

WELFARE — AT A SOCIAL DISTANCE

PROJECT REPORT

SEPTEMBER 2021

Solidarity in a crisis? Trends in attitudes to benefits during COVID-19

Robert de Vries, Ben Baumberg Geiger, Lisa Scullion, Kate Summers,
Daniel Edmiston, Jo Ingold, David Robertshaw and David Young

Welfare at a (Social) Distance is a major national research project investigating the benefits system during the COVID-19 pandemic, funded by the Economic and Social Research Council as part of UK Research and Innovation's rapid response to COVID-19. For more information about the project or to contact the team, please visit distantwelfare.co.uk

SUMMARY

There were good reasons to think COVID-19 would increase public support for welfare: it was a time of apparent increased solidarity in the face of a collective crisis; of clearly ‘deserving’ claimants; and of increases in direct experiences of the benefits system. And yet, the limited evidence collected so far suggests that attitudes have not changed.

In this report, we explain this puzzle in the UK, using two datasets which are uniquely suited to the challenge: (i) bimonthly data collected by YouGov from 2019-2021, which provides comparable, high-resolution information on attitude changes across the pandemic and (ii) a nationally representative survey we conducted as part of the *Welfare at a (Social) Distance* project in June 2021, which explored COVID-19-related attitudes in unique detail.

We found that COVID-19 prompted little change in public welfare attitudes. Attitudes did become slightly more generous during the first wave of the pandemic, only to rebound quickly in the Summer of 2020. The second COVID wave prompted another small increase in generosity. However, this appears unlikely to have endured.

Overall, comparing May 2021 with the pre-pandemic period, the public were less anti-welfare than before – but only slightly. The extent of the overall change from pre-pandemic to June 2021 differs between measures (from six percentage points for whether people out of work get too much support; to no change in whether benefit conditions are strict enough), and is stronger for Conservative voters than Labour voters. In the context of a considerable softening in attitudes 2013-19, however, all of the pandemic-associated changes are small.

There are two possible explanations for this muted change: firstly, that the pandemic has not meaningfully changed attitudes; and second, that attitudes have changed, but in ways which bracket COVID claimants away from pre-pandemic claimants, which we term ‘COVID exceptionalism’. We investigate which of these are true using three different approaches within In our novel survey of welfare attitudes conducted in June 2021, we found that:

- Half of Britons (50.6%) think that COVID-19 claimants are more likely than pre-pandemic claimants to be deserving (with only a small minority saying COVID-19 claimants are less likely to be deserving);
- People were more than twice as likely to say that pre-pandemic claimants (vs. COVID-19 claimants) were to blame for being unable to find a job, and more than three times as likely to say pre-pandemic claimants were to blame for losing their jobs.
- When asked about a hypothetical vignette claimant, people were nearly twice as likely to say that a pre-pandemic claimant (vs. a COVID-19 claimant) was at least a little to blame for being out of work.

In other words, we consistently found that the public believes COVID-19 claimants are considerably more deserving of benefits than pre-pandemic claimants. We did not show that pre-pandemic claimants were viewed harshly: pre-pandemic claimants were viewed quite sympathetically, and COVID claimants simply more so.

We also asked people to describe in free text the main differences between people claiming before vs. during COVID-19. We analysed these responses using Structural Topic Models. These show that COVID claimants were generally identified as the unfortunate victims of a situation beyond their control, while pre-COVID claimants were much more likely to be labelled by respondents as potential 'scroungers'. It is also worth noting, however, that many respondents focused on the broader impact of the pandemic at a more abstract level, or on the more human cost of the economic disruption on the living standards of both existing and new claimants alike.

These results pose a puzzle: why did a large influx of claimants seen as 'more deserving' fail to produce a correspondingly large increase in welfare generosity? The reason seems to be that general welfare attitudes are much more closely tied to perceptions of pre-pandemic claimants than to perceptions of COVID-19 claimants – that is, we find support for 'COVID exceptionalism'. COVID-19 claimants are considered to be different from 'normal' claimants, and welfare attitudes are to some extent insulated from sympathetic perceptions of this new type of claimant.

It would be easy to conclude that despite COVID-19, the public has little appetite for a more generous welfare system – but this would be wrong. Before the pandemic, attitudes had become more pro-welfare than the UK has seen in 20-30 years. Furthermore, if universal increases in welfare generosity are framed as specifically COVID-related (such as the £20 uplift in Universal Credit), support for increasing benefits in general is higher still. Rather than COVID-19 leading to an automatic transformation of attitudes, we can instead see it as offering discursive opportunities for a more generous benefit system, which political actors may choose to exploit or to counter, and whose resonance will depend on wider contextual factors. As an expert roundtable in June 2020 noted, "*There is nothing inevitable about how the welfare state will be reshaped post-crisis... Political control matters*" (Burchardt, 2020). The impact of COVID-19 on welfare attitudes and generosity is therefore all to play for.

CONTENTS

Summary

1. Introduction

- 6 1.1 Why would COVID-19 change attitudes?
- 7 1.2 Challenging expectations of increased generosity
- 9 1.3 The emerging evidence (and its limits)

2. Data sources

3. Results

- 12 3.1 Trends in attitudes
- 16 3.2 COVID exceptionalism?
- 24 3.2 Reconciling COVID exceptionalism with relatively stable attitudes

4. Conclusions

References

Appendices

- 32 Appendix A: The YouGov panel
- 33 Appendix B: The full set of YouGov welfare policy trackers
- 37 Appendix C: Structural Topic Model of open text responses

1. INTRODUCTION

COVID-19 presents a puzzle for those studying public attitudes to working-age benefits. A priori, COVID-19 seemed to be a moment where support for social security would rise: a time of apparent increased solidarity in the face of a collective crisis; of clearly ‘deserving’ claimants with no control over their need; of positive and prominent public discourses about social security; and of increases in direct experiences of the benefits system, particularly from social groups that were previously unlikely to claim. In most countries, the generosity of benefits increased, which may have also created a feedback loop to more positive attitudes.

Yet the limited evidence so far suggests that attitudes have been stable. Early analysis of UK trend data shows no systematic change (Curtice et al., 2020; Hicks, 2020), while published Dutch research shows possible declines in support for state intervention (Reeskens et al., 2021). Experimental evidence from Germany, the Netherlands, Italy and Spain suggests that priming people to think about COVID-19 reduces their willingness to support higher taxes to pay for more generous unemployment benefits (Daniele et al., 2020). Alongside this, we are starting to see attempts to undo COVID-19-related increases in benefit generosity, at least in some countries.¹

One explanation for this is that we were simply wrong to expect attitudes to become more positive. Previous studies on the effect of economic crises on benefits attitudes have yielded mixed results (see below). Moreover, the most recent international economic crisis in 2008–9 did not produce a groundswell of support for more generous benefits (Roosma, 2020), and in many countries the legacy of the crisis was public sector austerity, including cuts to working-age benefits. In this interpretation, attitudes have not changed, and it was historically short-sighted to have expected them to do so.

There is an alternative explanation, however: that there has been a change in benefit attitudes that the tentative initial findings have missed. This may be because the initial findings are simply wrong; for example, some surveys were forced to move from face-to-face/telephone modes to online surveys, which we know will affect the answers people give, distorting trends. But it may also be that standard attitude measures are not well-designed for capturing rapid attitudinal change. The apparent stasis captured by general attitude measures may conceal substantial changes in more specific attitudes and beliefs engendered by the COVID-19 pandemic.

In this report, we investigate recent trends in benefit attitudes in the UK. The UK is an ideal case study: compared to many other countries it has historically seen more negative attitudes and more negative media representations of benefit claimants (Larsen and Dejgaard, 2013), providing a stronger contrast to the more positive

¹ E.g. in Australia; see theconversation.com/the-50-boost-to-jobseeker-will-take-australias-payment-from-the-lowest-in-the-oecd-to-the-second-lowest-after-greece-155739 [accessed 1/6/2021]

representations during COVID-19. The UK also has ideal data to explore these questions. In this report, we draw upon (i) bimonthly trend data collected by YouGov from 2019-2021, which unlike previous studies has a consistent survey mode and multiple time points; and (ii) a nationally representative survey we conducted as part of the Welfare at a (Social) Distance project in June 2021, uniquely exploring COVID-19-related attitudes in unique level of detail.

We explain the data and methods below, but first we review in more detail previous research on attitude change during the pandemic, and during previous crises.

1.1 WHY WOULD COVID-19 CHANGE ATTITUDES?

At the outset of the COVID-19 pandemic, there were at least four reasons to think that attitudes to benefits would become more positive: social solidarity, perceived deservingness of claimants, experiences of claiming, and path dependency.

First, social solidarity seemed to rise in the face of an all-consuming collective crisis, which we may expect to increase support for the welfare state (Burchardt, 2020). Many UK commentators made the analogy to the 1942 Beveridge Report, which laid the foundation for the modern welfare state in the midst of wartime Britain (e.g. *Financial Times*, 2020). This was not just idle speculation; there was tangible evidence to suggest that solidarity was rising at the start of the pandemic, ranging from the creation of formal and informal support groups² to widespread public expressions of support for the NHS and other key workers.³ Emerging research evidence seemingly confirmed an increase in generalised trust in others by May 2020 (Parsons and Wiggins, 2020) and a feeling that COVID-19 has increased people's concern for each other (More in Common, 2020a; b).

Second, it seems likely that benefit claimants during COVID-19 were seen as more 'deserving' of support. Perceived deservingness is a crucial determinant of public attitudes to benefits, and reflects a mixture of claimants' perceived control (or blameworthiness), attitudes (e.g. gratitude), reciprocity (having paid into the system), identity (are claimants 'one of us'?) and need (how much genuine hardship would claimants endure without help) (which constitute the 'CARIN' model of deservingness; see van Oorschot and Roosma, 2017). In economic crises in general, it is often argued that claimants will be seen as more deserving – partly because of identity (middle-class claims tend to become more common), but primarily because claimants have less control over their situation (Roosma, 2020; Uunk and van Oorschot, 2019). This effect seems likely to have been even stronger during COVID-19, where large swathes of the economy were required to cease activity for the public good. Perceptions of deservingness also depend on media coverage and political rhetoric (although to a debated extent; Baumberg et al., 2012), and anecdotally it seems that discourses were unusually positive during COVID-19, particularly by UK standards.

