on the discussion in chapter four, and the cases addressed so far in this chapter, we cannot confirm the hypothesis that I set out earlier in thesis: repetitions in adjacent intonation groups or non-adjacent intonation groups contribute to the recovery of effects at the implicit level, while repetitions within intonation groups contribute to the recovery of content at the explicit level. If there is any distinction which the boundary-internal/boundary-external distinction cross cuts, it cannot be the explicit/implicit distinction. For one, I have

shown that although the kinds of expressions addressed in this chapter do principally contribute to the identification of explicit content when repeated in the same intonation boundary, each type of expression addressed has also been associated with a kind of 'tipping' point whereby it can sometimes be optimally relevant to treat the repeated tokens as communicating something emotional. This would fall clearly on the implicit side of communication. What seems to play a part in the adoption of this different processing strategy is how marked and ostensive those repetitions are, and, as we have seen elsewhere in this thesis, markedness is not really an inherent property of an expression or a communicative act - it's calculated against the backdrop of the context. As such, it becomes very hard to maintain the hypothesis. Instead, there is an interaction between repetition, any paralinguistics, intonation boundaries *and* context which plays a part in a hearer judging whether it is optimally relevant to recover explicit or implicit effects, or some combination of the two. The intonation grouping itself does not guarantee what type of effects will be recovered in advance of deployment in an actual context.

5.6 The repetition of procedural items within groups or units

5.6.1 Introducing repeated procedural items repeated within a group

There are expressions which are considered at least 'borderline linguistic' and which, when repeated within the same intonation group, yield an expressive interpretation or reveal something about the internal state of the speaker as a matter of course. A conceptual interpretation is never recovered in these cases, e.g.:

(28a) Julie: Someone's keyed your car.
John: |Shitshitshit|.

(28b) [Kelly wakes up to find that her pet chicken has escaped for the third time and has gone missing.]
Neighbour: I think the chicken's gone walkies again.
Kelly: |Fuckfuckfuckfuckfuck|.

Moreover, certain interjections can be repeated in the same way, and also reveal something about how someone is feeling:

(29a) | Owowow! | That fucking hurt!

(29b) | Wowwowwow! | Did you see that huge flash?

These repetitions are clearly intended to communicate something about the internal state of the speaker. Expressive effects and/or representations of emotional states can be categorised as implicit content. As such, we definitely have to abandon the hypothesis about boundary-internal repetitions solely constraining the identification of explicit content. This again raises the question of what the nature of the interaction between intonation boundaries and repetition actually is, and I return to this in my concluding chapter. At this point, it is at least possible to conclude that some procedural expressions can be repeated within intonation boundaries - *yes* and *no* can be treated procedurally within Relevance Theory, as can the interjections *ow* and *wow* immediately above (see Wharton, 2009). The repeated expletives would also trigger emotion-reading procedures, and cannot be treated as communicating a concept. We have a picture, then, where (some) procedural expressions can be repeated within an intonation boundary to yield effects either at the explicit level, the implicit level, or both. Procedural expressions have been neglected in any relevance-theoretic discussion of repetition.

Although procedural expressions such as expressive *shit* and *fuck*, *yes*/*no* and even *well* can be repeated within an intonation boundary, not all expressions which have been analysed as procedural are necessarily repeatable in this way. The following would likely be interpreted as indicative of speaker disfluency:

(30a) |*She she she* is my friend.|

(30b) |Look, | it's *my my* book. |

(30c) |I spent all my wages. |However however, | I managed to go out.|

(30d) |*But but but* I didn't spend all my wages!|

This raises the question of why certain procedural expressions are repeatable within the same intonation boundary, while others might not be. To shed light on this question, I would like to turn to the stylistic repetition of a non-linguistic phenomenon which can, in some, but not all, cases, be given a procedural account - facial emoji.

As the examples addressed in chapter four and chapter five have involved linguistic or borderline linguistic expressions, it has been possible to talk about *intonation* boundaries. Emoji are nonverbal, and are pictorial, and, as such, cannot be analysed in the same way. However, it is possible to talk about them occurring within units. They often appear repeated in clusters with no linguistic or non-linguistic material in between, and, as I have observed, they mirror punctuation in that they often occur at the end of sentences/clauses/tweets/statuses. As such, it is going to be possible to draw comparisons between what happens when they are repeated within a unit, or within some kind of grouping. As I discuss how particular (repeated) emoji are interpreted, I will return to how procedural items such as expletives and interjections are processed in order to draw comparisons between the two.

5.6.2 The interpretation of emoji

Traditionally, emoticons are said to be pictograms which are made up of ASCII characters which are considered to indicate particular emotions, e.g., :) for happy and :(for sad. Many online anecdotal sources credit Scott E. Fahlman of Carnegie Mellon with their inception. He introduced emoticons to computer science online bulletin boards in 1982 to help people understand whether messages were meant to be taken seriously or not (Ptasynski *et al*., 2011, p 1159). However, there are earlier instances of emoticon use. Abraham Lincoln was anecdotally said to have produced a handwritten one in a letter, and there are examples from the 1800s (Ptasynski *et al*., 2011, p 1159). The word *emoticon* is a blend of *emotion* and *icon*, and the use of emoticons is prevalent in electronic means of communication such as texting, Facebook messenger, Whatsapp and informal emails. It is generally thought that emoticons are used to make up for the lack of paralinguistic features in such electronic media, e.g., prosody, gesture, and facial expressions (Walther & D'Addario, 2001; Ptasynski et al, 2011).

It is still possible, in all of these message formats, to physically type out a face icon made up of ASCII characters. However, these days, applications usually provide pre-constructed symbols that can be used instead of traditional pictograms, or, as with my mobile phone, these applications sometimes autocorrect manually typed emoticons into graphic symbols

instead. These facial emoticon symbols are now accepted as a part of the ever-growing set of pictorial characters which are known as emoji, and, as such, I will refer to them as emoji. Below are examples of face emoji.

Figure 9 A range of face emoji

A glance at the available emoji on any smartphone or messaging application reveals that most emoji are not faces - they are instead symbols that represent people, places, and a variety of other objects, e.g.:

Figure 10 A selection of non-face emoji

The word emoji is a 'false friend' in that it has nothing to do with the English word *emotion*. *Emoji* is loaned from Japanese, and means something like 'picture character' (Ryoko Sasamoto, personal communication, 2014). The first emoji was developed in the late 1990s by Shigetaka Kurita of NTT DoCoMo, a mobile phone operator in Japan (Vulliamy, 2015). Many of the first emoji were visual representations of facial expressions. Anecdotally, and drawing on experience, it seems there are over 800 emoji in regular existence across varying platforms.

Emoji can often be repeated in particular contexts within the same unit or 'grouping':

(31a) Morning baby 💋💋💋💋💋

(31b) I'm so fucking tired mate 😭😭😭😭

(31c) It's cold 😬😬😬

(31d) I fell asleep today.
😴😴😴zZzZzZzZzZ

We see, then, that both face and non-face emoji can be repeated, and we can see that for the data I have collected, these do seem to sit at the right-periphery of clauses or utterances, at least for English. Figure (31d) contains repeated emoji that appear on a new line (and so perhaps could be construed as a new 'unit' or 'turn'), which suggests that repeated emoji can function together as a unit of relevance.

Many emoji can be used iconically to represent entities in the real world. @emoticoncaselaw is a humorous Twitter account which recounts the legal facts of highly memorable landmark cases. Consider figures 11 and 12 below:

Figure 11 Tweet from emojicaselaw setting out R v Brown using emoji

Figure 12 Tweet from emojicaselaw setting out Donaghue v Stevenson using emoji

In R v Brown, a group of sadomasochists were arrested and tried for occasioning actual bodily harm (R v Brown [1994] 1 AC 212). Their sadomasochistic practices were interpreted as an offence because their activities involved creating wounds with tools. Meanwhile, in Donoghue v Stevenson, 'Patsy' purchased a bottle of ginger beer for 'Delia' in a cafe. 'Delia' drank half the ginger beer and found there was a snail at the bottom. This made her ill and, after some legal wrangling in contract law, she was compensated (Donoghue v Stevenson [1932] UKHL 100). As can be seen from the examples above, many emoji iconically communicate concepts in representations that are part of narrative structure, and this can be accounted for with the relevance-theoretic notion of interpretive resemblance, which we encountered in chapter one, §1.2. The tool icons seen in figure 11 stand in for the concepts of tools in the narrative which explains what sexual practices the men in R v Brown engaged in. The hat/streamers icon represents the parties they threw. In Donoghue v Stevenson (see figure 12), the snail icon represents the concept of the snail that was found in the bottle, while the icon of the two women represents the two friends that were involved in the case. These clearly conceptual non-face cases of emoji are not the focus of this section. This section, instead, focuses on uses of very common face emoji, and examines how they can be interpreted.

It has been thought that facial emoji serve to replace paralinguistic features of communication that are absent in electronically mediated contexts such as email or

Whatsapp. Ip (2014) concluded that facial emoji are important in conveying emotion online, while Read (2005, cited in Ptasynski *et al*. 2011, p 1161) showed that facial emoji can be useful in sentiment analysis to work out whether reviewers are happy with a product, or whether response to, say, a government policy is positive or not. Tossell *et al*. (2011, p 659) consider that facial emoji provide socio-emotional context online, and are the primary means of replacing nonverbal cues online. They (*ibid*.) suggest that these media can be disruptive for communicating content and intent behind messages, and facial emoji can help to clarify what is meant. Derks *et al*. (2008) also said that facial emoji enhance written communication in the same way visual cues or body language support verbal communication. Rezabek and Cochenour (1998, p 201) consider facial emoji to be visual cues that represent feelings and emotions. Danet *et al*. (1997) note that facial emoji allow for the expression of emotions. Walther and D'Addario (2001) consider that facial emoji can be emblematic of actual facial expressions, and, indeed, there is some research which suggests that certain facial emoji activate parts of the brain that are involved in processing facial expressions (Yuasa *et al*., 2011; Churches *et al*., 2014). In the literature, it seems the consensus is that facial emoji exist to (somehow) communicate particular emotions, and that they stand in for, and/or function in a similar way to paralinguistic features such as facial expressions.

