

LETTER

Risk and Preferences for Government Healthcare Spending: Evidence from the UK COVID-19 Crisis

Jack Blumenau¹ , Timothy Hicks^{1*}  and Raluca L. Pahontu² 

¹School of Public Policy, University College London, London, UK and ²Department of Political Economy, King's College London, London, UK

*Corresponding author. Email: t.hicks@ucl.ac.uk

Authors listed in alphabetical order to indicate equal contributions from each.

(Received 7 December 2021; revised 5 July 2022; accepted 14 September 2022)

Abstract

The onset of the COVID-19 pandemic constituted a large shock to the risk of acquiring a disease that represents a meaningful threat to health. We investigate whether individuals subject to larger increases in objective health risk – operationalized by occupation-based measures of proximity to other people – became more supportive of increased government healthcare spending during the crisis. Using panel data that track UK individuals before (May 2018–December 2019) and after (June 2020) the outbreak of the pandemic, we implement a fixed-effect design that was pre-registered before the key treatment variable was available to us. While individuals in high-risk occupations were more worried about their personal risk of infection and had higher COVID-19 death rates, there is no evidence that increased health risks during COVID-19 shifted either attitudes on government spending on healthcare or broader attitudes relating to redistribution. Our findings are consistent with recent research demonstrating the limited effects of the pandemic on political attitudes.

Keywords: risk; insurance; healthcare spending; difference-in-differences

The onset of the COVID-19 pandemic led to the biggest public health emergency in the Western world in over a century. While its consequences were extremely wide-ranging, one feature of the pandemic that is particularly theoretically interesting is that it constituted a large and extremely salient shock to both real and perceived health risks for individuals. From around late January 2020, a series of fairly unprecedented ‘lockdowns’ unfolded around the world, leaving hundreds of millions of people sitting in their homes, fearful of a very poorly understood and apparently lethal virus. News media was full of coverage of healthcare settings overwhelmed with extremely sick patients, of ambulances unable to reach people in time, and of morgues running out of space for bodies. What, if any, were the consequences of these events for mass preferences regarding government funding of health systems?¹

We use data from a large and rich panel survey in the UK to answer this question. In doing so, we add to a rapidly growing literature that has sought to understand how the pandemic has affected mass political attitudes across a range of ‘developed’ democracies, including core political attitudes (see, for example, Ares, Bürgisser and Häusermann 2021; Blumenau et al. 2021) and trust in governing elites (see, for example, Bol et al. 2021; Esaisson, Sohlberg and Ghersetti 2021). As such, we also contribute to a broader literature that has sought to understand the

¹Our theoretical focus is on health systems that are predominantly government funded, so that any increases in health expenditure must predominantly run through government action.

© The Author(s), 2022. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

consequences of public health emergencies, more generally, on public policies and mass politics. For example, scholarship on the longer-term effects of the Black Death and 1918 influenza pandemic has shown lasting impacts on economic inequality, public spending and voting behaviour (see, for example, Gingerich and Vogler 2021; Grantz et al. 2016).

Our conceptualization of the COVID-19 pandemic as entailing a massive shock to health risks additionally allows us to contribute to an influential literature in political economy that has demonstrated that risk can be an important determinant of mass attitudes regarding welfare policy. For example, attempts to insure against the risks associated with highly specific labour market skills have been shown to drive broad welfare state attitudes (Iversen and Soskice 2001), and a higher likelihood of lost employment has been shown to be associated with support for more generous unemployment benefits as a way of insuring against those risks (Rehm 2009). Indeed, shocks to individuals' health and human capital have been shown to increase support for the welfare state and the Left (Pahontu, Hooijer and Rueda 2020).

We contribute to this literature by exploring the effects of exposure to health risks, prompted by the pandemic, on individuals' support for government healthcare spending. We study the UK, one of the countries hardest hit by the pandemic, using a generalized difference-in-difference (DiD) design in which we follow a panel of survey respondents in the period before (May 2018–December 2019) and during (June 2020) the COVID-19 pandemic. Our inferences are based on exploiting the relatively sudden onset of the health crisis, together with a novel measure of the health risk that individuals faced as a consequence of the pandemic based on the physical proximity between people across occupations.

We find no evidence that individuals more exposed to this health risk changed their preferences for government healthcare spending (more than low-risk individuals), nor do we find any heterogeneous treatment effects by their ability to work from home. We do, however, find evidence that those who had higher *objective* health risks as a consequence of the pandemic also reported being more worried about catching COVID-19 and expressed higher levels of support for restrictions aimed at limiting the spread of COVID-19 infections. This provides reassuring evidence that our null finding on preferences over healthcare spending is unlikely to have been driven by people failing to notice the changing health risks they faced. The pandemic substantially increased the health risks faced by many UK citizens, but we find that these risks did not lead to greater support for increased government spending on healthcare.