² blogs.lse.ac.uk/politicsandpolicy/covid19-mutual-aid-solidarity

³ clapforourcarers.co.uk

Third, the proportion of the population who had some experience of claiming rose dramatically. Attitudes may have become more generous among those who found themselves claiming, and among their family and friends (Hedegaard, 2014). Attitude change is not just a matter of self-interest, but – in a country where ‘benefit myths’ are widespread (Geiger, 2018) – direct or indirect experiences of claiming seem likely to increase sympathy (Edmiston, 2018; Geiger, 2017). Beyond the population of new claimants (and their friends/family), large numbers of people experienced a form of welfare through the Coronavirus Job Retention Scheme (‘furlough’) or the Self Employment Income Support Scheme (‘SEISS’). This could potentially lead to increased empathy with ‘conventional’ welfare claimants, though furlough/SEISS were administered outside the traditional benefits system which may have muted this effect. Even more broadly, new claimants during the pandemic were generally more economically advantaged than pre-existing claimants (Edmiston et al., 2020), so the perceived chances of claiming may have risen among higher socioeconomic groups. Overall, COVID-19 seems likely to have created a broader constituency of people feeling that they have some stake in the benefits system (of critical importance for welfare state policy; Rehm et al., 2012)

Finally, benefits in many countries played a key role in dealing with the social and economic impact of the pandemic. Benefits were commonly extended in scope (e.g. Finland, Spain, Portugal, Germany) or – particularly in the case of Anglo-Saxon countries – made more generous (see e.g. Béland et al., 2021:5-8; Moreira and Hick, 2021:11). In the UK, the basic element of UC was temporarily increased by £20/week, the housing entitlement made more generous, eligibility for self-employed people was broadened, and the conditions of claiming made less onerous. As policies change, they also change the political context for the possible policies that follow (‘path dependency’). In particular, increased generosity in-and-of-itself may make attitudes more positive (Burchardt, 2020); as Pierson (1996:144) notes, *“there is a profound difference between extending benefits to large numbers of people and taking benefits away”*.

1.2 CHALLENGING EXPECTATIONS OF INCREASED GENEROSITY

In the previous section, we have outlined a number of reasons to expect COVID-19 to yield more generous attitudes to benefits. However, there are also reasons to challenge such expectations; reasons that might lead us to expect attitudes to remain the same, or even to become more negative. For example, any improvements in solidarity may have been fleeting. Although social trust initially rose, by June 2020 British people were reporting that COVID-19 had led to declines in social trust and feelings of social division, partly due to fears of contagion and perceived rule infringements by others (More in Common, 2020a; b). The same may also be true for benefit attitudes. Even at the outset of the pandemic, political scientists like Joakim Palme were warning that “the political effects of crises tend to wash away fairly quickly” (in Burchardt, 2020; see also Kenworthy and Owens, 2011).

One source of knowledge about the potential effects of the pandemic is previous research on the effects of recessions on attitudes, where many of the same

mechanisms apply (notably: less perceived responsibility for unemployment, and a broadening of the claimant population). If we look at the empirical evidence on whether recessions make benefits attitudes more positive, then the answer is mixed,⁴ and often ‘no’ (Kenworthy and Owens, 2011; Roosma, 2020; Uunk and van Oorschot, 2019). Even within the UK, the economic crisis of 2008–9 led to more negative rather than positive benefits attitudes (Taylor-Gooby, 2013), and was followed by a period of particularly negative media coverage of benefits (Baumberg et al., 2012). There is some evidence that recessions make 18–25-year-olds more favourable about redistribution – the worldviews of young adults being most susceptible to change, as well as being more affected by at least some recessions (Bell and Blanchflower, 2011)– but this seems to be restricted to people’s ingroup (Cotofan et al., 2021), and moreover this cohort is unlikely to be large enough to influence population-level attitudes (Kenworthy and Owens, 2011).

Why is this evidence so mixed, given the compelling arguments of the previous section? One explanation is that economic crises are associated with financial strain, at both an individual and societal level (Burchardt, 2020; Roosma, 2020; Uunk and van Oorschot, 2019). This can cause attitudinal shifts that go against the greater perceived deservingness of claimants in crises, as individual and collective financial insecurity can foster austerity, self-interest, competition and resentment (see also Durr, 1993; Hoggett et al., 2013). Uunk and van Oorschot (2019) partially disentangle these two mechanisms using Dutch time-series data. They find that unemployment levels are associated with more positive welfare attitudes (belief that unemployment benefits were not sufficient), but economic declines are associated with less positive attitudes. Household and state financial strain during COVID-19 may therefore act to counteract the effects we describe above. Claimants may be seen as in some sense more ‘objectively’ deserving; but their situations are being judged by a less altruistic public.

Another reason for attitudinal stability may be that benefit claims during recessions are seen as exceptional – crisis-associated claimants may be seen by much of the population as more ‘like us’ and blameless, but this is contrasted to pre-existing claimants who are seen as ‘other’ and blameworthy. This is exactly what Erler (2012) found in her analysis of US newspaper reporting (building on Gilens’ earlier work). People in poverty were represented as more deserving during the 2008–9 Great Recession, but this was often explicitly in contrast to ‘normal’ poverty. As Erler concludes (p185),

“The discourse that emerged reified class differences by portraying the ‘newly poor’ as fundamentally different and more deserving of policy action than those who were in poverty prior to the economic crisis. While this new discourse is more likely to employ a structure/context narrative, it is

⁴ There are however some exceptions to this (most influentially, Blekesaune, 2007). As Laenen and van Oorschot (2020) summarise the literature on economic crises, “*Most prior studies in the field would consider change in welfare attitudes to be the most likely scenario... The direction of such crisis-induced transformations, however, continues to be debated: do welfare attitudes become more positive, or more critical, in times of economic downturn?*”

a discourse that is largely reserved for individuals who have the markers of middle-class identity.”

COVID-19 likely felt like an even greater departure from normal times than the Great Recession, particularly given the institutional separateness of furlough/SEISS (above), and anecdotally we can see similar contrasts of COVID-19 versus ‘normal’ claimants in UK media coverage.⁵ We may therefore see ‘Covid exceptionalism’ (Summers et al., 2021b), where claimants during the pandemic are bracketed away from other benefit claimants, leaving deeper-lying attitudes unchanged.

Finally, while public attitudes sometimes follow policy changes (because it is harder to cut a benefit than to increase/extend it; see above), benefit attitudes are also ‘thermostatic’ (Curtice, 2010; Curtice et al., 2020; Wlezien, 1995). The analogy of the thermostat captures how attitudes can move in the opposite direction to policy: if we make a room hotter, then we are more likely to want the temperature to go down. In this case, (i) after making benefits more generous, the public may be less supportive of raising benefits yet further; and (ii) after temporarily relaxing controls on applications and removing conditionality (Summers et al., 2021a), the public may want stricter controls.

1.3 THE EMERGING EVIDENCE (AND ITS LIMITS)

In the face of these multiple mechanisms working in different directions, empirical evidence is crucial. At the time of writing, four studies have examined changes in welfare attitudes in Europe over the course of the pandemic. None of these have found a positive change in benefit attitudes. Ares et al. (2021) find no changes 2018–2020 in preferences for benefits retrenchment/expansion (or redistribution/left-right ideology) in Germany, Spain and Sweden. Reeskens et al. (2021) find declining support 2017–2020 for state intervention in the Netherlands (that is, whether ‘the state should take more responsibility to ensure that everyone is provided for’). Similarly in the UK, people in July 2020 were slightly more likely than in 2019 to say that *“if welfare benefits weren’t so generous, people would learn to stand on their own two feet”* (Curtice et al., 2020) while there was less desire for higher taxes/more spending (Hicks, 2020), although the extent of change was small, and both studies interpret this as evidence against COVID-19-induced change.

However, there are several caveats which limit the conclusions we can draw from these studies. First, it is well-known that survey mode – e.g. face-to-face vs. online – affects how people respond to certain questions (Atkeson et al., 2014). Face-to-face pre-pandemic surveys had to change survey mode during COVID-19 (Curtice et al., 2020), potentially biasing their results in unknown directions. Second, three of the four studies look only at two time points (Ares et al., 2021; Curtice et al., 2020; Reeskens et al., 2021), comparing one pre-pandemic moment (2017/2018/2019) to one moment during the pandemic (May/June/July 2020), with no follow-ups beyond the initial lockdown. This prevents us from understanding the dynamics of attitude change

⁵ See e.g. BBC News 15/12/2020, [“Universal credit: ‘I felt guilty claiming benefits for the first time’”](#)

during the pandemic. Third, and most importantly, none of the previous studies ask directly about people's perceptions of the COVID-19 wave of benefit claimants. This prevents any investigation of the possible 'Covid exceptionalism' described above.

In this report, we present novel UK data that responds to all three limitations – it uses a consistent survey mode, analysing attitudes every two months from 2019–2021, and supplemented by a detailed survey specifically addressing attitudes towards the COVID-19 wave of claimants. We explain these data, and our results, in more detail below.

2. DATA SOURCES

We employ two sources of data:

- The first source is YouGov's 'Welfare Tracker' series 2019-2021. These unique repeated cross-sectional surveys have been conducted approximately every eight weeks from June 2019 to June 2021, with 1,600-1,700 responses per wave. The dataset analysed here includes 13 waves and 21,600 responses. Cross-tabulations (although not the underlying dataset) are publicly available.⁶ In this report, we use this data to analyse trends in welfare attitudes across the pandemic. We describe the specific measures below.
- Our second data source is a novel survey we commissioned (again through YouGov) as part of the Welfare at a (Social) Distance project, which included a number of measures intended to compare public attitudes towards people claiming benefits during the pandemic versus those claiming prior to COVID-19. This was conducted in May/June 2021 and included a sample of 3,429 people; ethical approval was given by the University of Salford. The full dataset and questionnaire will shortly be publicly available from the UK Data Service (SN6989),⁷ and we provide full question wording when we discuss the results of each measure below.