Let's leave aside for a moment the question of how the facial emoji involved in the communication of emotions achieve relevance. I would like to pick up on an important observation made by Dresner and Herring (2010): it is not actually clear that all facial emoji contribute to an emotional interpretation. The reason for this is that it is easy to show that their use sometimes results in an interpretation that cannot be construed as emotional. The case that they use to support this argument is that of the 'winky', commonly rendered ;), or  in today's pictorial emoji.

Dresner and Herring (2010, p 252) note that this facial emoji is associated with joking or sarcasm, and these are attitudes, not emotions. They (*ibid*.) say that it is possible to joke or be sarcastic when you are happy and when you are sad, and that this facial emoji can be used to communicate sarcasm when one of a range of emotions is experienced by the speaker. In essence, it would seem that there are facial emoji which do not communicate an emotion,

but communicate an attitude. The below examples show this proposed disconnect between emotions and attitudes:

(32a) [A Whatsapp message between two people in a relationship.]

I've missed you so much. I can't wait for you to get home. 

(32b) [A Whatsapp message between two people in a relationship. They have been discussing the pros and cons of the sender giving up work to focus on their career.]

This job's crap  but you're going to look after me full time now, right? 

In (32a), we can say that the sender is happy that the receiver will be coming home soon, and, perhaps the repeated winking emoticons are involved in the recovery of a flirtatious attitude. In (32b), it is safe to say the sender is sad, but is making a joke about her partner supporting her while she gives up work. The data presented in this section so far should already cast doubt on the claim that facial emoji always communicate emotions. Within Relevance Theory, Yus (2014) has noted that facial emoji are associated with a range of functions, and not all of these are related to emotions or affect. For example, he (2014) associates some facial emoji with sarcasm, joking, and the intensity of propositional attitudes. Yus, however, also (2014) notes that certain facial emoji *are* involved in the communication of emotions, and that they can contribute to the computation of the intensity of emotions. This insight is essentially correct. However, as Yus' paper serves to produce a taxonomy of varying uses/interpretations of facial emoji, it is outside the scope of his paper to explain exactly how facial emoji contribute to any 'intensification' of emotional states. Whilst it is not a criticism, his examples do not feature any repetition of entire facial emoji, and so an account of how these might contribute to the interpretation of emotional states is needed. Given how some other non-linguistic or borderline linguistic devices are analysed to lead to the recovery of emotional states in Relevance Theory, it may be that some facial emoji trigger procedures in the emotion-reading domain(s).

It seems, then, that there are cases of individual facial emoji that are interpreted conceptually, and there are cases that are interpreted procedurally. Since the same facial

emoji can be involved in all cases, it is clear that a conceptual or procedural interpretation is imposed in context, following the relevance-theoretic comprehension procedure. Consider the following examples taken from text messages.

(33a) I am 

(33b) Don't be 

(33c) No, I don't want to 

(33d) I'll miss my grandad so much

In (33a), it seems the facial emoji is standing in for the word *sad*, and communicates a concept SAD*, just as the word would. Of course, there are extra weak effects associated with the rebus-like nature of the utterance, but a concept is still communicated, nevertheless. In (33b), the same applies. Thus there are cases of facial emoji that are interpreted conceptually. In (33c), the sad facial emoji could be interpreted as communicating a higher-level explicature, an attitude towards whatever it is that the sender does not want to do. This is part of explicit content. Just as we saw with *yes* and *no*, the facial emoji in this case might activate a procedure for the recovery of that higher-level explicature. In (33d), however, it is clear that the facial emoji communicates an emotional state, and I think that it activates a procedure for the recovery and calibration of an emotional state. To understand how and why this works, especially in repetition cases, I would like to think about how expletives and interjections are interpreted. In any case, we can see that there are conceptual and procedural instances of (the same) facial emoji, and it would appear to be expectations of relevance in context which help us to judge what sort of interpretation to impose on any given occasion.

As noted by Wharton (2003, 2009), expletives have been categorised as interjections by some researchers. Interjections are expressions that are considered to sit on the border of linguistics, both literally and figuratively. Interjections can be analysed as borderline linguistic expressions in that they can be stylised or regular in form, but do not communicate concepts or content, and may have something in common with response cries, e.g., *ouch* or *wow*

(Wharton, 2003, p 42). Interjections are also at the border of language in that they tend to occur on the peripheries. They are often syntactically independent of surrounding structure. Wharton (*ibid*.) notes that researchers associate many interjections with the expression of emotions. Wharton (2003, 2009) sees interjections as communicating, among other things, something like feelings, sensations or internal states rather than just something propositional, or a particular type of higher-level explicature. Wharton (2003, p 57) says that there is something non-representational, and expressive about certain interjections, and that how this communication takes places is not via conceptual representations.

Wharton (2009, p 109) treats interjections as expressions that can communicate attitudinal information, but can also communicate emotions that are not directed towards something, or at least not directed towards anything propositional - something more like feelings or sensations. This is, for him (2009, pp. 102-103), something very vague, and may involve a tiny increase in a very wide range of assumptions. Interjections are partly-natural and partly coded, meaning they can *say* and *show* when analysed individually. It is unsurprising that individual interjections can be seen at least partially as cases of *showing* because the interpretations they communicate are too vague and idiosyncratic to be something that could ever stabilise as a word meaning, or be paraphrased conceptually (Wharton, 2009, p 103). It is suggested that interjections may activate emotion-reading procedures during their interpretation (Wharton, 2009, 2015). What Wharton (2009, p 133) says is that an interjection can activate particular internal states which constrain the search for relevance. Certain coded facial expressions can do the same. Wharton (2003) did not wish to treat expletives as interjections. These are clearly much more linguistically integrated and productive than most interjections, and perhaps have a stronger element of linguistic encoding or *saying* than interjections proper, so this is reasonable. However, I do not think Wharton would object to me saying that expletives can also activate procedures for the recovery of particular feelings or emotions. When produced in isolation, as with some individual interjections, expletives would provide both direct and indirect evidence for the layer of information that the speaker wishes to 'point out':

(34a) Julie: Someone's keyed your car.

John: Shit!

(34b) [Kelly wakes up to find that her pet chicken has escaped for the third time and has gone missing.]
Neighbour: I think the chicken's gone walkies again.
Kelly: Fuck!

Let's assume that the cases reproduced above are communicative. I think it is reasonable to say that both Julie and the neighbour would be encouraged by the expressive expletives to recover something about John and Kelly's emotions, or a sub-component of an emotional state, e.g., a feeling or sensation. After all, both incidents are serious enough that John and Kelly might be worried and experiencing a sick kind of sensation in their stomachs - similar feelings might even then be induced in Julie and the neighbour via mirror neurons and mechanisms for empathy. I think it is reasonable to say that these individual expletives *show* these sensations and feelings as well as perhaps *say*, or encode, a means of recovering them. Individually, interjections and expletives involve at least some *showing*, especially where expressive prosody may also co-occur, as we would expect in the cases immediately above. Individual human facial expressions are natural signals that can also *show* someone's emotional state (Wharton, 2009), and can be analysed as activating emotion-reading procedures. If, as mentioned above, it is right that facial emoji can activate the same types of processes or procedures that are involved in the interpretation of actual facial expressions, then it seems reasonable to conclude that there are likely some similarities between what and how some interjections, some expletives, some facial expressions, and some facial emoji communicate. Facial emoji that are clearly not conceptual activate a procedure for the recovery of emotional states or particular subcomponents thereof, and these are not comments on or attitudes towards particular propositions. I believe that such facial emoji are cases of deliberate and communicative *showing*, and can be 'natural' just as genuine smiles or frowns are.

It should be pointed out that the analysis supplied here only applies to these items when produced in isolation, and not when they are repeated in immediate succession in some kind of group or unit. With repeated *yes* or *no* particles, the first particle served to *say*, that is, yield a linguistically encoded procedure for the recovery of a higher-level explicature. However,

subsequent particle tokens served to provide input to processes for computing the degree to which a higher-level explicature was committed to, for example. These were analysed as cases of *showing* in that the speaker displayed the repetition, and, the hearer being bound by the prosodic grouping, encouraged the audience to use the number of tokens for evidence about computing the strength of some aspect of the higher-level explicature in question. The same sort of analysis, with a few changes, is going to apply to repeated facial emoji which are clearly intended to communicate a particular emotional state.

In Wharton's (2009) discussion of interjections, *showing* and *saying*, and natural and non-natural codes, he explains that there are different types of coding that have been involved in human *linguistic* communication, and a particular distinction that has been made is the analogue-digital distinction (Sebeok, 1972, cited in Wharton, 2009, p 117). Wharton (2009, p 117) explains:

> 'The analogue-digital distinction exists in a variety of guises. Essentially, it is a distinction between codes or systems in which the repertoire of signals is either - in the case of analogue codes - graded, blended, or continuous, or - in the case of digital codes - discrete or discontinuous. In a graded system the boundaries between the signals cannot be demarcated, whereas within a digital one they can.'

Wharton (2009) helps us visualise analogue coding with the example of measuring gas pressure in a given system. He (2009, p 117) says:

> 'There is a variable of some physical quantity...This variable is related to another variable, say the needle in a pressure gauge, in such a way that the variations in the former are in a proportional relationship to the variations in the latter. As the pressure in the system rises, so the needle on the pressure gauge rises; as the pressure falls, the needle falls. The movement of the needle is analogous to the rising and falling of the pressure, and the continuous fluctuation of pressure is reflected in the continuous movement of the needle.'