From Health Risk to Preferences over Government Healthcare Spending

Why might we expect that heightened health risk should be met with increased demand for government healthcare spending? While there are many possible mechanisms that could connect features of the COVID-19 pandemic to attitudes over healthcare spending,² our theoretical interest in this article is in assessing the extent to which health risks are drivers of such preferences *via an insurance logic*. The vast majority of UK healthcare expenditure (78 per cent as of 2018) is channelled through the public sector (Office for National Statistics 2020a, Section 5), and only around 11 per cent of residents have any kind of private health insurance (The King's Fund 2014, 4). As such, for the majority of people in the UK, the primary mechanism through which healthcare expenditure can be increased is via increases in government expenditure on the health system.

For the risk–insurance logic to operate, it is necessary that the increased expenditure – financed by increased taxes that beneficiaries of the National Health Service (NHS) would have to expect to pay – should actually purchase some kind of insurance for those increased health risks. Our main argument in this regard relates to concerns about health system capacity. News media reporting on the unfolding pandemic (in the UK and beyond) provided heavy emphasis on the extent to which the NHS was struggling to cope with the sheer volume of

²For example, increased salience of the operation of the health system, increased perceptions of need for the operation of the health system and sympathy and gratitude towards health system workers.

patients that it needed to treat: wards were full; private hospitals were being (temporarily) de facto nationalized; early on, there was a highly salient national effort to design and source ventilators; and staffing shortages were feared given the viral threat to healthcare workers. As a consequence, a reasonable inference for British residents to make was that the NHS urgently needed a lot more resources in order to cope with these demands. Increasing public healthcare expenditure provided an obvious way to ensure the (future) availability of those additional resources. Moreover, the risk–insurance logic should operate fairly clearly here. People who faced greater risk of COVID-19 infection were also those for whom the capacity constraints of the healthcare system were likely to cause most concern, as they increased the risk of leaving them with inadequate care should they have become infected with COVID-19. As such, those at high risk of infection should have had the strongest incentive, other things equal, to want health expenditure to increase.

There is another possible mechanism through which the risk–insurance logic might operate. Increased healthcare expenditure might be expected to be channelled into medical research that could yield improved medical treatments or preventions (for example, vaccines) for the increased health risks *directly* caused by the SARS-CoV-2 virus.³ It is an open question – and one that is unanswerable with our data – as to the extent to which respondents in our British panel survey had expectations that supporting a fairly generic increase in health spending would have meant increased funding in those particular categories. On the one hand, the vaccine research effort was plausibly seen as distinct from the provision of healthcare (and so its associated funding streams). On the other hand, treatment-oriented research was, at points, being reported as being carried out within NHS hospitals, suggesting that health expenditure might also be helpful there.

Our goal in this article is not to distinguish between these different mechanisms. Rather, it is to argue that there *are* reasons to think that the risk–insurance logic, generally conceived, may operate in this context and then to test that proposition empirically.

Research Design

Our research design combines individual-level data on attitudes towards taxation and spending on healthcare from the British Election Study (BES) panel survey with data on the objective health risks faced by individuals during the pandemic. We use four waves of the BES fielded between May 2018 and December 2019, as well as one wave fielded during the pandemic in June 2020. Our inferences are primarily informed by the respondents who appear in more than one wave of the BES data, one of which is the pandemic wave – which is to say, 8,025 individuals.⁴ We operationalize health risks using measures of the physical proximity between people in different occupations. As we describe in more detail in the following, this key treatment variable is matched to respondents in the BES on the basis of Standard Occupational Category (SOC) codes. All our analyses were pre-registered with the Open Science Framework in January 2021, before the BES release and therefore before we had access to our central treatment variable (Blumenau, Hicks and Pahontu 2021).

Measuring Public Health Expenditure Preferences

Our main outcome variable comes from a question that asks respondents to place themselves on an eleven-point scale, where the minimum value corresponds to the statement ‘Government should cut taxes a lot and spend much less on health and social services’ and the maximum corresponds to ‘Government should increase taxes a lot and spend much more on health and social services’. Following BES convention, we label this variable $taxSpendSelf_{i,t}$, where i indexes individuals and t indexes survey waves. We assign integer values to the response categories, dropping

³We note that the period for which we have panel survey data is entirely before the availability of approved COVID-19 vaccines.

⁴For further details of the sample that are relevant to our design, see Section A of the Online Appendix.

‘Don’t know’ responses, and then scale this variable to have a mean of 0 and a standard deviation of 1. Our estimated effects can therefore be interpreted in standard deviations of the outcome, where higher values constitute more ‘left-wing’ positions.

Using this BES question as our dependent variable may pose potential problems. In particular, the question wording includes reference to a spending trade-off with taxation. Also, the question is not limited to asking about healthcare spending, but also mentions ‘social services’. We have reasons, however, not to be too concerned.