Both data sources use the YouGov online survey platform and opt-in panel. Opt-in panels – particularly YouGov, which is the largest in the UK – have been a common choice for pre-pandemic attitudinal studies (e.g. Aarøe and Petersen, 2014; Harell et al., 2016; Jensen and Petersen, 2017; Kootstra, 2016). Their value was even greater during the pandemic, as they enabled rapid data collection during lockdowns when face-to-face research was impossible. They also contained pre-pandemic attitudinal measures on benefits (the 'Trackers'; see above), which could be continued during COVID-19 using identical methods. Further detail on the YouGov panel is given in **Appendix A**.

To create nationally representative estimates, YouGov select panel members to participate according to known population totals, and then further apply non-response weights to the achieved sample. There are two limitations to this – the panel under-represents those with weaker English language skills or digital access/skills, and there will be biases where the weighting variables do not fully capture factors that influence both survey response and the phenomenon under study. Indeed, non-response in opt-in panels resulted in prediction errors for the UK 2015 General Election (Sturgis et al., 2016:67). Nevertheless, YouGov has performed well in predicting election results since, and the data can be considered broadly representative of the UK population (see **Appendix A**).

⁶ yougov.co.uk/topics/politics/explore/academic_area/Welfare_Policy

⁷ beta.ukdataservice.ac.uk/datacatalogue/studies/study?id=8689

3. RESULTS

3.1 TRENDS IN ATTITUDES

We focus on two YouGov Welfare Tracker questions that capture key policy-related attitudes (deservingness and support for more generous benefits):

Deservingness

“Thinking about people who receive welfare benefits, including disability benefits, out of work benefits and benefits to support people in low paying work, what proportion of people receiving benefits do you think are genuinely in need and deserving of help?”

All or almost all people who receive benefits are genuinely in need and deserving of help

The majority of people...

Around half of people...

Only a minority of people...

Hardly anyone...


Level of benefits

“And thinking about the level of benefits, do you think they are too high, too low, or is the balance about right?”

These measures also differ in how ‘thermostatic’ they are (see above); benefit generosity is a relative measure (compared to current perceived levels of generosity, which changed at the start of COVID-19), whereas deservingness can be considered an absolute measure. For analysis, we collapsed responses to each question into two categories, contrasting respondents who said that half or fewer claimants were deserving / who said benefits were too high versus those giving any other valid response.

These trends are visualised in **Figure 1** below (the light blue area in each graph represents 95% confidence intervals around each estimate). To contextualise the trends, the grey areas represent the first and second waves of the pandemic in the UK.⁸

⁸ We have defined the first wave of the pandemic as running from the 16th of March 2020 (when the Prime Minister issued the statement that “now is the time for everyone to stop non-essential contact and travel”) to the 15th of June 2020 (when non-essential shops were allowed to re-open). We have defined the second wave as running from the 14th

Figure 1: Trends in anti-welfare attitudes 2019–21

Source: YouGov Welfare Tracker surveys (total N=21,602)

These trends show that attitudes became somewhat more generous as the UK entered the first wave of the pandemic. Attitudes appear to have rebounded quickly, however, with attitudes in summer 2020 similar to those a year earlier. The onset of the second wave of the pandemic appears to have prompted another shift towards more generous attitudes, though this is less noticeable in the question on whether benefit payments are too high than for the deservingness measure (in general, opinions on benefit levels appear much more stable than those concerning deservingness).

For both measures, the changes associated with the two waves of the pandemic are relatively small in absolute terms – not more than five percentage points. For context, it is worth comparing these pandemic-associated changes with the ‘Christmas effect’ in each graph. The ‘Christmas effect’ is an apparent softening of welfare attitudes around the Christmas period which has been observed, for example, in research on sanction rates (National Audit Office, 2016). It is notable that, for each measure, the change in attitudes associated with each wave of the pandemic is not markedly different in scale from this ‘Christmas effect’.

In order to determine whether the differences may have arisen due to sampling error, we combined responses from the individual waves into five separate periods:

- **Pre-pandemic:** data collected prior to 16th March 2020, excluding December 2019 (to exclude the Christmas effect) (N=6,769)
- **First wave:** data collected 16th March 2020 to 15th June 2020 (N=3,237)
- **Summer 2020:** data collected 15th June 2020 to 14th October 2020 (N=3,286)

- **Second wave:** data collected 14th October 2020 to 29th March 2021 (N=4,962)
- **After second wave:** May 3rd 2021 data collection (N=1,653)

We used linear probability regression models to compare responses to each question across the five periods. **Table 1** shows the percentage point change in responses between these time periods, along with 95% confidence intervals:

Table 1: Change in anti-welfare attitudes between time periods

	Half or fewer are deserving		Benefit levels are too high	
Pre-pandemic level	34.2%	[32.9 to 35.5]	17.8%	[16.8 to 18.8]
Changes during COVID				
Pre-pandemic to First wave	-3.2	[-5.4 to -0.9]	-3.1	[-4.8 to -1.3]
First wave to Summer 2020	+4.2	[1.5 to 6.9]	+2.4	[0.4 to 4.5]
Summer 2020 to Second wave	-4.4	[-6.9 to -1.9]	-1.9	[-3.8 to 0.05]
Second wave to After second wave	+0.2	[-2.7 to 3.3]	-0.3	[-2.5 to 2.0]
Total changes to May 2021	-3.1	[-6.0 to -0.2]	-2.8	[-5.0 to -0.6]

Table shows percentage point change with 95% Confidence Intervals.
Source: YouGov Welfare Tracker surveys (total N=21,602)

Table 1 shows that attitudes during the first wave of the pandemic were significantly more generous (people were less anti-welfare) than during the pre-pandemic period – but attitudes during Summer 2020 then became significantly less generous. Attitudes during the second COVID-19 wave were significantly more generous than those during the Summer of 2020 (although the change for the benefit levels question is more tentative). Finally, attitudes after the second wave did not differ significantly from attitudes during the second wave.

Comparing the single post-2nd-lockdown wave with the pre-pandemic period, the public were less anti-welfare than they were before – but only slightly (by about 3 percentage points on both measures).

Although we focus here on only two of the YouGov Welfare Tracker measures, eight other Trackers are available. Trends for these measures are given in **Table A5**. These show a similar pattern to that reported above: more pro-welfare attitudes during the first lockdown, followed by a rapid bounceback in summer 2020, followed by a further, smaller shift towards more pro-welfare attitudes during the second lockdown. The extent of these changes differs between measures however, with some measures showing larger changes than in **Table 1** (for the view that people out of work get too much support, or that the benefits system works well), and others showing no change over the period (that benefit conditions are not strict enough,⁹ or whether disabled people get too much support). It is noteworthy that despite the increased

⁹ This is possibly a response to the relaxing of eligibility criteria associated with COVID-19

generosity of UC during COVID-19, questions about whether benefits are too high/people get too much support – which are relative to existing benefit levels – nevertheless move in a pro-welfare direction.

In summary, the two waves of the pandemic were associated with a statistically significant increase in welfare generosity. However, these changes were small in scale, and the rapid rebound observed in the summer of 2020 suggests that they are unlikely to be long-lasting.

Is there a ‘floor effect’ preventing attitudes from becoming more generous during the pandemic?

Pandemic-related trends cannot be understood without putting them in the context of longer-term attitudinal shifts. In April 2013, YouGov asked respondents in a nationally representative survey the benefit levels Tracker question we analyse above (are benefits too high, too low, or about right; n=1,991). In 2013, 37.3% of respondents [95% CI: 34.7% to 40.1%] said that benefit levels were too high, far more than the 17.8% [95% CI: 16.8% to 18.8%] who said this in the 2019/20 pre-pandemic period. This shows very large long-term shift towards more generous welfare attitudes before COVID-19, as has been found elsewhere (Hudson et al., 2020). The pandemic-associated changes we describe above are therefore small relative to this general softening of welfare attitudes.


One interpretation of recent trends, therefore, is that COVID saw a continuation of existing pro-welfare attitudes – as Curtice et al. (2020:23) put it, “*public opinion had already become somewhat more receptive to a more interventionist role for the state... that revised mood seems still to have been in place in the early months of the pandemic.*” Indeed, it might be the case that there was little space for attitudes to become more positive at a population-level, in a context when attitudes have already changed. However, it is possible that COVID-19 had a more profound effect on the attitudes of people who began the pandemic with more anti-welfare views; for example, political conservatives.

To test this, **Figure 2** shows trends in perceptions of benefit levels across the pandemic according to political affiliation, contrasting respondents who voted Conservative with those who voted Labour in the 2019 General Election. Unsurprisingly, Conservative voters begin with considerably harsher welfare attitudes than Labour voters, and this remains the case throughout the pandemic. More importantly, changes during COVID-19 partly support our expectation: the attitudes of Conservative voters do appear to respond more strongly to the first and second lockdowns than those of Labour voters.¹⁰ This may suggest a ‘floor effect’ among Labour voters – so few believed that benefit levels were too high prior to the pandemic that there is little scope for further aggregate increases in generosity. Nevertheless, even the changes among Conservatives are relatively small in size.

¹⁰ If we test this statistically, then Conservatives saw a significantly greater change in attitudes than Labour voters in the first wave. However, for subsequent shifts, the differences between Labour and Conservative voters are not statistically distinguishable. This is likely because the softening of attitudes among Conservatives did not occur until January 2021 – i.e. in the midst of the second wave of the pandemic. A full contrast of these trends is given in **Appendix A, Table A6**

We conducted the same analysis for the deservingness measure (not shown), which followed a very similar pattern.