Wharton (2009, pp. 118-119) explains that language is typically viewed as a digital system. It is combinatorial, and is made up of lots of discrete units in particular configurations. This is what gives language its power to communicate so many different propositions. However, while we interpret the words digitally, as it were, some of its packaging is more analogue in its nature. Wharton (2009, p 119-120) speaks of people 'reading off' emotional states from speakers' tone of voice or facial expressions. He suggests that smiles and tones of voice can have analogous relationships to the amount of affect, for example, that someone wishes to convey. We are able, as Wharton (*ibid*.) suggests, to pick up small variations in someone's smile, or prosodic behaviours and use them to *calibrate* the degree of the aspect of the interpretation they seem to have a proportional relationship to. The extent of someone's smile, and how ostensive it is can lead us to *calibrate* a very happy emotional state. If someone speaks rapidly and in the higher portion of their range, this ostensive behaviour can encourage us to *calibrate* an extremely excited emotional state for a speaker. Prosody and openly displayed natural facial expressions (and presumably other nonverbal natural elements) would seem to then activate some kind of procedure which requires a hearer to calibrate emotional states or sub-components thereof (which the hearer themselves may come to experience through something like a mirror neuron system) via this sort of analogue means. If facial emoji are interpreted similarly to actual facial expressions, then there is every reason to assume that they also lead an audience to calibrate these types of representations.

If we think back to what we learned about repeated adjectives, repeated intensifiers, and the repeated polarity particles *yes* and *no*, we saw that the first instance of such expressions was responsible for either yielding a concept or a procedure via linguistic encoding, and the subsequent tokens provided evidence (via *showing*) for the processing path that should be adopted and, in many cases, how much effort should be invested in developing a particular representation in a particular direction. With analogue systems, when more of the first variable enters a system, we see a corresponding change in the output display, the second variable. The louder the volume, the more the red LED monitor 'goes up'. The more steam, the more the needle gauge swings to the right. The more input, and the more ostensive the input, to an emotion calibration procedure, the more the audience must calibrate the the representation, or even, perhaps, the degree of activation in a particular network of neurons.

I suggest that some repeated interjections, expletives, and repeated facial emoji interpreted in an emotional way *within the same unit* are inputs that provide evidence for the degree of calibration in the emotion-reading system. The first facial emoji activates the procedure and requires the audience to calibrate a representation or sensation or emotional state of some kind - the subsequent facial emoji do not re-activate the procedure but *show* how the output of the procedure is to be calibrated. The more tokens, the higher the read out, metaphorically speaking. Greater token numbers represent more highly ostensive cases of *showing*. If we compare the cases below, it seems clear that the more repetitions, the greater the disgust that is felt:

(35a) Urgh not peas 

(35b) Urgh not peas 

(35c) Urgh not peas 

The same would seem to work for repeated expletives and interjections. The more token numbers we have, the more the audience is encouraged to *expend effort on further calibrating* an emotional representation, using the repeated tokens as evidence for how calibration should proceed. (The intonation boundaries in the examples below can also be exploited to further show the repetition, as I have posited for other cases.)

(36a) |FUCKfuckfuckfuck! |That is bloody PAINful! |

(36b) | We're going to be LATE.|SHITshitshitshitshit.|

(36c) | OWowowowow. | What did you POKE me for? |

(36d) | YUKyukyukyuk. |That's disGUSting.|

We saw in chapter four that intensification can occur at the implicit level of content in that a hearer can be encouraged to expend effort on extending or (re)activating the range of assumptions she entertains in an interpretation. Earlier in this chapter, I showed that intensification can occur at the level of explicit content if a hearer is encouraged to expend effort additionally adjusting a particular concept or fine tuning a higher-level explicature further. It seems that facial emoji, and likely certain interjections and expletives, when repeated, encourage an audience to expend effort on further calibrating the output or result

of a procedure. This could be via extending the range and degree of assumptions that are activated, as Wharton (2009) suggested for interjections, but it could also be that an audience is, in the right context, encouraged *to further calibrate* a representation of an emotional state, or *expand the reach and/or activation levels* of a particular network of neurons, perhaps mirror neurons, which might lead an audience to feel something like a sensation. Recall John and Julie arguing about the leak in chapter three, where I also mentioned feelings, and mirror neurons briefly. Some of Julie's remarks might have had the more or less intended effect of making John feel nervous or sick or worried - the type of thing that makes you feel a knot in your stomach. In this way, we can say that certain non-propositional emotional interpretations can be intensified, and this through the repetition of procedurally interpreted facial emoji, among other means. Again, intensification, whatever it is, appears to be a processing phenomenon.

Certainly, repeated facial emoji such as those in (35b) and (35c) can be said to *show* in that the only optimally relevant way to take them is as an input for calibrating an analogue representation of an emotional state. Is there any other sense in which they *show*? Do they have anything in common with other repetitions addressed in this thesis as far as *showing* is concerned? The answer is yes. The writer who employs repeated facial emoji in the same unit also draws attention to what is being repeated. The token number in (35c) makes the repetition itself particularly salient. We can also argue that the repeated tokens occur within the same 'unit'. Nothing intervenes between tokens. The audience can infer that the repetitions appear to occur within a unit. The perception of a unit can be interpreted as bounding the repeated tokens, and bounding, as we have seen throughout this chapter, can be used to further increase the ostensiveness of *showing*. I therefore also analyse repeated facial emoji such as those in (35b) and (35b) as cases of *indeterminate showing* whose effects cannot be paraphrased propositionally. Once attention is drawn to repeated facial emoji, there is an expectation that there are extra or different effects on offer for the demand made in this, and, in line with our expectations of relevance in context, the optimally relevant way to treat them is as providing relatively direct evidence for how calibration is to proceed. Is there any emphasis? I would suggest that in cases where a speaker produces a markedly high number of repeated emoji, this could be interpreted as very ostensive *showing* behaviour, and, thus, emphasis.

In this section, we have seen that facial emoji can communicate emotions but not always, and that, when they do, they can behave in similar ways to certain interjections and swears. They achieve relevance by activating procedures that are similar to those involved in processing genuine facial expressions that lead to the calibration of particular emotional states or sensations. When these facial emoji are repeated, the result is additional calibration. An increased number of tokens *shows* how the output of the procedure is to be calibrated. However, a writer using repeated emoji also *shows* in that he draws attention to his repetition to suggest that attention should be paid to it in pursuit of further effects. This fills a gap in the relevance-theoretic literature on how facial emoji can communicate emotions or components thereof. Certain interjections and expletives may activate the same type of procedure - one that requires a calibration rather than a different kind of output, and it is this insight which now allows us to explain why certain items interpreted procedurally can be repeated within an intonation group, while others cannot. Let's return to examples (30a-30d).

5.6.3 Why not everything linguistic and procedural can be repeated within an intonation group

(30a) |*She she she* is my friend.|

(30b) |Look,| it's *my my* book. |

(30c) |I spent all my wages. |However however, | I managed to go out.|

(30d) |*But but but* I didn't spend all my wages!|

The italicised expressions here can be analysed as encoding procedures. In the case of *she*, the procedure is for identifying a referent. I treat *my* as encoding a procedure for the computation of a relevant relation between two entities - a relation that can be possession, but need not be. *However* and *but* are treated within Relevance Theory as encoding procedures for the elimination or deletion of particular assumptions (Blakemore, 2002).

The reason why these expressions cannot be repeated within intonation groups is to do with what type of contextual inputs the linguistically encoded procedures 'look for'. *She* looks to

the context for an optimally relevant, single entity that can be referred to as female (regardless of actual gender). *My* looks to the context to establish a single relationship between an individual or object and the individuals or objects it stands in relationship to. *However* and *but* look to the context to find a single, highly salient assumption that can be eliminated. Loosely put, these are 'one time' procedures that seek one input from the context that can only be used once per intonation group, or, unit of relevance. Once that input is taken and computed over, an optimally relevant output results (e.g., the concept of a referent), and, as a hearer's expectation of relevance is then satisfied, it would put the hearer to undue effort to have her activate and undertake that procedure again in the same context of the same intonation group. The information from the context will have also been 'used up', and would not be as relevant a second or third time around. Considered from another angle, though these expressions are procedural, they compute over the propositional and the conceptual. In other words, the inputs and outputs they are associated with are best thought of in digital, discrete terms - these 'unrepeatable' procedures are not associated with calibrations as with the repetitions of interjections, expletives, and facial emoji. They are one-offs, and this is why they cannot be re-used within (intonation) groups. The claims I have made in this final section can only be taken as holding for the linguistic expressions I explicitly addressed. However, it is my expectation that the findings will hold for other linguistic procedural expressions within English and in other languages. The type of input a linguistically encoded procedure computes over should constrain whether and how it can be repeated.