Figure 2: View that the level of benefits is too high, split by political affiliation



Source: YouGov Welfare Tracker surveys (total N=13,118)

3.2 COVID EXCEPTIONALISM?

The evidence so far suggests that there were limited changes in people's attitudes to benefits during COVID-19, and that even these limited changes were not always sustained after the lockdowns ended. As we argue above, these findings are consistent with two differing accounts: first, that the pandemic has not meaningfully changed people's welfare attitudes; and second, that they have changed, but in ways which bracket COVID claimants (people who began claiming during the pandemic) away from claimants more generally, which standard measures do not capture. We term this 'COVID exceptionalism' (Summers et al., 2021b), and in this section we investigate it using our novel survey of welfare attitudes conducted in June 2021.

We examined possible COVID exceptionalism using three different approaches. First, we asked respondents directly about differences in deservingness between COVID-19 and pre-pandemic claimants (adapting the wording of the YouGov deservingness Tracker, which focuses particularly on ‘need’):

“Compared to people who were claiming benefits before COVID-19 (i.e. before March 2020).... Do you think that people that claimed welfare benefits during COVID-19 were more likely to be in need and deserving of help, or less likely, or about the same?”

Responses were given on a five-point scale from “Far more...” to “Far fewer people who claimed during COVID-19 were genuinely in need and deserving of help” (as well as ‘other’ and ‘don’t know’ options). As **Table 2** shows, around half of respondents (50.6% [95%CI: 48.8 to 52.5%]) felt that COVID-19 claimants were more likely than pre-COVID claimants to be genuinely in need/deserving. Around a quarter of respondents felt that both groups were equally likely to be deserving, and only a small minority (7.1%) felt that COVID-19 claimants were less likely to be deserving.

Around one-fifth of responses were ‘Don’t knows’. Results from the other measures we report in this section suggest that respondents are most likely to resort to ‘Don’t know’ when asked to make sweeping generalisations of fact (e.g. what proportion of a group of claimants are to blame for losing their jobs), rather than entirely subjective opinions. If we exclude respondents who chose the ‘Don’t know’ option, a clear majority of the remaining respondents (61.7%) felt that COVID-19 claimants were more likely to be deserving than pre-pandemic claimants.

Table 2: Were those claiming during COVID-19 (vs. before COVID-19) perceived as being more likely to be genuinely in need and deserving of help?

Response	%	95% confidence interval
Far more people who claimed during COVID-19 were genuinely in need and deserving of help	26.3	[24.7 to 28.0]
Slightly more...	24.3	[22.8 to 25.9]
About the same...	24.5	[22.9 to 26.1]
Slightly fewer...	3.7	[3.0 to 4.4]
Far fewer...	3.4	[2.8 to 4.1]
Don’t know / Other	17.8	[16.3 to 19.4]

Source: WASD general population survey June 2021, n=3,476.

Second, rather than asking about the relative deservingness of pre/during-COVID-19 claimants, we asked respondents to judge pre/during-COVID-19 claimants separately. We asked specifically for judgements of blameworthiness. Blame/control is central to wider deservingness judgements (van Oorschot and Roosma, 2017), and the perceived blameworthiness of COVID-19 claimants have has been a matter of public

debate in the UK.¹¹ To investigate this in more detail, we looked at both blame for losing a job (retrospective blame), and blame for being unable to find a job and leave benefits (prospective blame). Respondents were randomly assigned to see either the retrospective or prospective version:

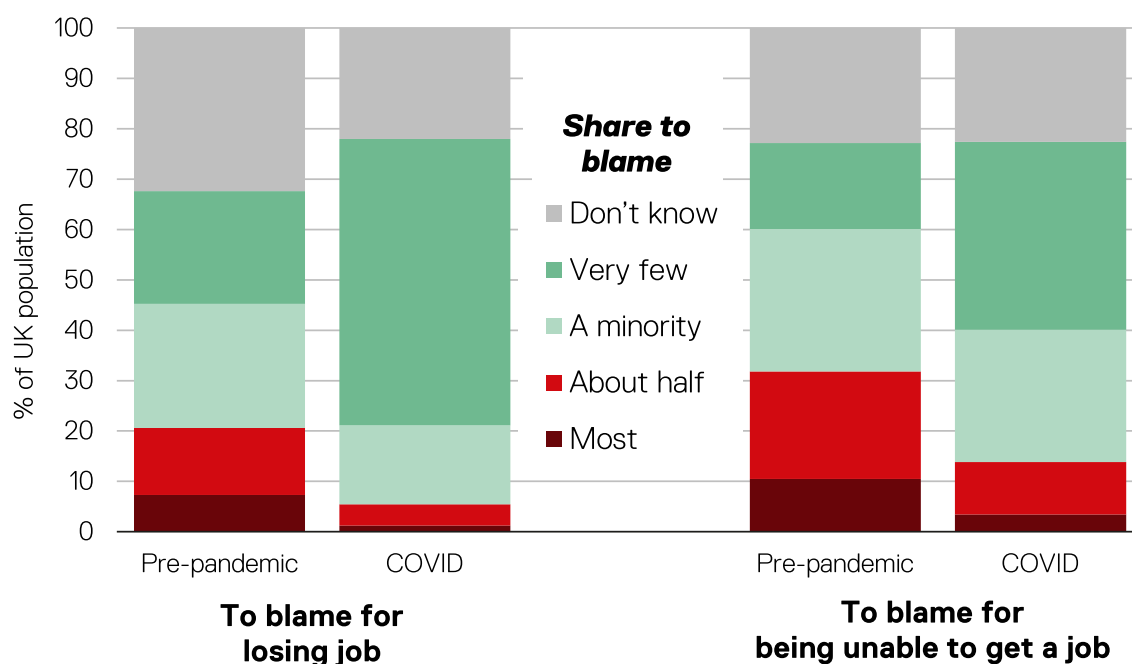
“On average, how much do you think each of the following groups are to blame for [losing their jobs/being unable to get a job and leave benefits]?”

1. People [who lost their job/on benefits and unable to get a job] before Covid-19 (before Feb 2020)
2. People [who lost their job/on benefits and unable to get a job] during Covid-19 (since Mar 2020)

In both cases, response options were that ‘most / about half / a minority / very few are to blame’ (plus a ‘don’t know’ option). As for the previous question, in the face of a question that asked for broad generalisations, large minorities selected ‘don’t know’.

Figure 3 shows that COVID-19 claimants were much less likely to be blamed for either losing their jobs or for being unable to find a job and leave benefits. Only 5.4% of respondents felt that half or more of COVID-19 claimants were to blame for losing their jobs, compared to 20.5% of pre-pandemic claimants (a difference of 15.1 percentage points [95% CI: 12.4 to 17.7pp]). Similarly, 13.8% of respondents thought that half or more of COVID-19 claimants were to blame for being unable to find a job, compared with 31.8% who said the same about pre-pandemic claimants – an 18.0 percentage point difference [95% CI: 15.0 to 21.0pp].

Figure 3: Were people claiming during COVID-19 perceived as being less to blame for losing their jobs/being unable to find a job?



¹¹ A media headline claimed that nearly half of the public ‘blame job losses unemployed people’ during COVID-19, which has been argued to misinterpret the underlying research; fullfact.org/economy/coronavirus-job-loss-guardian

Third, responses to this direct question on blame may still be affected by social desirability bias (some respondents may believe that one group of claimants is less blameworthy, but not wish to say so explicitly). To address this, we included a vignette at the end of the questionnaire (after several unrelated questions), in which a hypothetical claimant was randomly described as claiming (i) pre-pandemic or (ii) during COVID-19. The vignettes are not only more tangible, but due to randomisation also permit stronger causal inference as to the effect of COVID-19 on responses (Mutz, 2011). An example full-text of the vignette is given below:

John is a 35 year old man living in England. In April 2019, before the COVID pandemic, John’s company went bust and John lost his job. John then started claiming Universal Credit, and three months later (in July 2019, before the COVID pandemic), had not been able to get a job.

Do you think John is to blame for being out of work three months later (in July 2019, before the COVID pandemic)?

As well as varying the time period, we also varied the claimant’s gender/name (John/Liz), age (35/60), and whether we provided an explanation for job loss (“John/Liz’s company went bust and John/Liz lost his job” / “John/Liz lost his job”). This latter manipulation allowed us to determine whether vaguely-described claimants were subject to greater COVID-19 effects than those more explicitly cued as ‘deserving’ (due to having lost their job for a reason outside their control).

As shown in **Table 3**, the majority of respondents felt that the vignette claimant was not at all to blame for losing their job and still being out of work three months later. However, consistent with the previous results, respondents who read about a COVID-19 claimant were 18.8 percentage points more likely to consider the claimant to be entirely blameless [95% CI: – 22.3 to – 15.2ppt].

Table 3: Were vignette claimants who lost their jobs during COVID perceived as less blameworthy?

	Pre-COVID claimant		COVID-19 claimant	
	%	95% CI	%	95% CI
Not at all to blame	59.1	[56.4 to 61.7]	77.9	[75.4 to 80.1]
A little to blame	18.7	[16.6 to 20.9]	10.3	[8.8 to 12.2]
Mostly to blame	4.4	[3.3 to 5.7]	1.8	[1.2 to 2.8]
Completely to blame	1.0	[0.6 to 1.6]	0.9	[0.4 to 1.7]
Don’t know	16.9	[14.9 to 19.1]	9.1	[7.5 to 10.9]

Source: WASD general population survey June 2021, n=3,480

Examining the other elements of the vignette, we found that respondents were not substantially more or less sympathetic to young subjects, male subjects, or subjects

who had lost their jobs due to the closure of their employer (vs. no explanation given). Against our expectations, pre-pandemic claimants were not perceived more similarly to COVID-19 claimants if a clear explanation for their job loss was given (a 2.9 percentage point difference in the marginal effect of COVID-19 claimant vs. pre-pandemic claimant; 95% CI: – 3.8 to 9.7pp), nor did the pre-pandemic vs. COVID-19 contrast differ according to age or gender.

Taken together, these three measures provide strong evidence that COVID-19 claimants were seen as substantially less blameworthy for being out of work, and as being substantially more likely to be deserving of help. It is worth noting here that we do not show that pre-COVID claimants were viewed harshly: only a minority of respondents blamed pre-COVID claimants for their situation. In other words, pre-COVID claimants were viewed quite sympathetically, and COVID claimants simply more so.

Exploring differences between COVID-19 and pre-pandemic claimants

To explore how COVID exceptionalism reflected the different dimensions of deservingness (and the other considerations raised in the literature review), we asked the following open text question:

“What do you think are the main differences between people claiming benefits before COVID-19 and people claiming during COVID-19, if any?”

Around a quarter of participants (26%) responded to this question with some version of ‘There are no differences’. A further 13% gave some version of a ‘don’t know’ or other non-response (such as a non-meaningful string of characters). After excluding these respondents, we analysed remaining responses using Structural Topic Models (STMs) using the `stm` package in R. STMs are used to identify patterns in corpora of text documents (such as open-text survey responses) (Roberts et al., 2014). The STM groups these responses into ‘topics’ through a machine-learning procedure: similar responses using overlapping sets of words were grouped together under the same topic, whereas responses using dissimilar words were grouped under different topics. The identification of these topics allows us to define the ‘types’ of answers respondents gave to our question about the differences between COVID-19 and pre-pandemic claimants. Further details on the STM procedure are given in **Appendix C**.

We found that dividing responses into seven categories (topics) was the most parsimonious solution (see **Appendix C**). **Table 4** provides (i) the estimated proportion of responses falling into each category, (ii) the seven words most strongly associated with each category, and (iii) illustrative real responses strongly associated with each category. Our summary of the topics is as follows:

- Responses in the most popular category (Topic 7) focused primarily on the idea that COVID-19 claimants had different characteristics than pre-COVID claimants: they were people who may have had “established careers” and who may never have expected to claim benefits. A number of responses falling into this category also

noted that COVID-19 claimants had a “genuine reason to claim” because they had lost their jobs “through no fault of their own.”

- The distinction between ‘deserving’ COVID-19 claimants and ‘undeserving’ pre-pandemic claimants came through most strongly in Topic 6. Responses in this category focused strongly on the idea that COVID claimants had been employed beforehand, thereby demonstrating a willingness to work (they were “obviously...trying to work and look after themselves and their families”). By contrast, pre-COVID were identified as potential “spongers” who “did not want to work.”
- Responses in Topic 4 (the second most popular category) tended to make a less clear distinction between COVID-19 and pre-pandemic claimants in terms of deservingness. Instead, they tended to focus on the idea of the pandemic as an extreme exogenous event with devastating economic consequences. The pandemic was a “naturally caused event” that caused “absolute devastation”, especially to sectors like hospitality and tourism. However, there was an implicit differentiation here between the unavoidable job losses caused by COVID-19 and the potentially avoidable reasons for job loss pre-pandemic.
- This distinction was made more explicit in responses falling under Topic 2. Responses in this category often contrasted the variable reasons for unemployment pre-pandemic (“bump[s] in the road”, “complex societal factors”) with the massive decrease in job availability during the pandemic.
- Responses falling under the remaining topics (1, 3, and 5) focused less on differential reasons for unemployment:
 - » Topics 1 and 3 focused strongly on the impact of the pandemic on living standards, with responses under Topic 3 tending to emphasise the challenge of living on benefits, and responses under Topic 1 focusing on the sudden loss of income experienced by many COVID-19 claimants.
 - » Responses falling under Topic 5 were quite variable. However, there was a strong emphasis on the potential for fraud; among non-COVID-19 claimants (people “gaming the system” in more normal times) but also among COVID-19 claimants – with a number of respondents fearing that a lack of oversight might lead to rampant fraud among new pandemic claimants.

Taken together, this suggests that many respondents did draw a clear distinction between COVID and pre-COVID claimants. With a few exceptions (concerns about fraud among ‘unsupervised’ new claimants), COVID claimants were generally identified as the unfortunate victims of a situation beyond their control – as people who would otherwise be gainfully employed were it not for the economic devastation wrought by the pandemic. While many respondents also described pre-COVID claimants in sympathetic terms, they were much more likely to be labelled by respondents as potential ‘scroungers’ or ‘spongers’, who were less likely to have a genuine reason to claim.

It is also worth noting, however, that many respondents, when asked to describe the main differences between COVID-19 and pre-pandemic claimants, did not focus on deservingness per se. Rather, many respondents focused on the broader impact of the pandemic at a more abstract level, or on the more human cost of the economic disruption on the living standards of both existing and new claimants alike.

Table 4: Perceived differences between COVID-19 and pre-pandemic claimants, derived from the STM analysis

Topic	%	Key words ¹² (highest probability of being included)	Illustrative responses
7	19%	peopl*, claim, lost, mani*, unemploy*, genuin*, lose, support	“The people who were claiming benefits before COVID-19 might not have jobs, and the people claiming benefits during COVID-19 may have lost their jobs because of COVID-19”
4	18%	pandem*, busi*, furlough, differ*, close, help, time, forc*	“One is due to the normality of a given economic system or government action, a deliberate effect. The pandemic was a naturally caused event that effected immediate responses and a relaxation of any rules that exist in normal times”
5	15%	need, think, will, find, redund*, system, just, live	“During COVID, benefit fraud investigations have been largely suspended – and fraud’s endemic as a result. Pretty much every one was claiming – so easy to scam the system with a hard luck story. And no one dares say anything because they don’t want to be cancelled on Twitter if they speak the uncomfortable truth”
6	14%	work, due, get, loss, fault, lockdown, may, claimant	“Those who found themselves unemployed due to the covid economic crisis were unfortunate to be in the affected industries. They worked before covid and likely want to get back into work ASAP. Those on benefits before don’t want work before or after covid”
1	12%	benefit, lot, covid-, incom*, less, made, now, probabl*	“Some that claimed during the pandemic were working people who had a steady income possibly good income and then lost their job or were furloughed therefore dramatically reducing income. Those who claimed benefits before would have noticed no change to their benefits or income”
3	12%	covid, never, increas*, chang*, lack, paid, self, also	“i have an impression that the covid hit an amount of claimant very badly due to lack of other income sources, school meals, school clubs etc”
2	10%	job, employ, like, avail*, affect, hour, long, first	“Job availability fell so benefits increased – this is what it is for. The main difference was job availability for covid. Potentially pre covid there were greater complex societal factors leading to unemployment such as lack of education, stability, sound mental and physical health combined with lack of job availability. This didn’t go away but covid worsened job availability and therefore more people were unemployed”

Source: n=2,018 free text responses (excluding don’t knows, ‘no differences’, and unintelligible responses)

¹² The STM modelling procedure reduces words to stems before analysis; for example, ‘increased’ and ‘increasing’ become ‘increas*’, with the asterisk indicating a stemmed word.

3.2 RECONCILING COVID EXCEPTIONALISM WITH RELATIVELY STABLE ATTITUDES

Our results show that people claiming benefits during COVID-19 are seen as considerably more deserving than pre-pandemic claimants. However, our analysis of trends in attitudes shows little change in overall welfare generosity over the course of the pandemic: the proportion of the British public who think that benefit levels are too high has remained relatively unchanged by the COVID-19 crisis. This represents something of a puzzle: why did a large influx of ‘more deserving’ claimants fail to produce a corresponding increase in welfare generosity? Several strands of evidence help us to understand this puzzle.

First, general welfare attitudes are much more strongly related to perceptions of pre-pandemic claimants than to perceptions of COVID-19 claimants. Using data from our novel YouGov survey, we used logistic regression to jointly predict whether respondents thought that benefit levels were too high from binary indicators of whether they saw half or more of (i) pre-pandemic claimants, and (ii) COVID-19 claimants as being to blame for losing their jobs or being unable to find a job. Respondents who saw pre-pandemic claimants as highly blameworthy (vs. those that saw them as less blameworthy) were 28.3 percentage points more likely to think that benefits were too high [95% CI: 23.6 to 33.1pp]. By contrast, the difference based on perceptions of the blameworthiness of COVID-19 claimants was only 15.0pp [95% CI: 8.7 to 21.2pp].¹³ This almost two-fold difference suggest that opinions on general benefit levels are to some extent insulated from perceptions of COVID-19 claimants, and are instead being strongly determined by existing perceptions of pre-pandemic claimants.

A second strand of evidence supporting this conclusion is that support for more generous benefits is substantially stronger when an increase in generosity is tied specifically to the COVID-19 pandemic. In the UK (as elsewhere; Moreira and Hick, 2021), the main benefit (UC) was made more generous at the start of COVID-19, initially for a year, and then (after sustained pressure) for a further six months.¹⁴ To examine attitudes to the UC uplift, we adapted two Health Foundation/Ipsos MORI questions:

In April 2020 the Government announced a £20 a week increase in Universal Credit during the pandemic. For a single, childless Universal Credit claimant (aged 25 or over), this means the standard allowance before housing costs has increased from £73 to £94 per week.¹⁵ To what extent do you support or oppose this increase during the pandemic?

¹³ These are average marginal effects

¹⁴ At the time of writing this is shortly due to end, and most non-UC claimants saw no increases

¹⁵ A split-half sample saw an alternative 2nd sentence (which matched the original wording, but is less clear), “For a single Universal Credit claimant (aged 25 or over), this means the standard allowance has increased from around £318 to £410 per month (i.e. from £73 to £94 per week).” There were no meaningful differences in responses between the two version of this question.

The increase in Universal Credit is temporary and is due to end in September this year (2021). To what extent would you support or oppose making this £20 increase permanent beyond September 2021?