Chapter Six: Repetition, Emphasis and Intensification: Findings and Future Directions

6.1 Aims, 'hunches' and hypotheses

This thesis began with a picture of repetition that was somewhat messy to say the least. As I showed in chapter one, no one really agreed what repetition was. It had many names and many faces. Stylistic repetition was often conflated with accidental and incidental repetitions of form, and with grammatical reduplication. I also demonstrated that this conflation worked the other way - a form of stylistic repetition was misdiagnosed as reduplication proper: the cases of X-x addressed in chapter three. This particular conflation appeared to arise, mainly, from failing to afford pragmatic inference sufficient a role in the interpretation of stylistic repetition. Moreover, there was conflation concerning the nature of the effects of repetition. The effects of repetition were called emphatic, intensifying, or some mixture of the two. It was not clear to me that emphasis and intensification were effects, and it seemed obvious that, whatever they were, they could not be the same phenomenon, and that unitary and cognitively-driven accounts of both emphasis and intensification were lacking. Finally, existing treatments of stylistic repetition, even within Relevance Theory (Sperber & Wilson, 1995), chiefly dealt with only a small subset of verbal cases of repeated material. A key aim of this thesis was to extend our understanding of stylistic repetition to a larger set of data than has been previously treated. This has been achieved, as this thesis analysed adjacent, fairly adjacent, and nonadjacent repetitions, and repetitions within an intonation boundary (and, so, repetition at below the level of the clause). I also analysed stylistic repetitions of borderline and non-linguistic communicative phenomena, e.g., some interjections, expletives, and facial emoji. In this way, I was able to provide a deflationary account of repetition, showing that it occurs in many forms of communicative behaviour, and is, in some sense, not that 'special'.

At the start of this thesis, my aim was simply to clean up some of the aforementioned mess. In addition to extending the number of cases of repetition that we could explain, I wanted to develop unitary, working accounts of emphasis and intensification, establish the relationship between repetition, emphasis, and intensification, and explore the interaction that I was beginning to notice between repetition and intonational grouping. Taking inspiration from work on expressives (Blakemore, 2011; Wharton, 2015), and, admittedly, developing a hunch, I surmised that emphasis was something to do with the ostensiveness of particular speaker behaviours in communication. Intensification, I thought, concerned what speakers undertook during the processing of communicative stimuli, and was a processing phenomenon. I

therefore considered that neither emphasis nor intensification could be treated as types of effects that were the output of utterance interpretation. Gathering my data, I wondered if the placing of intonation boundaries with respect to certain repetitions (e.g., epizeuxis) would guarantee effects at a particular level of content, and I set about investigating this in chapters four and five. I also had a feeling that I needed to account for how repetitions are recognised at 'distance' and wanted to think about what 'distance' might be from a relevance-theoretic perspective. What I never explicitly set out to do was develop an account of repetition that appeared to suggest, in one very particular respect, that stylistic repetition was actually a unitary phenomenon after all. The picture of repetition painted at the start of the thesis was one that suggested that it could never be treated in a unitary fashion. However, if my findings in this thesis are along the right lines, it appears that this is exactly what has emerged. I now summarise my main findings concerning these initial hunches and hypotheses, and make suggestions for future research. There were a great many phenomena that, for practical reasons, simply could not be included in this study.

6.2 Repetition and *showing*

6.2.1 A unitary account of stylistic repetition as cases of *indeterminate showing*

Given the range of phenomena I have addressed, repetitions that ostensively communicate stylistic effects cannot be treated as a unitary phenomenon from the point of view of their form. I addressed repetitions that were both linguistic, and borderline or non-linguistic, which were of different lengths, and where repeated tokens were of varying numbers and placed at different distances from their original. In terms of non-propositional effects, the effects communicated were either poetic (in the sense of Sperber & Wilson's (1995) and Pilkington's (2000) poetic effects) or emotional in nature, or some combination of the two. Strictly viewed, it might be tempting to say that these repetitions are not unitary from the point of view of their interpretations either. However, what became clear in each of my analyses was that they do have something in common from the point of view of their interpretations: the import of each repetition I treated is indeterminate. It is not possible to propositionally paraphrase what the communicator deploying such a repetition communicates; it is vague, nebulous and, to borrow a much loathed metaphor, hard to 'put into words'. This insight provides some

support for the major conclusion I draw regarding the role of stylistic repetitions in communication. As I have suggested several times in this thesis, the type of things that stylistic repetitions communicate are not the type of things that we could acquire and store and attach to an expression or chunk in our lexicon. The interpretations are just too nebulous, context-specific, and, often, idiosyncratic. Since we cannot encode these aspects of meaning linguistically, we have to find another means of communicating them. One of the ways that we can communicate vague, non-propositional effects is to *show* in order to suggest the path to relevance to our audiences. The suggestion that I have developed in this thesis is that the cases of stylistic repetition which I have analysed can all be considered cases of Sperber and Wilson's (2015) *indeterminate showing*, and, if this is correct, they can be treated as a unitary phenomenon from the point of view of the way that they achieve relevance.

Taking inspiration from my analysis of onomatopoeia, I suggested that a communicator reproducing a form that he has already ostensively produced in communication exploits the fact that, generally, our communicative behaviors are different from one another, and different in terms of internal composition. I suggested that communicators can 'play on' this and, sometimes, with the support of other behaviours such as prosodic highlighting and careful intonation grouping, can *show* through their repetition. What they effectively *show* is the form of what has been produced before (or, in the cases addressed in chapter five, what is currently being *shown*). The audience then has to reason from what is *shown* to what the speaker or writer wanted to communicate (an indeterminate interpretation). There is certainly a sense that a speaker who deliberately displays his repetition and draws attention to it seeks to make manifest that he deliberately did not produce a different communicative behaviour in a context where he might reasonably have been expected to do so. Where repetitions occur across different intonation groups, it seems that the result of the *showing* is to suggest to the audience that they should 'go back' and revisit and rework their original interpretation (as in the epizeuxis and 'long distance' repetitions of chapter four). Where the repetitions occur inside an intonation group or some other 'unit', the repeated material seems to bear some sort of iconic relationship either to the amount or type of adjustment required in pursuit of ad hoc concept construction or the construction of a higher-level explicature, or provide some kind of analogue input to emotion-reading procedures that require the calibration of an emotional state.

6.2.2 A natural account of stylistic repetition

While we can repeat both verbal and nonverbal communicative behaviours, repetition as *showing* is, itself, actually a nonverbal behaviour. And a very natural one at that - perhaps, in some sense, as natural as pointing, which is surely the human prima facie case of *showing*. I have said that I consider the term *natural* to mean 'exploiting what the mechanisms and principles underpinning the interpretation process are set up to do'. We are cognitively driven towards maximising relevance, generally, and optimising it when it comes to communication. We are also 'set up' to seek to recover interpretations that more or less resemble the ones our communicator has in mind, rather than duplicate thoughts (see chapter two, §2.7). This links with the fact we are primed to seek relevance in stimuli on our own responsibility, at some great effort at times, if we are presented with very ostensive communicative stimuli that suggest that a useful improvement to our cognitive environment is on offer. All other things being equal, if I point at something, even something you can't see well, you'll (try to) look at it because I point, and you'll keep looking at it and thinking about it until you offset that effort with sufficient effects. This may be mostly on your own responsibility. Communication would be very weak. Stylistic repetition is 'natural' in this way. It exploits how we communicate and how we seek positive changes to our cognitive environment. It also, to an extent, exploits what we cannot do with language, i.e., linguistically encode indeterminate interpretations. Essentially, if I want you to pay further attention to a stimulus in order to 'wring out' (more of) an indeterminate import, I'll *show* it to you, and I'll *show* it again and again. I'll repeat it. By virtue of my ostensive *showing*, you will pay attention to what I repeat, recovering a vague and nebulous interpretation, and/or, an emotional one, mostly on your own responsibility and creating quite significant cognitive mutuality between communicators, which in turn, can yield information about the assumed relationship between persons.

6.2.3 Updating relevance-theoretic accounts of showing

Although this thesis was intended to be an in-depth treatment of repetition, it turned out to also be, in some ways, a re-assessment of *showing*. Our current conception of *showing* seems correct in that it concerns presenting relatively direct evidence for what we want to

communicate (see chapter two, §2.8.1, and Wharton (2009)). However, the picture painted above means that we need to make an update to our explicit thinking on *showing*, as I already suggested in chapter two. Essentially, if I am correct, stylistic repetition is a nonverbal communicative behaviour which allows us to *show* not only visible, nonverbal behaviours, but *show* (aspects of) utterance form. Verbal utterances are invisible, transient sounds, and, so, are not part of the visual mode. Moreover, my analyses suggest that cases of *showing* are, actually, very often complex clusters of ostensive *showing* behaviours, e.g., open display, prosodic highlighting, and intonational grouping all at the same time - just think about Miranda's *plunge* case in chapter four. Multiple behaviours from different modalities can come together to really ***EMPHASISE*** what is being *shown*. This ought to also be reflected, if right, in future work on *showing* within Relevance Theory.

6.3 The role of intonation boundaries in the interpretation of stylistic repetitions

Let me now turn to further insights about intonation boundaries. Above, I suggested that, for cases of stylistic repetition, there does seem to be some sort of connection between whether material is repeated inside or outside of an intonation group, and how that repetition achieves relevance. I had wondered if intonation boundaries might somehow mandate whether or not effects were to be sought at the explicit or implicit level of content. This hypothesis is, perhaps unsurprisingly, not confirmed. The discussion of apposition in chapter four suggests that the point of many apposite expressions in separate adjacent boundaries is to scaffold particular narrowed concepts in pursuit of *explicit* content which then, in turn, can lead to particular implicit effects. I have shown that cases of epizeuxis, which concern repeated material in adjacent intonation groups, can achieve relevance at the level of explicit and/or implicit content, as Sperber and Wilson (1995) originally suggested (see chapter four, §4.2). While the repetition of repeated *yes*/*no* particles, gradable adjectives and the intensifiers *really* and *very* do interact with the recovery of explicit content when repeated within the same intonation group, there can come a point which cannot be specified in advance of context construction whereby they are more optimally relevantly interpreted as communicating effects on the implicit side of things. Instead, what seems to be happening is that repeated material, the placement of intonation boundaries, and the context interact to

suggest the path for relevance, i.e., the level of content at which effects should be sought. In line with the psycholinguistic literature reviewed in chapter four, §4.3.4, it seems that intonation boundaries serve a general bounding function and serve to delimit what I termed *units of relevance*. Where a stimulus such as a repetition occurs within a group, the optimally relevant way to treat that stimulus is to allow oneself to be restricted by the boundary to fixate on the repeated material, and keep collecting effects until the perception of a boundary, whereby the suggestion is to 'wrap up' the collection of effects either at the explicit or implicit level, depending on what is optimally relevant. Where repetitions occur in more or less adjacent but distinct intonation groups, the only optimally relevant way to treat them is as an encouragement to 'go back' and process that unit of relevance again, adding to and expanding the initial interpretation - whatever level(s) its effects were recovered at. My analyses also suggest that intonation boundaries can be exploited as part of a cluster of behaviours for further *showing* a repetition. Just as I can draw a circle around an amount on a bank statement to make it clearer to you what I want you to focus on, I can carefully place (and, heavily cue, as we saw in the discussion of *a fox, a fox*) my intonation boundaries to increase the salience of what might be a fairly subtle case of *showing* in the face of other competing and potentially very relevant stimuli, e.g., a ravenous fox eyeing up your beloved chickens in your garden.

6.4 A continuum from display to emphasis

What we see from the discussion so far in this chapter, and the thesis as a whole, is that ostensive speaker (or communicator) behaviours can cluster to further *show*, or increase the salience of what is *shown*, to ensure that more attention is attracted to it in order to suggest that there are extra or additional non-propositional effects on offer. The conclusion I am led to is one that was suggested by briefly considering the nature of expressives in chapter two, §2.8.4. Where the ostensive *showing* behaviour or cluster of ostensive *showing* behaviours of a speaker is very ostensive, I think we have emphasis. A speaker overlaying an expression with extremely high pitch height, loud volume and a marked pitch excursion, for example, could be said to be emphasising a word. A speaker who repeats using heavily cued boundary placement, prosodic signals, slow tempo and, say, like Miranda, ostensive glances to camera would be highly ostensive in their *showing*, and, as such, their ostensive communicative

behaviour would be emphatic. In terms of the repetition cases I addressed, I would suggest that there is a continuum of more subtle to completely unsubtle cases of *showing*. For example, spoken epizeuxis such as *a fox, a fox* exhibits emphatic speaker behaviour, while 'long distance' repetitions are potentially much more subtle, which explains why I said that, for the most part, the cases I addressed in chapter four did not feel emphatic in any sense. 'Distance' played a role of in the subtlety of some of these repetitions, where distance is understood in terms of processing effort and (re)activation. The findings in my thesis also suggest, then, why some researchers say repetition is emphatic, and why some don't - it depends what cases they are actually looking at. What all this suggests is a continuum between subtle cases of *showing*, or mere display, through less subtle cases, which I would say are akin to what has been called highlighting (Wilson & Wharton, 2005), right through to cases of emphasis - extremely ostensive *showing* behaviour. The way to move up or down the display-emphasis continuum depends on a judgment of ostensiveness against the background of a given context. Emphasis is not an inherent property of *showing* behaviour. It is emergent, just as relevance is. This picture of emphasis is a unitary one that meshes with the non-expert intuitions about what emphasis is, which I set out in chapter one, §1.6. Finally, Bolinger's (1972, p 17) comment concerning the 'noisiness' of some repetitions now makes sense in the context of this thesis. By 'noisiness', I think Bolinger meant something akin to 'very ostensive', or 'noticeable'.

6.5 Intensification as a processing phenomenon

What about intensification? I had an early intuition that intensification would prove to be a processing phenomenon. Although future work is required, this perspective seems to have been vindicated to an extent. X-x leads a hearer to expend more effort on further narrowing a concept (among other things). Adjacent and non-adjacent repetitions lead hearers to 'go back' and expend effort on extending the context for interpretation and reactivating and adding to weakly communicated implicatures. Repeated gradable adjectives and intensifiers can lead a hearer to expend more effort on further narrowing an ad hoc concept, while repeated *yes* and repeated *no* encourage the hearer to expend more effort on computing the degree of commitment to or strength of a higher-level explicature. Repeated expletives, interjections, and some facial emoji can lead an audience to expend more effort on calibrating

a representation of an emotional state of a communicator. Considered purely in this sense, intensification is not unitary - at least from the point of view of the level of content that its outputs emerge at. However, if we look at the analyses, what all these cases have in common is that the hearer expends more effort on carrying out more of a particular processing path or strategy during the interpretation process. This is what I think intensification is, and, so, in this sense, it can be given a unitary account.

6.6 The relationship between repetition, emphasis, and intensification

What is, then, the relationship between repetition, emphasis, and intensification? Put very simply, it can be summarised thus: the stylistic repetitions addressed in this thesis are cases of *indeterminate showing*. Some cases are very subtle, others, much more ostensive. Very ostensive cases of *showing* which involve very attention-attracting and flag-waving speaker behaviours can be said to exhibit emphasis. Emphasis, because it is so ostensive, suggests to the audience that there are suitable effects on offer, and that substantial effort should be invested in 'digging deep' to recover these effects. If a repetition exhibits a large amount of emphasis in terms of speaker behaviour, the hearer will be encouraged to undertake a correspondingly large amount of intensification during processing in pursuit of these effects. This is the connection between the three phenomena.

What we have seen in this thesis is that not all repetitions feature very emphatic speaker showing behaviour but still do lead to intensification - in particular, I think here of very 'long distance' repetitions. There is, then, a disconnect between emphasis and intensification. I believe emphasis always results in intensification of a kind, but that intensification could emerge without emphasis. The fact that not all repetitions feature emphatic communicative behaviour, and the fact that intensification might be able to obtain without emphasis surely explains, at least in part, why emphasis and intensification were conflated in the literature reviewed in chapter one.

6.7 Future directions

I divide this final section into phenomena that I was simply unable to integrate into the thesis, and research that I would like to undertake because puzzles remain and/or further research could be used to confirm and develop some of my findings.

6.7.1 The cutting room floor

To draw the conclusions I drew in this thesis, it was necessary to address exactly the repetitions which were included. A researcher can always include more phenomena. However, the more you address, the more you run out of space, and the less in-depth the accounts you do manage to integrate are. As such, there are a number of repetition phenomena that I could not investigate. While writing chapter four, I collected a wealth of data of repetitions from a host clause appearing in a parenthetical, and featuring prosodic highlighting, e.g., *I've got some lovely, and I MEAN LOVEly, sprigs of thyme*. The TV cooks Nigella Lawson and Jamie Oliver appeared to be quite the fans of such utterances, and it would be useful to investigate these with respect to the communication of stylistic effects in food programmes. Moreover, this work could have practical applications by explaining how experiential elements of cooking are 'shared' with viewers, and how cognitive and affective mutuality are increased between presenters and viewers. I also found examples of multiple repetitions connected by conjunction, and involving prosodic cues which suggested that processing should come to an end. But then it didn't. The utterances went on, and on, and on. And on. I was also unable to investigate the repetition of sound and rhyme, and consider the role of repetition in long-running jokes, and drawn out humorous skits found in cartoons such as *American Dad* and *Family Guy*.

There are also prosodic issues that I could not investigate in more depth in this thesis. I set out, naively, with no expectation of working with prosody, and, in the main, I only really looked at the placement of intonation boundaries and nuclei with respect to repetition. I am sure that particular intonation contours, as well as particular nucleic placements with respect to group-internal repetitions, would also affect how certain repetitions are interpreted. Moreover, I would like to conduct more thorough prosodic analysis of the repetitions I

addressed in chapter four. I would like to know just how much of a spoken 'form' can be repeated. Pitch height and pitch contour shape, clearly, can be repeated. I expect that I could show that intensity, loudness, and other prosodic features could all be repeated together to scaffold the perception of a spoken repetition.

6.7.2 Augmentatives and diminutives, intensification, and style

In terms of areas that I would like to investigate, I think that intensification needs much more study. What I would like to look at are cases of intensification which are considered morphological, and might be analysed to somehow linguistically encode intensification of a kind. In particular, I should like to look at the interaction between morphology, semantics and pragmatics concerning morphemes that have been called augmentative or diminutive, as I mentioned briefly in chapter five. If these can be said to encode any kind of intensification, and if intensification is a processing phenomenon, we might expect that their semantics is at least part-procedural and scaffolds the type and/or extent of adjustment required to develop explicit content. I would like to undertake a relevance-theoretic review of the semantics and pragmatics of augmentatives and diminutives. What I expect to find is that augmentation and diminuation, underneath the surface, turn out to be two sides of the same coin in terms of their contribution to conceptual adjustment in lexical pragmatics.

In terms of intensification and 'long distance' repetition, the data and my suggestions in chapter four potentially lend themselves nicely to empirical, psycholinguistic study. The prediction that follows from my account is that the audience of a 'long distance' repetition will be 'tripped up' by an ostensive repetition, and fixate on it while he or she 'goes back' to rework their initial interpretation. Psycholinguistic measures of surprise and fixation could be employed in an attempt to confirm this.

Finally, I end this thesis approximately where I started it - with style. Repetition is a stylistic phenomenon. It communicates stylistic effects, and it concerns a communicator packaging his communicative behaviours in a particular way so as to suggest the path to relevance in particular contexts. Repetition as *indeterminate showing* can communicate that a communicator is unable or unwilling to provide different evidence for the path to relevance

than he might reasonably have been expected to provide given the context, particularly in contexts where it may be manifest that what the speaker wants to communicate might be quite indeterminate. Think back to that rose garden in chapter two. Thus, it could be argued that there is a relationship between acts of *showing* and style - what a communicator makes manifest via *showing* about his interests and resources, and his audience's interests and resources, and, thus, about the relationship between communicators. I only really investigated one 'stylistic device' in detail in this thesis: repetition. I should like to further examine the relationship between *showing* and style in future work by examining other stylistic phenomena that may be treated as *showing*. Nevertheless, while the abovementioned issues are worthy of attention, I hope to have provided in this thesis a unitary account of stylistic repetition as *indeterminate showing*, cognitively-grounded and unitary accounts of emphasis and intensification, support for the case that intonation boundaries generally demarcate units of relevance, and a notion of 'distance' that may be useful in future pragmatic/stylistic analysis of other stylistic devices.