Mirroring earlier findings,¹⁶ responses to our June 2021 survey show that the increase in UC was overwhelmingly popular – 63.3% [95% CI: 61.1% to 64.9%] of respondents supported or strongly supported the £20 uplift during the pandemic. This is not simply support for a temporary increase in generosity: in the follow-up question, half of all respondents (50.0%; 95% CI: 47.4% to 51.2%) supported making this increase permanent.¹⁷ Both of these figures are substantially higher than the 31.9% who said that benefits are too low in the YouGov Tracker reported above.

We also asked respondents directly whether they thought that COVID-19 claimants should be treated more generously than pre-pandemic claimants, using the following question:

Do you think people who claimed benefits during COVID-19 (i.e. since Mar 2020) should receive higher payments, lower payments, or the same amount as people who claimed benefits before COVID-19?

Only 8.7% [95% CI: 6.9% to 11.4%] of respondents believed that COVID-19 claimants should receive higher payments – not substantially more than the 5.2% [95% CI: 3.7% to 7.1%] who felt that pre-pandemic claimants should receive more. The vast majority (71.0% [95% CI: 66.8% to 74.1%]) felt that COVID-19 and pre-pandemic claimants should receive the same amount. These results run against the grain of COVID exceptionalism. However, they are consistent with wider evidence that British people support the principle of flat-rate benefits (Sefton, 2005).

Taken together, these results suggest that although general welfare support among the British public during the pandemic is not strongly determined by perceptions of COVID-19 claimants (and there is little support for differential treatment), support for more generous benefits is higher if this is framed as COVID-related.

¹⁶ While we used the same questions as the Health Foundation/Ipsos MORI November 2020 survey, there are some slight methodological differences that make it hard to directly compare them - partly that their results are GB while ours are UK, but primarily that the Ipsos MORI survey was conducted by phone rather than online, and this may well affect people's responses. Bearing this caveat in mind, the Ipsos MORI survey shows slightly greater support for the £20 uplift than our survey (74% support vs. 9% oppose, compared to 63% v. 11% in our survey), and slightly more support for making this permanent (59% support vs. 20% oppose, compared to 50% vs. 23% in our survey). See p37-38 of ipsos.com/sites/default/files/ct/news/documents/2020-12/health-foundation-covid-19-december-2020-slides.pdf

¹⁷ The full responses for our questions are (n=3,429):

- Support for the £20 uplift during the pandemic: 33.4% strongly agree, 30.0% tend to agree, 17.1% neither agree nor disagree, 7.0% tend to disagree, 4.5% strongly disagree, 8.0% don't know
- Support for making this permanent: 28.8% strongly agree, 20.8% tend to agree, 18.1% neither agree nor disagree, 13.5% tend to disagree, 9.8% strongly disagree, 9.0% don't know

4. CONCLUSIONS

Our results suggest that the pandemic has not engendered a meaningful increase in general solidarity with welfare claimants in the UK. Welfare attitudes did appear to soften during the first and second waves of the pandemic – but attitudes rebounded following lockdowns, with only small changes being sustained. These findings are similar to previous evidence that shows limited change in the UK (Curtice et al., 2020; Hicks, 2020) or the Netherlands (Reeskens et al., 2021). However, our findings show this even more convincingly than previous research: the YouGov Trackers used here were conducted using identical methods before and during the pandemic, with repeated measures every eight weeks rather than at only two time points.

These findings may seem surprising, given a number of reasons to expect attitudes to become more pro-welfare. The pandemic was an all-consuming collective crisis, creating solidarity that some argued might lead to a ‘new Beveridge’ moment (e.g. Financial Times, 2020). Experiences of the welfare system grew much more widespread, and both direct and indirect experiences usually make welfare attitudes more generous (Hedegaard, 2014). It might also be expected that COVID-19 claimants would be seen as much more deserving of support (van Oorschot and Roosma, 2017), with many driven out of work in a manner that was transparently outside their control.

One possible explanation for attitude stability is that, contra these expectations, COVID-19 claimants were judged as undeserving, a view taken by some UK commentators.¹⁸ However, using a purpose-designed survey in June 2021, we have shown that this is not the case. Although respondents were generally sympathetic to pre-pandemic claimants, COVID-19 claimants were more likely to be seen as genuinely in need and deserving of help than pre-pandemic claimants, and were considered much less blameworthy (whether for losing their jobs or for being unable to find a job and leave benefits). These effects were large and were consistent across methods (whether we asked people to compare pre-pandemic versus COVID-19 claimants directly; whether evaluating them alongside one another; or when contrasting responses using a randomised vignette design).

Instead, the most plausible explanation is what we term ‘COVID exceptionalism’ (Summers et al., 2021b), where COVID-19 claimants are seen as exceptional, and are mentally bracketed away from other benefit claimants. (This mirrors a phenomenon that has been seen during recessions, in which media coverage contrasts less blameworthy recession-era claims from more blameworthy ‘normal’ claims; see Erler, 2012). When asked to describe the main differences between COVID-19 and pre-pandemic claimants, responses commonly identified COVID-19 claimants as categorically different: as working people who would never have had to claim benefits

¹⁸ [theguardian.com/commentisfree/2021/feb/25/its-heartless-to-blame-job-losses-on-unemployed-people-but-its-how-much-of-britain-thinks](https://www.theguardian.com/commentisfree/2021/feb/25/its-heartless-to-blame-job-losses-on-unemployed-people-but-its-how-much-of-britain-thinks)

if not for the pandemic. We also found that overall welfare generosity was much more strongly predicted by perceptions of pre-pandemic claimants than by perceptions of those claiming during the pandemic.

It would be easy to conclude that despite COVID-19, the public has little appetite for a more generous welfare system – but this would be wrong. Before the pandemic, attitudes had become more pro-welfare than the UK has seen in 20-30 years. Furthermore, we show that if universal increases in welfare generosity are framed as specifically COVID-related (such as the UC uplift), support for increasing benefits in general is higher still. Yet public support is no guarantee of the continuation of COVID-19-era policy changes, and we must remember that there is more to welfare policy than public attitudes. Welfare expansions such as the Beveridge Report in the UK (Hudson et al., 2016) or the New Deal in the US (Newman and Jacobs, 2008) did not arise because of overwhelming public demand. Rather than COVID-19 or other crises leading to an automatic transformation of attitudes, we can instead see them as offering discursive opportunities for a more generous benefit system, which political actors may choose to exploit or to counter, and whose resonance will depend on wider contextual factors. As an expert roundtable in June 2020 noted, “*There is nothing inevitable about how the welfare state will be reshaped post-crisis... Political control matters*” (Burchardt, 2020). The impact of COVID-19 on welfare attitudes and generosity is therefore all to play for.

REFERENCES

- Aarøe L and Petersen MB.** (2014) Crowding out culture: Scandinavians and Americans agree on social welfare in the face of deservingness cues. *The Journal of Politics* **76**: 684–697.
- Ares M, Bürgisser R and Häusermann S.** (2021) Attitudinal polarization towards the redistributive role of the state in the wake of the COVID-19 crisis. *Journal of Elections, Public Opinion and Parties* **31**: 41–55.
- Atkeson LR, Adams AN and Alvarez RM.** (2014) Nonresponse and Mode Effects in Self – and Interviewer-Administered Surveys. *Political Analysis* **22**: 304–320.
- Baumberg B, Bell K and Gaffney D.** (2012) Benefits Stigma in Britain. London: Elizabeth Finn Care/Turn2us. turn2us.org.uk/about_us/media_centre/news_archive/benefits_stigma.aspx
- Béland D, Dinan S, Rocco P, et al.** (2021) Social policy responses to COVID-19 in Canada and the United States: Explaining policy variations between two liberal welfare state regimes. *Social Policy & Administration* **55**: 280–294.
- Bell DNF and Blanchflower DG.** (2011) Young people and the Great Recession. *Oxford Review of Economic Policy* **27**: 241–267.
- Blekesaune M.** (2007) Economic Conditions and Public Attitudes to Welfare Policies. *European Sociological Review* **23**: 393–403.
- Burchardt T.** (2020) Does COVID-19 represent a ‘new Beveridge’ moment, a crisis that will wash away, or a call to action? Report of a roundtable discussion on theories of welfare. *SPDO research note 2*. London: Centre for Analysis of Social Exclusion (CASE), LSE.
- Cotofan M, Dur R and Meier S.** (2021) Does growing up in a recession increase compassion? The case of attitudes towards immigration. *Discussion Paper*. London: Centre for Economic Performance, LSE.
- Curtice J.** (2010) Thermostat or weathervane? Public reactions to spending and redistribution under New Labour. In: Park A, Curtice J, Thomson K, et al. (eds) *British Social Attitudes: the 26th Report*. London: Sage.
- Curtice J, Abrams D and Jessop C.** (2020) Coronavirus and Public Attitudes: Plus ça change? London: NatCen. <https://www.natcen.ac.uk/media/1989002/Coronavirus-and-public-attitudes.pdf>
- Daniele G, Martinangeli AFM, Passarelli F, et al.** (2020) Wind of Change? Experimental Survey Evidence on the COVID-19 Shock and Socio-Political Attitudes in Europe. *Covid Economics* **44**: 72–153.
- Durr RH.** (1993) What Moves Policy Sentiment? *The American Political Science Review* **87**: 158–170.