List of References

Aitchison, J. (1994). “Say it again, Sam”. In A. Fischer (Ed.), *Swiss Papers in English Language and Literature 7 Repetition* (pp. 15-34). Tuebingen: Gunter Narr Verlag.

Aitken, M. (2009). The English possessive marker in a framework of relevance. *Journal of Language and Communication Studies*, *43*, 119-135.

American Rhetoric. (2014). Rhetorical figures in sound: epizeuxis. Retrieved 21st April, 2014, from http://americanrhetoric.com/figures/epizeuxis.htm

Aristotle. (1963). *Poetics*. Boston, MA: Dutton.

Aronoff, M., & Fudemann, K. (2005). *What is morphology?* Oxford: Blackwell Publishing Ltd.

Attridge, D. (1994). The movement of meaning: phrasing and repetition in English poetry. In A. Fischer (Ed.), *Swiss Papers in English Language and Literature 7 Repetition* (pp. 61-83). Tuebingen: Gunter Narr Verlag.

Autism Speaks. (2015). What is autism? Retrieved 5th April, 2015, from https://www.autismspeaks.org/what-autism/video-glossary/glossary-terms

Bach, K. (1984). Conversational implicature. *Mind & Language*, *9*(2), 124-162.

Bakker, E. J. (1993). Discourse and performance: involvement, visualization, and “presence” in Homeric poetry. *Classical Antiquity*, *12*(1), 1-29.

Baldick, C. (2008). *The Oxford dictionary of literary terms*. Oxford: Oxford University Press.

Baldwin, S. R., & Coady, J. M. (1978). Psycholinguistic approaches to a theory of punctuation. *Journal of Reading Behaviour*, *10*(4), 363-375.

Bar-Hillel, J. (1971). Out of the pragmatics wastebasket. *Linguistic Inquiry*, *2*(3), 401-407.

Barker, C. (2002). The dynamics of vagueness. *Linguistics and Philosophy*, *25*(1), 1-36.

Barth-Weingarten, D. (2011). The fuzziness of intonation boundaries: some theoretical considerations and a practical solution. *InLiSt*, *51*, 1-22.

Bartsch, R., & Venneman, T. (1973). *Semantic structures: a study in the relation between syntax and semantics*. Frankfurt: Athaenum Verlag.

Bazzanella, C. (2011). Redundancy, repetition and intensity in discourse. *The Discourse of Redundancy, Special Issue of Language Sciences*, *33*(2), 243-254.

Beltrama, A., & Bochnak, M. R. (2015). Intensification without degrees cross-linguistically. *Natural Language & Linguistic Theory*, *3*, 843-879.

Blakemore, D. (1987). *Semantic constraints on relevance*. Oxford: Blackwell Publishers Ltd.

Blakemore, D. (2002). *Relevance and linguistic meaning: the semantics and pragmatics of discourse markers*. Cambridge: Cambridge University Press.

Blakemore, D. (2004). Discourse and relevance theory. In D. Schiffrin, D. Tannen & H. E. Hamilton (Eds.), *The handbook of discourse analysis* (pp. 100-118). Oxford: Blackwell Publishers Ltd.

Blakemore, D. (2006). Divisions of labour: the analysis of parentheticals. *Lingua*, *116*, 1670-1687.

Blakemore, D. (2008). Apposition and affective communication. *Language and Literature*, *17*(1), 37-57.

Blakemore, D. (2009). Parentheticals and point of view in free indirect style. *Language and Literature*, *18*(2), 129-157.

Blakemore, D. (2011). On the descriptive ineffability of expressive meaning. *Journal of Pragmatics*, *43*(14), 3537-3550.

Blass, R. (1990). *Relevance relations in discourse: a study with special reference to Sissala. Cambridge studies in linguistics 55*. Cambridge: Cambridge University Press.

Blodgett, A. R. (2004). *The interaction of prosodic phrasing, verb bias, and plausibility during sentence comprehension*. (Unpublished PhD thesis), Ohio State University, Ohio.

Bolinger, D. (1972). *Degree words*. The Hague: Mouton & Co.

Brody, J. (1986). Repetition as a rhetorical and conversational device in Tojolabal (Mayan). *International Journal of American Linguistics*, *52*(3), 255-274.

Brown, G., Currie, K. L., & Kenworthy, J. (1980). *Questions of intonation*. London: Croom Helm.

Burge, T. (1974). Demonstrative constructions, reference and truth. *Journal of Philosophy*, *71*(7), 205-223.

Burnett, H. (2012). The puzzle of absolute adjectives. *UCLA Working Papers in Linguistics*, *16*, 1-50.

Carston, R. (1998). Informativeness, relevance and scalar implicature. In R. Carston & S. Uchida (Eds.), *Relevance theory: applications and implications* (pp. 179-236). Amsterdam: John Benjamins.

Carston, R. (2002). *Thoughts and utterances: the pragmatics of explicit communication*. Oxford: Blackwell Publishers Ltd.

Carston, R. (2009). Relevance theory: contextualism or pragmaticism? *UCL Working Papers in Linguistics*, *21*, 17-24.

Carston, R. (2016). The heterogeneity of procedural meaning. *Lingua*, In press.

Carston, R., & Wearing, C. (2014). Metaphor, hyperbole and simile: a pragmatic approach. *Language and Cognition*, *3*(2), 283-312.

Chafe, W. (1976). Givenness, contrastiveness, definiteness, subjects, topics and point of view. In C. N. Li (Ed.), *Subject and topic* (pp. 27-55). New York: Academic Press.

Chafe, W. (1994). *Discourse, consciousness and time. The flow and displacement of conscious experience in speaking and writing*. Chicago: Chicago University Press.

Chomsky, N. (1965). *Aspects of the theory of syntax*. Cambridge, Massachusetts: The M.I.T. Press.

Chomsky, N. (1986). *Knowledge of language: its nature, origins and use*. Praeger: New York.

Clark, B. (1991). *Relevance theory and the semantics of non-declaratives*. (Unpublished PhD Thesis), University College London, London.

Clark, H., & Marshall, C. (1981). Definite reference and mutual knowledge. In A. Joshi, B. H. Weber & I. A. Sag (Eds.), *Elements of discourse understanding* (pp. 11-63). Cambridge: Cambridge University Press.

Cosmides, L., & Tooby, J. (2000). Evolutionary psychology and emotions. In M. Lewis & J. Haviland Jones (Eds.), *Handbook of emotions* (pp. 91-115). New York: Guilford.

Cruse, D. (1980). Antonyms and gradable complementaries. In D. Kastovsky (Ed.), *Perspektiven der lexicalischen Semantik: Beitraege zum Wuppertaler Semantikkolloquium Dezember 1977* (pp. 14-25). Bonn: Bouvier.

Crystal, D. (1975). *The English tone of voice. Essays in intonation, prosody and paralanguage*. London: Arnold.

Curl, L., Local, J., & Walker, G. (2003). Repetition and the prosody-pragmatics interface. *Journal of Pragmatics*, *38*(10), 1721-1751.

Danesi, M. (2012). *Popular culture: introductory perspectives.* Plymouth: Rowman & Littlefield Publishers, Inc.

Danet, B., Ruedenberg-Wright, L., & Rosenbaum-Tamari, Y. (1997). "Hmmm...where's that smoke coming from?" Writing, play, and performance on internet relay chat. *Journal of Computer-Mediated Communication*. *2*(4).

Davenport, M., & Hannahs, S. J. (2013). *Introducing phonetics and phonology*. London: Routledge.

Delin, J. (1987). Towards a model for generating cleft sentences. In J. Verschueren (Ed.), *Pragmatics at issue: selected papers of the international pragmatics conference, Antwerp, August 17-22, 1987, 1* (pp. 113-132). Amsterdam: John Benjamins Publishing Company.

Derks, D., Bos Arjan, E. R., & von Grumbkow, J. (2008). Emoticons in computer-mediated communication: social motives and social context. *Cyberpsychology and Behaviour*, *11*(1), 99-101.

Duffy, C. A. (2004). *New selected poems 1984-2004*. London: Picador.

Duffy, C. A. (2006). *The rapture*. London: Picador.

Duffy, C. A. (2012). *The bees*. London: Pan Macmillan.

Farnsworth, W. (2011). *Farnsworth's classical English rhetoric*. Boston: David R. Godine.

Fischer, A. (Ed.). (1994). *Swiss papers in English language and linguistics, repetition 7.* Tuebingen: Gunter Narr Verlag.

Gallai, F. (2013). *Understanding discourse markers in interpreter mediated police interviews*. (Unpublished Phd thesis), University of Salford, Salford.

Gargani, A. (2014). *Poetic comparisons. How similes are understood*. (Unpublished PhD thesis), University of Salford, Salford.

Garner, B. (2009). *Garner's modern American usage*. Oxford: Oxford University Press.

Gazdar, G. (1979). *Pragmatics: implicature, presupposition and logical form*. New York: Academic.

Gerleman, G. (1951). The Song of Deborah in the light of stylistics. *Vetus Testamentum*, *3*, 168-180.

Ghomeshi, J., Jackendoff, R., Rosen, N., & Russell, K. (2004). Contrastive focus reduplication in English (The salad-salad paper). *Natural Language & Linguistic Theory*, *22*, 307-357.

Gil, D. (2005). From repetition to reduplication in Riau Indonesian. In B. Hurch (Ed.), *Studies in reduplication* (pp. 31-64). Berlin: Mouton de Gruyter.

Giora, R. (1985). Towards a theory of coherence. *Poetics Today*, *6*, 699-716.

Giora, R. (1991). On the cognitive aspects of the joke. *Journal of Pragmatics*, *16*, 465-485.

Giora, R. (1997). Discourse coherence and theory of relevance: Stumbling blocks in search of a unified theory. *Journal of Pragmatics*, *27*, 17-34.

Giora, R. (1998). Discourse coherence is an independent notion: a reply to Deirdre Wilson. *Journal of Pragmatics*, *29*, 75-86.

Grice, H. P. (1989). *Studies in the way of words*. London: Harvard University Press.

Gussenhoven, C. (2002). *Intonation and interpretation: Phonetics and phonology*. Retrieved from https://perso.limsi.fr/mareuil/control/gussenhoven.pdf

Gussenhoven, C. (2015). *Foundations of intonational meaning: anatomical and physiological factors*. Retrieved from http://gep.ruhosting.nl/carlos/FoundationsOfIntonationalMeaning_2015.pdf

Halliday, M. A. K., & Hasan, R. (1976). *Cohesion in English*. London: Longman.

Hernandez, A.C. (2007). Gapping as a syntactic dependency. *Lingua*, *117*(12), 2106-2133.

Higashimori, I., & Wilson, D. (1996). Questions on relevance. *UCL Working Papers in Linguistics*, *8*, 111-124.

Higginbotham, J. (1985). On semantics. *Linguistic Inquiry, 16*, 547-93.

Higginbotham, J. (1988). Contexts, models and meanings: a note on the data of semantics. In R. Kempson (Ed.), *Mental representations: the interface between language and reality* (pp. 29-48). Cambridge: Cambridge University Press.

Hobbs, J. (1979). Coherence and coreference. *Cognitive Science*, *3*, 67-90.

Huddleston, R., & Pullum, G. (2010). *The Cambridge grammar of the English language*. Cambridge: Cambridge University Press.

Intro to RST. (n.d.). Retrieved 5th August, 2015, from http://www.sfu.ca/rst/01intro/intro.html

Ip, A. (2014). *The impact of emoticons on affect interpretation in instant messaging*. Retrieved from http://www.amysmile.com/doc/emoticon_paper.pdf

Jackendoff, R. (1997). *The architecture of the language faculty*. Cambridge, MA: MIT Press.

Jasinski, J. (2001). *Sourcebook on rhetoric*. London: Sage Publications Ltd.

Jucker, A. H. (1994). Irrelevant repetitions: a challenge to relevance theory. In A. Fischer (Ed.), *Swiss Papers in English Language and Literature Repetition 7* (pp. 47-60). Tuebingen: Gunter Narr Verlag.

Karttunen, L., & Peters, S. (1979). Conventional implicature. In C. Oh & D. Dinneen (Eds.), S*yntax and semantics, vol. 11: presupposition* (pp. 1-56). New York: Academic Press.

Keefe, R. (2006). Vagueness: philosophical aspects. In A. Barber & R. J. Stainton (Eds.), *Concise encyclopaedia of philosophy of language and linguistics* (pp. 807-810). London: Elsevier.

Kennedy, C. (1999). Gradable adjectives denote measure functions not partial functions. *Studies in the Linguistic Sciences*, *29*(1), 65-80.

Kennedy, C. (2007). Vagueness and grammar: the semantics of relative and absolute gradable adjectives. *Linguistics and Philosophy*, *30*(1), 1-45.

Kennedy, C., & McNally, L. (2005). *Scale structure, degree modification, and the semantics of gradable predicates*. Retrieved from https://www.upf.edu/pdi/louise-mcnally/_pdf/publications/kennedymcnally.pdf

Kjellmer, G. (2008). Self-repetition in spoken English discourse. In T. Nevalainen, I. Taavitsainen, P. Pahta & M. Korhonen (Eds.), *The dynamics of linguistic variation: corpus evidence on English past and present* (pp. 37-60). Amsterdam: John Benjamins Publishing Company.

Klein, E. (1991). Comparatives. In A. von Stechow & D. Wunderlich (Eds.), *Semantik: ein internationales Handbuch der zeitgenoessischen Forschung* (pp. 673-691). Berlin: Walter de Gruyter.

Koguchi, K. (2009). *Repetition in Dicken's A Tale of Two Cities: an exploration into his linguistic artistry*. Hiroshima: Keisuisha Co. Ltd.

Kortmann, B., & Loebner, S. (2002). *Understanding semantics*. London: Taylor & Francis Ltd.

Krifka, M. (2006). *Basic notions of information structure*. Retrieved from http://amor.cms.hu-berlin.de/~h2816i3x/Publications/Krifka_InformationStructure.pdf

Lamm, C., & Majdandzic, J. (2015). The role of shared neural activations, mirror neurons, and morality in empathy - a critical comment. *Neuroscience Research*, *90*, 15-24.

Leach, E. (1976). *Culture and communication: the logic by which symbols are connected. An introduction to the use of structuralist analysis in social anthropology.* Cambridge: Cambridge University Press.

Leech, G. N., & Short, M. H. (1981). *Style in fiction: A linguistic introduction to English fictional prose*. London: Longman.

Lenk, U. (1998). *Marking discourse coherence: functions of discourse markers in spoken English*. Tuebingen: Gunter Narr Verlag.

Levinson, S. (1983). *Pragmatics*. Cambridge: Cambridge University Press.

Liebert, W., Redeker, G., & Waugh, L. R. (1997). *Discourse and perspective in cognitive linguistics*. Amsterdam: John Benjamins B.V.

LiteraryDevices. (2013). Epizeuxis. Retrieved 5th November, 2014, from http://literarydevices.net/epizeuxis/

Luo, Y., Yan, M., & Zhou, X. (2013). Prosodic boundaries delay the processing of upcoming lexical information during silent sentence reading. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *39*(3), 915-930.

MacMahon, B. (2007). The effects of sound patterning in poetry, a cognitive pragmatic approach. *Journal of Literary Semantics*, *36*, 103-120.

Malmkjaer, K. (2003). *Linguistics encyclopaedia*. London: Routledge.

Mann, W., & Thompson, S. (1986). Relational propositions in discourse. *Discourse Processes*, *9*, 57-90.

McNally, L. (2011). *The relative role of property type and scale structure in explaining the behaviour of gradable adjectives*. Retrieved from https://www.upf.edu/pdi/louise-mcnally/_pdf/publications/mcnally-lncs10_v2.pdf

McNally, L. (2016). *Modification*. Retrieved from https://www.upf.edu/pdi/louise-mcnally/_pdf/publications/Cambridge_Handbook_article_Modification_v2.pdf

Miranda. (2014). 'Je regrette nothing' *Miranda*, Series 3, Episode 4, BBC1, January, 2014.

Monty Python. (1970). *Monty Python's Flying Circus*, Episode 25, BBC, December, 1970.

Nadarajan, S. (2006). A crosslinguistic study of reduplication. *Arizona Working Papers in SLAT*, *13*, 39-541.

Nan. (2014). *Catherine Tate's Nan*, Series 1, Episode 1, BBC1, January, 2014.

Norrick, N. (1987). Functions of repetition in conversation. *Text - Interdisciplinary Journal for the Study of Discourse*, *7*(3), 245-264.

Payne, J., & Huddleston, R. (2002). Nouns and noun phrases. In R. Huddleston & G. Pullum (Eds.), *The Cambridge grammar of the English language* (pp. 323-523). Cambridge: Cambridge University Press.

Pike, K. (1967). *Language in relation to a unified theory of the structure of human behavior.* The Hague: Mouton.

Pilkington, A. (2000). *Poetic effects: a relevance theory perspective*. Philadelphia: John Benjamins B.V.

Plath, S. (1965). *Ariel*. London: Faber and Faber.

Poe, E. A. (2002). *Complete tales and poems*. New Jersey: Castle Books.

Potts, C. (2007). The expressive dimension. *Theoretical Linguistics*, *33*(2), 165–197.

Preminger, A., Brogan, T. V. F., & Warnke, F. J. (1993). *The new Princeton encyclopaedia of poetry and poetics*. Princeton: Princeton University Press.

Ptasynski, M., Rzepka, R., & Araki, K., & Momouchi, Y. (2011). *Research on emoticons: review of the field and proposal of research framework*. Retrieved from http://arakilab.media.eng.hokudai.ac.jp/~ptaszynski/data/E5-6.pdf

Pullum, G. (2006, 28 Nov). Very very good pope. [Weblog]. Retrieved from http://itre.cis.upenn.edu/~myl/languagelog/archives/003830.html

Quirk, R., Greenbaum, S., Leech, G., & Svartvik, J. (1985). *A comprehensive grammar of the English language*. London: Longman.

Rand, A. (2010). *Anthem*. USA: Pacific Publishing Studio.

Rayo, A. (2011). *A plea for semantic localism*. Retrieved from http://web.mit.edu/arayo/www/plea.pdf

Recanati, F. (1989). The pragmatics of what is said. *Mind and Language*, *4*, 294-328.

Redman, S. (1997). *English vocabulary in use*. Cambridge: Cambridge University Press.

Rey, G. (1980). Functionalism and the emotions. In A. Rorty (Ed.), *Explaining emotions* (pp. 163-198). Los Angeles: University of California Press.

Rezabek, L. L., & Cochenour, J. J. (1998). Visual cues in computer-mediated communication: supplementing text with emoticons. *Journal of Visual Literacy*, *18*(2), 201-215.