- Edmiston D.** (2018) The poor “sociological imagination” of the rich: Explaining attitudinal divergence towards welfare, inequality, and redistribution. *Social Policy & Administration* **52**: 983–997.
- Edmiston D, Baumberg Geiger B, de Vries R, et al.** (2020) Who are the new COVID-19 cohort of benefit claimants. Welfare at a (Social) Distance, Rapid Report #2. <http://hub.salford.ac.uk/welfare-at-a-social-distance/wp-content/uploads/sites/120/2020/09/WaSD-Rapid-Report-2-New-COVID-19-claimants.pdf>
- Erler HA.** (2012) A New Face of Poverty? Economic Crises and Poverty Discourses. *Poverty & Public Policy* **4**: 183–204.
- Financial Times.** (2020) Virus lays bare the frailty of the social contract [editorial]. *Financial Times*.
- Geiger B.** (2017) False Beliefs and the Perceived Deservingness of Social Security Benefit Claimants. In: van Oorschot W, Roosma F, Meuleman B, et al. (eds) *The Social Legitimacy of Targeted Welfare: Attitudes to Welfare Deservingness*. Cheltenham, UK: Edward Elgar.
- Geiger BB.** (2018) Benefit ‘Myths’? The Accuracy and Inaccuracy of Public Beliefs about the Benefits System. *Social Policy & Administration* **52**: 998–1018.
- Harell A, Soroka S and Iyengar S.** (2016) Race, prejudice and attitudes toward redistribution: A comparative experimental approach. *European Journal of Political Research* **55**: 723–744.
- Hedegaard TF.** (2014) The Policy Design Effect: Proximity as a Micro-level Explanation of the Effect of Policy Designs on Social Benefit Attitudes. *Scandinavian Political Studies* **37**: 366–384.
- Hicks T.** (2020) Has the UK electorate become more redistributively left-wing since Covid19? *Covid PACER [blog]*. <https://covidpacer.wordpress.com/2020/10/21/has-the-uk-electorate-become-more-redistributively-left-wing-since-covid19/>
- Hoggett P, Wilkinson H and Beedell P.** (2013) Fairness and the Politics of Resentment. *Journal of Social Policy* **42**: 567–585.
- Hudson J, Lunt N, Hamilton C, et al.** (2016) Nostalgia narratives? Pejorative attitudes to welfare in historical perspective: survey evidence from Beveridge to the British Social Attitudes Survey. *Journal of Poverty and Social Justice* **24**: 227–243.
- Hudson N, Grollman C, Kolbas V, et al.** (2020) Public attitudes in the context of COVID-19 and Brexit. In: Curtice J, Hudson N and Montagu I (eds) *British Social Attitudes: The 37th Report*. London: The National Centre for Social Research.
- Jensen C and Petersen MB.** (2017) The Deservingness Heuristic and the Politics of Health Care. *American Journal of Political Science* **61**: 68–83.
- Kenworthy L and Owens LA.** (2011) The surprisingly weak effect of recessions on public opinion. In: Grusky DB, Western B and Wimer C (eds) *The great recession*. Russell Sage Foundation, 196–219.

- Kootstra A.** (2016) Deserving and Undeserving Welfare Claimants in Britain and the Netherlands: Examining the Role of Ethnicity and Migration Status Using a Vignette Experiment. *European Sociological Review* **32**: 325–338.
- Laenen T and van Oorschot W.** (2020) Change or Continuity in Europeans' Welfare Attitudes? In: Laenen T, Meuleman B and van Oorschot W (eds) *Welfare State Legitimacy in Times of Crisis and Austerity: Between Continuity and Change*. Edward Elgar, 249–266.
- Larsen CA and Dejgaard TE.** (2013) The institutional logic of images of the poor and welfare recipients: A comparative study of British, Swedish and Danish newspapers. *Journal of European Social Policy* **23**: 287–299.
- More in Common.** (2020a) The Impact of COVID-19 on British Society. moreincommon.com/media/ucvfntgq/more-in-common-the-new-normal-united-kingdom.pdf
- More in Common.** (2020b) The New Normal? A 7-country comparative study on the impacts of COVID-19 on trust, social cohesion, democracy and expectations for an uncertain future. moreincommon.com/media/bfwlsrxu/more-in-common-the-new-normal-comparative-7-country-en.pdf
- Moreira A and Hick R.** (2021) COVID-19, the Great Recession and social policy: Is this time different? *Social Policy & Administration* **55**: 261-279.
- Mutz DC.** (2011) *Population-based survey experiments*, Princeton, New Jersey: Princeton University Press.
- National Audit Office.** (2016) Benefit sanctions: detailed methodology. *HC 628*. London: National Audit Office. nao.org.uk/report/benefit-sanctions/
- Newman KS and Jacobs E.** (2008) Brothers' Keepers?: The limits of social solidarity in the New Deal Era. In: Rosenthal H and Rothman D (eds) *What do we owe each other?: rights and obligations in contemporary American society*. New Brunswick, New Jersey: Transaction Publishers.
- Parsons S and Wiggins RD.** (2020) Trust in government and others during the COVID-19 pandemic: Initial findings from the COVID-19 Survey in Five National Longitudinal Studies. London: UCL Centre for Longitudinal Studies. cls.ucl.ac.uk/wp-content/uploads/2020/10/Trust-in-government-and-others-during-the-COVID-19-pandemic-%E2%80%93-initial-findings-from-COVID-19-survey.pdf
- Pierson P.** (1996) The New Politics of the Welfare State. *World Politics* **48**: 143–179.
- Reeskens T, Muis Q, Sieben I, et al.** (2021) Stability or change of public opinion and values during the coronavirus crisis? Exploring Dutch longitudinal panel data. *European Societies* **23**: S153–S171.
- Rehm P, Hacker JS and Schlesinger M.** (2012) Insecure Alliances: Risk, Inequality, and Support for the Welfare State. *American Political Science Review* **106**: 386–406.
- Roberts, M. E., Stewart, B. M., Tingley, D., et al.** (2014) Structural topic models for open-ended survey responses. *American Journal of Political Science*, **58**(4): 1064–1082.

- Roosma F.** (2020) The Perceived Abuse of Welfare Benefits in Times of Crisis: Change or Stability in the Achilles' Heel of Welfare State Legitimacy? In: Laenen T, Meuleman B and van Oorschot W (eds) *Welfare State Legitimacy in Times of Crisis and Austerity: Between Continuity and Change*. Edward Elgar.
- Sefton T.** (2005) Give and Take: Attitudes to Redistribution. In: Park A, Curtice J, Thomson K, et al. (eds) *British Social Attitudes: The 22nd Report*. 2005/2006 ed.: SAGE Publications Ltd.
- Sturgis P, Baker N, Callegaro M, et al.** (2016) Report of the Inquiry into the 2015 British general election opinion polls. London: NCRM, British Polling Council, & Market Research Society. eprints.ncrm.ac.uk/3789/1/Report_final_revised.pdf [accessed 13/8/2018]
- Summers K, Lisa Scullion, Geiger BB, et al.** (2021a) Claimants' experiences of the social security system during the first wave of COVID-19. Welfare at a (Social) Distance. distantwelfare.co.uk/publications
- Summers K, Robertshaw D, Gibbons A, et al.** (2021b) Covid Exceptionalism? Attitudes towards working-age benefits during the COVID-19 pandemic. Social Policy Association Conference 2021, 7-9th July 2021. Social Policy Association Conference 2021, 7-9th July 2021. distantwelfare.co.uk/publications
- Taylor-Gooby P.** (2013) Why Do People Stigmatise the Poor at a Time of Rapidly Increasing Inequality, and What Can Be Done About It? *The Political Quarterly* **84**: 31-42.
- Uunk W and van Oorschot W.** (2019) Going with the Flow? The Effect of Economic Fluctuation on People's Solidarity with Unemployed People. *Social Indicators Research* **143**: 1129-1146.
- van Oorschot W and Roosma F.** (2017) The Social Legitimacy of Targeted Welfare and Welfare Deservingness. In: Van Oorschot W, Roosma F, Meuleman B, et al. (eds) *The Social Legitimacy of Targeted Welfare: Attitudes to Welfare Deservingness*. Cheltenham, UK: Edward Elgar, 3-36.
- Wlezien C.** (1995) The Public as Thermostat: Dynamics of Preferences for Spending. *American Journal of Political Science* **39**: 981-1000.

APPENDICES

APPENDIX A: THE YUGOV PANEL

The YouGov panel represents a diverse group of people recruited from a variety of sources, at the current time numbering 400,000 (UK) active users – about 0.7% of the total UK population. To generate approximately representative samples from this non-probability panel, YouGov offer incentives to a sub-sample of the panel to take part, who are designed to be representative of the national adult population (which they term ‘active sampling’). Non-response weights are also calculated to ensure that the final sample match these known population totals.¹⁹ YouGov provides more general descriptions of its panels here and via the YouGov ESOMAR statement [accessed 11/8/2020].

It is not possible to provide a conventional response rate (as a proportion of the YouGov panel members invited to participate), because participants are allocated to surveys at the point they log in to the YouGov site, rather than at the point that they are invited to participate – something that has been noted by other political scientists using YouGov data (e.g. Kootstra, 2016). For the same reason, however, this non-response is likely to be unrelated to interest in the benefits system – participants will not be aware of the topic of the survey in question, which avoids a major contributor to non-response bias (Groves et al., 2006). Across different surveys, about 1 in 5 of those invited to participate will ultimately do so, on average 19 hours after receiving the invitation email.

Note that there are two limitations to the representativeness of our weighted YouGov surveys:

- While being broadly representative of the population, the YouGov panel inevitably under-represents those with weaker written English language skills (and therefore under-represents first-generation migrants) and who struggle to access the internet via a computer/smartphone.
- Weighting ensures representative results where the weighting variables fully capture those factors that influence both participation in the survey and the phenomenon under investigation. To the extent that they fail to do this, biases can result.

It is worth noting that even ‘gold standard’ social research surveys – those using random samples of the population, with high response rates – must contend with threats to representativeness, as non-respondents may differ from respondents. Overall, our judgement is that for most purposes, weighted YouGov data can be treated as broadly representative of the population – but there are particular issues in capturing some groups of unsuccessful claimants, as mentioned in the main report.

¹⁹ Normal YouGov weights are based on age, gender, social class, region and level of education. For political work they also weight by how respondents voted at the previous election, how respondents voted at the EU referendum and their level of political interest. The known totals are taken from large random surveys (the Labour Force Survey, the National Readership Survey and the British Election Study) and administrative data (the Census, official ONS population estimates, electoral results).