Romero-Trillo, J. (2014). *Yearbook of corpus linguistics 2014: new empirical and theoretical paradigms*. London: Springer.

Rosch, E. (1973). As a category, it's natural! Natural categories. *Cognitive Psychology*, *4*, 328-350.

Rotstein, C, & Winter, Y. (2004). Total adjectives vs. partial adjectives: scale structure and higher-order modifiers. *Natural Language Semantics*, *12*, 259-288.

Rubino, C. (2005). Reduplication. In M. Haspelmath, M. S. Dryer, D. Gil & B. Comrie. (Eds.), *The world atlas of language structures* (pp. 114-117). Oxford: Oxford University Press.

Sanders, T., & Maat, H. P. (2006). Cohesion and coherence: linguistic approaches. *Reading*, *99*, 440-466.

Sasamoto, R., & Jackson, R. (2015). Onomatopoeia – showing-word or saying-word? Relevance theory, lexis, and the communication of impressions. *Lingua*, In press.

Sassoon, G. W., & Toledo, A. (2011). *Absolute and relative adjectives and their comparison classes*. Retrieved from https://staff.fnwi.uva.nl/u.endriss/teaching/lolaco/2011/papers/sassoon.pdf

Saussure, F. (1974). *Course in general linguistics*. London: Fontana/Collins.

Schafer, A. J. (1997). *Prosodic parsing: the role of prosody in sentence comprehension*. (Unpublished PhD thesis), University of Massachusetts, Massachusetts.

Scott, K. (2013). This and that: a procedural analysis. *Lingua*, *131*, 49-65.

Scott, K. (2015). Pronouns and procedures: reference and beyond. *Lingua*, In press.

Sedivy, J. C., Tannenhaus M. K., Chambers, C. G., & Carlson, G. N. (1999). Achieving incremental semantic interpretation through contextual representation. *Cognition*, *71*, 109-147.

Shakespeare, W. (2000). *Romeo and Juliet*. Hertfordshire: Wordsworth Editions Limited.

Shannon, C. E. (1948). A mathematical theory of communication. *The Bell System Technical Journal*, *27*, 379-423.

Siegel, M. E. A. (1979). Measure adjectives in Montague grammar. In S. Davis & M. Mithun (Eds.), *Linguistics, philosophy, and Montague grammar* (pp. 223-262). Austin: University of Texas Press.

Snowden, J. S., Neary, D., & Mann, D. M. A. (2002). Frontotemporal dementia. *The British Journal of Psychiatry*, *180*(2), 140-143.

Sperber, D., & Wilson, D. (1995). *Relevance: communication and cognition* (2nd ed.). Oxford: Blackwell Publishers Ltd.

Sperber, D., & Wilson, D. (1998). The mapping between the mental and the public lexicon. In P. Carruthers & J. Boucher (Eds.), *Thought and language* (pp. 184-200). Cambridge: Cambridge University Press.

Sperber, D., & Wilson. D. (2006). A deflationary account of metaphor. *UCL Working Papers in Linguistics*, *18*, 171-203.

Sperber, D., & Wilson, D. (2012). *Meaning and relevance*. Cambridge: Cambridge University Press.

Sperber. D., & Wilson, D. (2015). Beyond speaker's meaning. *Croatian Journal of Philosophy*, *15*(44), 117-149.

Speyer, A., & Fetzer, A. (2014). The coding of discourse relations in English and German argumentative discourse. In H. Gruber & G. Redeker (Eds.), *The pragmatics of discourse coherence: theories and applications* (pp. 87-120). Amsterdam: John Benjamins B.V.

Stede, M. (2012). *Discourse processing*. No location: Morgan & Claypool.

Styler, W. (2016). *Using Praat for linguistic research*. Retrieved from http://savethevowels.org/praat/

Syrett, K., Bradley, E., Kennedy, C. & Lidz, J. (2005). *Shifting standards: children's understanding of gradable adjectives*. Retrieved from http://home.uchicago.edu/~ck0/docs/sbkl-galana05.pdf

Syrett, K., Kennedy, C., & Lidz, J. (2009). Meaning and context in children's understanding of gradable adjectives. *Journal of Semantics*, *27*(1), 1-35.

Tannen, D. (1983). 'I take out the rock-dok!': how Greek women tell about being molested and create involvement. *Anthropological Linguistics*, *25*, 359-373.

Tannen, D. (1987). Repetition in conversation: towards a poetics of talk. *Language*, *63*(3), 574-605.

Tossell, C. C., Kortum, P., Shepard, C., Barg-Walkow, L. H., Rahmati, A., & Zhong, L. (2011). A longitudinal study of emoticon use in text messaging from smartphones. *Computers in Human Behaviours*, *28*, 659-663.

Ulatowska, H. K., Streit Olness, G., Hill, C. L., & Keebler, M. (2000). Repetition in narratives of African Americans: the effects of aphasia. *Discourse Processes*, *30*(3), 265-283.

Unger, P. (1975). *Ignorance*. Oxford: Clarendon Press.

Van der Gaag, C., Minderaa, R. B., & Keysers, C. (2007). Facial expressions: what the mirror neuron system can and cannot tell us. *Social Neuroscience*, *2*(3-4), 179-222.

Van Rooij, R. (2011). Vagueness and linguistics. In G. Ronzitti (Ed.), *Vagueness: a guide (logic, epistemology and the unity of science) series 19* (pp.123-170). London: Springer.

Verhaeghe, P. (2008). *On being normal and other disorders*. London: Karnac Books Ltd.

Vickers, B. (1994). Repetition and emphasis in rhetoric: theory and practice. In A. Fischer (Ed.), *Swiss Papers in English Language and Literature Repetition 7* (pp. 85-113). Tuebingen: Gunter Narr Verlag.

Vulliamy, E. (2015). *The Independent*. Emojis: as Japanese tire of emoticons, could their days be numbered elsewhere? Retrieved from http://www.independent.co.uk/life-style/gadgets-and-tech/emojis-as-japanese-tire-of-emoticons-could-their-days-be-numbered-elsewhere-a6771516.html

Walaszewska, E. (2013). Like in similes - a relevance-theoretic view. *Research in Language*, *11*(3), 323-334.

Wales, K. (2001). *A dictionary of stylistics*. Harlow: Pearson.

Walther, J. B., & D'Addario, K. P. (2001). The impacts of emoticons on message interpretation in computer-mediated communication. *Social Science Computer Review*, *19*(3), 324-347.

Waltraud, P. (1999). Verb gapping in Chinese: a case of verb raising. *Lingua*, *117*(3-4), 207-226.

Ward, G., & Birner, B. (2001). *Discourse and information structure*. Retrieved http://faculty.wcas.northwestern.edu/~ward/HDA01.pdf

Warner Bros Pictures. (2014). *Official trailer 1: 300 Rise of an Empire*. Viewed December, 2014 at https://www.youtube.com/watch?v=2zqy21Z29ps

Wharton, T. (2003). Interjections, language and the 'showing'/'saying' continuum. *Pragmatics and Cognition*, *11*(1), 39-91.

Wharton, T. (2009). *Pragmatics and non-verbal communication*. Cambridge: Cambridge University Press.

Wharton, T. (2015). *That bloody-so-and-so has retired: expressives revisited*. Retrieved from https://www.researchgate.net/publication/284015363_That_bloody_so-and-so_has_retired_Expressives_revisited

Wichmann, A. (2013). *Intonation in text and discourse: beginnings, middles, and ends*. Abingdon: Routledge.

Wilson, D. (1998). Discourse, coherence and relevance: a reply to Rachel Giora. *Journal of Pragmatics*, *29*, 57-74.

Wilson, D. (2004). Relevance and lexical pragmatics. *UCL Working Papers in Linguistics*, *16*, 343-360.

Wilson, D. (2011). The conceptual-procedural distinction: past, present and future. In V. Escandell-Vidal, M. Leonetti & A. Ahern (Eds.), *Procedural meaning: problems and perspectives* (pp. 3-31). Bingley: Emerald.

Wilson, D. (2014). Relevance theory. *UCL Working Papers in Linguistics*, *26*, 129-148.

Wilson, D., & Carston, R. (2007). A unitary approach to lexical pragmatics: relevance, inference and ad hoc concepts. In N. Burton-Roberts (Ed.), *Pragmatics* (pp. 230-259). London: Palgrave Macmillan.

Wilson, D., & Matsui, T. (1998). Recent approaches to bridging: truth, coherence and relevance. *UCL Working Papers in Linguistics*, *10*, 173-200.

Wilson, D., & Sperber, D. (1981). On Grice's theory of conversation. In P. Werth (Ed.), *Conversation and discourse* (pp. 155-178). London: Croom Helm.

Wilson, D., & Sperber, D. (1988). Mood and the analysis of non-declarative sentences. In J. Dancy, J. Moravcsik & C. Taylor (Eds.), *Human agency: language, duty and value* (pp. 77-101). Stanford: Stanford University Press.

Wilson, D., & Sperber, D. (2002). Relevance theory. *UCL Working Papers in Linguistics*, *14*, 249-287.

Wilson, D., & Wharton, T. (2005). Relevance and prosody. *UCL Working Papers in Linguistics*, *17*, 427-454.

Wit, E. C., & Gillette, M. (1999). *What is linguistic redundancy? Technical report*. Retrieved from http://www.math.rug.nl/~ernst/linguistics/redundancy3.pdf

Yus, F. (2014). Not all emoticons are created equal. *Linguagem em Discurso*, *14*(3), 511-529.

Yus, F. (2016). *Humour and relevance*. Amsterdam: John Benjamins B.V.

300: Rise of an Empire. (2014). Directed by Noam Murro [film]. USA: Warner Bros. Pictures.