APPENDIX B: THE FULL SET OF YUGOV WELFARE POLICY TRACKERS

Trends and breakdowns for the Welfare Trackers are freely available on the YouGov website at yougov.co.uk/topics/politics/explore/academic_area/Welfare_Policy

The YouGov welfare tracker series also includes the following (measures not relevant here are in lighter text):

Thinking about the current benefits system, which of the following best reflects your view?

- » The benefit system works well as it is and doesn't need any reforms
- » The benefit system works fairly well, and doesn't need any major reform
- » The benefit system works fairly badly, and needs some significant reforms
- » The benefit system works very badly, and needs major reforms
- » Don't know

In determining who is eligible for benefits, do you think it is too strict, not strict enough or gets the balance about right?

- » It is not strict enough and is too open to abuse and fraud
- » It is too strict and prevents some people who genuinely need help from receiving it
- » Gets the balance about right
- » Don't know

Thinking about the different groups who the benefit system supports, for each of the following groups do you think it offers too much support, too little support, or about the right amount?

People over the age of 65

People who are currently out of work

People who have a disability

People who are in low-paid work

People on low incomes who are bringing up children

- » The benefit system offers too much support for this group
- » The benefit system offers too little support for this group
- » The benefit system offers about the right amount of support for this group
- » Don't know

If the government was able to increase the amount of money it spends, which of the following areas do you think the government should INCREASE spending in the most? Please tick up to three

- » Education and schools
- » Defence
- » Pensions
- » Welfare benefits
- » Crime and policing
- » Transport
- » Environment and climate change
- » NHS
- » Overseas aid
- » Local government
- » None of these
- » Don't know

And if the government were to cut the amount of money it currently spends, which of the following areas do you think the government should REDUCE spending in the most? Please tick up to three

(Repeats the same list as above)

Mirroring **Table 1** in the main text, we examined trends in each of these measures; the results are shown in **Table A5** below.

Table A5: Change in anti-welfare attitudes between time periods

(Percentage point change with 95% Confidence Intervals)

	Pre-pandemic level	Pre-pandemic to First wave	First wave to Summer 2020	Summer 2020 to 2nd wave	2nd wave to After 2nd wave	Total changes Pre-pandemic to May 2021
Half or fewer claimants are deserving	34.2% [32.9 to 35.5]	-3.2 [-5.4 to -0.9]	+4.2 [1.5 to 6.9]	-4.4 [-6.9 to -1.9]	+0.2 [-2.7 to 3.3]	-3.1 [-6.0 to -0.2]
Benefit levels are too high	17.8% [16.8 to 18.8]	-3.1 [-4.8 to -1.3]	+2.4 [0.4 to 4.5]	-1.9 [-3.8 to 0.05]	-0.3 [-2.5 to 2.0]	-2.8 [-5.0 to -0.6]
Benefit conditions are not strict enough	39.1% [59.6 to 62.1]	-2.9 [-5.2 to -0.6]	+6.9 [4.2 to 9.7]	-4.2 [-6.7 to -1.7]	+2.1 [-1.1 to 5.2]	+1.9 [-1.2 to 4.9]
People out of work get too much support	18.4% [17.4 to 19.5]	-4.7 [-6.5 to -3.0]	+0.3 [-1.8 to 2.3]	-2.3 [-4.1 to -0.5]	+0.3 [-1.9 to 2.4]	-6.5 [-8.6 to -4.4]
People in low paid work get too much support	8.4% [7.7 to 9.2]	-1.5 [-2.9 to -0.1]	+0.2 [-1.3 to 1.7]	-0.3 [-1.6 to 1.0]	+0.1 [-0.1 to 2.5]	-0.1 [-2.5 to 1.0]
Parents on low incomes get too much support	13.4% [12.5 to 14.3]	-1.9 [-3.6 to -0.2]	+1.5 [-0.5 to 3.5]	-0.1 [-2.0 to 1.7]	-0.2 [-2.5 to 2.0]	-1.0 [-2.9 to 1.4]
Disabled people get too much support	5.8% [5.2 to 6.5]	-0.1 [-1.8 to 0.3]	+1.0 [-0.4 to 2.4]	+1.0 [-0.4 to 2.4]	-1.3 [-2.9 to 0.3]	-0.0 [-1.5 to 1.4]
Priority is to reduce benefits spending	21.4% [20.4 to 22.5]	-2.5 [-4.3 to -1.0]	+1.4 [-1.0 to 3.5]	-2.5 [-4.7 to -0.3]	+2.6 [-0.3 to 5.4]	-1.0 [-3.6 to 1.6]
PRO-welfare measures						
Priority is to increase benefits spending	13.4% [12.5 to 14.3]	+1.3 [-0.3 to 3.0]	+0.0 [-1.8 to 2.0]	+2.8 [1.0 to 4.8]	-1.0 [-3.6 to 1.6]	+3.2 [1.3 to 5.4]
The benefits system works well or fairly well as it is	16.9% [15.9 to 17.9]	3.1 [1.3 to 5.0]	-0.3 [-2.5 to 2.0]	+2.1 [0.1 to 4.2]	+0.1 [-2.4 to 2.7]	+5.1 [2.7 to 7.6]

Source: YouGov Welfare Tracker surveys (total N=21,602). **Note:** The question on the spending priorities was asked in different survey waves than the other measures reported in this table. One of these waves fell in December 2020, and due to the potential 'Christmas effect' data from this wave was not used when calculating the 'Second wave' figures in the table.

Table A6: Change in anti-welfare attitudes between time periods for Labour and Conservative voters

(Percentage point change with 95% Confidence Intervals)

	Half or fewer are deserving		Benefit levels are too high	
Pre-pandemic level				
Conservative	48.0%	[45.6 to 50.0]	30.1%	[28.1 to 32.1]
Labour	19.7%	[17.7 to 22.0]	7.7%	[6.4 to 9.2]
Changes during COVID				
Pre-pandemic to First wave				
Conservative	-6.7	[-10.5 to -2.9]	-6.2	[-9.5 to -2.8]
Labour	-1.6	[-5.2 to 1.9]	+0.1	[-2.5 to 2.6]
Difference	-5.0	[-10.3 to 0.2]	-6.3	[-10.5 to -2.0]
First wave to Summer 2020				
Conservative	+4.3	[-0.1 to 8.7]	+2.6	[-1.3 to 6.4]
Labour	+3.1	[-1.1 to 7.4]	+0.4	[-2.5 to 3.4]
Difference	+1.2	[-4.9 to 7.3]	+2.1	[-2.7 to 7.0]
Summer 2020 to Second wave				
Conservative	-1.1	[-5.1 to 2.9]	-0.5	[-4.0 to 3.1]
Labour	-4.8	[-8.9 to -0.1]	-2.1	[-4.7 to 0.5]
Difference	+3.7	[-2.0 to 9.4]	+1.7	[-2.8 to 6.1]
Second wave to After second wave				
Conservative	-1.4	[-6.4 to 3.6]	-0.1	[-5.2 to 3.6]
Labour	+1.1	[-3.6 to 5.8]	+1.1	[-2.2 to 4.5]
Difference	-2.5	[-9.3 to 4.3]	-2.0	[-7.5 to 3.6]
Total changes to May 2021				
Conservative	-4.9	[-9.7 to -0.1]	-4.9	[-9.2 to -0.1]
Labour	-2.2	[-6.7 to 2.2]	-0.4	[-3.7 to 2.7]
Difference	-2.7	[-9.2 to 3.9]	-4.4	[-9.8 to 0.1]

Source: YouGov Welfare Tracker surveys (total N=13,118)

APPENDIX C: STRUCTURAL TOPIC MODEL OF OPEN TEXT RESPONSES

A Structural Topic Model (STM) builds off of the broader tradition of probabilistic topic models. Topic models use word counts to establish topics. Each word used in a given document (open text responses to a survey question in our case) is assumed to have a probability of belonging to a particular topic. The topic model uses an inductive process to attempt to discern the appropriate number of topics to cover the documents in the analytical corpus.

There is no single ‘correct’ number of topics for a given corpus. Rather we based our model selection on a combination of model fit statistics and manual inspection of exemplar response within each topic. We examined models with 3, 5, 7, and 10 topics. Model fit statistics are reported in **Figure C1**, below.

Figure C1: STM diagnostic values by number of topics




Figure C1 plots four values against the specified number of topics:

- The held-out likelihood measures the probability that the words in a given topic would appear in a text allocated to a particular topic. Higher values indicate that the model is doing a better job of predicting what words are going to appear in a particular document.
- Residuals measure the amount of overdispersion in the residuals. In other words, it tells us whether the topics are precise or not. Imprecise topics are too general and do not really capturing the nuance in each document. Low scores imply

less dispersion. Typically more topics will, by construction, lead to lower residuals and so this needs to be offset against the demands of parsimony.

- Semantic coherence refers to the degree to which the most probable words in a topic co-occur. This is useful because it tells us whether the most salient words for a particular topic are correlated with each other. In cases where there is low correlation between highly probable words for a given topic indicates that there is something wrong with the topic – this is because the words appear to bifurcated – with some words frequently appearing in a topic but not in the same document as other words that also frequently appear in the same topic. Higher semantic coherence is better but it is also relatively easy to get high semantic coherence with a small number of topics.
- The lower bound gives us an indication of the change in fit between iterations of the model. Smaller improvements between models, that is as this line starts to flatten, implies model convergence.

According to these measures, the three-topic model appears to be the poorest fit to the data, and was therefore discounted. The five-topic model appears to fit the data best by a number of metrics. However, differences between this model and the seven and ten-topic models are very small indeed, and manual inspection of topic exemplars for the seven-topic model indicate a clearer semantic differentiation between topics than for the five-topic model. The 10 topic has a slightly higher held-out likelihood than the seven-topic model. However, it has lower semantic coherence, and only very slightly lower residuals, despite three additional topics. Therefore, on the basis of parsimony, we selected the seven-topic model.



University of
Salford
MANCHESTER

University of
Kent


UNIVERSITY OF LEEDS



**WELFARE — AT A
SOCIAL DISTANCE**

ISBN: 978-1-912337-52-1