First, if anything, we expect that after the pandemic began, the ‘tax’ and ‘social services’ components of the question will have *decreased* in salience relative to the ‘health’ component among respondents. Consequently, preferences regarding health spending should have become an even stronger driver of responses to our survey question during the pandemic. Secondly, the clearest inferential threat posed by this dependent variable is that our treatment (health risks) might have differential effects for the ‘taxation’ and ‘healthcare’ components of our outcome. However, given our research design, for the (null) findings we present to be explained by such offsetting effects, it would have to be the case that people who work in close proximity to others (our treated group) became more oppositional to tax increases (or, indeed, to increasing spending on ‘social services’) during the pandemic than people who work further from others. We see no reason to believe that this would be the case. Taken together, then, although this variable is imperfect, we nevertheless expect that it is likely to capture the most salient effects of health risk on preferences for government healthcare spending during the pandemic.

Measuring Health Risk

The sudden onset of the pandemic increased the health risk to most, if not all, parts of the population. However, the increased risk of illness was not equally distributed. How can we measure which people saw larger or smaller increases in health risk as a result of the pandemic? The central methodological assumption we make is that people who work in occupations that involve closer physical proximity to others were at higher risk from COVID-19 than people from occupations that involve more socially distant interactions.

To operationalize this intuition, we rely on data from the Office of National Statistics (ONS) measuring the average physical proximity to others (*OccProximityRisk_i*) for different occupational categories. Our measure was originally developed by the US Department of Labour’s Occupational Information Network (O*NET) and is based on survey respondents’ answers to the question, ‘How physically close to other people are you when you perform your current job?’. Responses, which were given on a five-point scale that varies from 0 (beyond 100 feet from another human) to 100 (very close contact to others, nearly touching), were then aggregated by O*NET to the occupation level on the basis of US Standard Occupational Classification (SOC) codes. The ONS subsequently mapped the occupational averages to UK SOC codes, which we use to link the measure to respondents in the BES (Office for National Statistics 2020b). In the following, we provide a series of validation checks for our measure of occupational health risk.⁵

Finally, we also require information on whether respondents worked from home during the pandemic, as this may be a significant factor in determining workers’ exposure to COVID-19-related health risks. We take this information from Wave 20 of the BES, which asked respondents: ‘Have you started working from home as a result of the coronavirus

⁵We note that average occupational physical proximity is likely to be very similar in the US and UK. First, the labour markets in the two countries are very similar on a broad range of indicators, including unemployment trends, average hours worked, ease of finding a job and wage growth (Forbes 2016). Secondly, and most crucially, the two countries share a great similarity in terms of occupational tasks, presumably a key determinant of whether an occupation requires people to work in close proximity to others or not. For example, Goos and Manning (2007) note that the task composition of occupations in the UK is expected to be similar to that in the US.

Table 1. The table shows the occupations with the highest and lowest proximity-based risk scores from the ONS data (*OccProximityRisk*)

Occupation	(<i>OccProximityRisk</i>)
Dental nurses	99.5
Dental practitioners	97
Midwives	97
Paramedics	97
Ambulance staff (excluding paramedics)	97
Physiotherapists	96.5
Actors, entertainers and presenters	95
Veterinarians	91
Ophthalmic opticians	90
Veterinary nurses	90
...	...
Physical scientists	39.3
Actuaries, economists and statisticians	39
Launderers, dry cleaners and pressers	35
Legal professionals n.e.c.	34
Solicitors	34
Barristers and judges	34
Marketing associate professionals	33
Advertising accounts managers and creative directors	33
Agricultural machinery drivers	26.5
Artists	21.5

outbreak?'. We define $workHome_i = 1$ for responses that are 'Yes' or 'I already regularly worked at home', and 0 otherwise. We again drop all 'Don't know' responses from the analyses.⁶

Occupational Health Risk Validation

Table 1 shows the occupations that have the highest and lowest proximity-based risk in our data. High-risk occupations include several types of medical professionals, as well as entertainers and veterinarians, all of whom clearly work in close proximity to others in the course of their work. By contrast, occupations with low proximity-based risk include legal professionals, scientists, actuaries, farmers and artists. This suggests plausible face validity for our measure, as these occupations conform with intuitive notions of which types of people were likely to be at risk of infection from COVID-19.

In Figure 1, we show that the *OccProximityRisk* measure correlates with health experiences of, and attitudes towards, COVID-19.⁷ The left-hand panel of Figure 1 shows, for each occupation, the proximity-based risk (*OccProximityRisk* [x-axis]) and COVID-19 death rates for the period between March and December 2020 (y-axis). Point sizes are proportional to the total number of deaths recorded by the ONS for each occupation. Occupational proximity-based risk is positively associated with occupational death rates from COVID-19. The regression line, weighted by the total number of deaths for each occupation, is positive and significantly different from 0 ($t = 5.09$). The magnitude of this relationship is non-trivial: occupations at the 90th percentile of our *OccProximityRisk* variable have a death rate that is 3.7 percentage points higher than that of occupations at the 10th percentile, which is an increase of 34 per cent.⁸

⁶We provide further detail on our coding choices in the Online Appendix Section A.

⁷In the case of labour market risks, subjective measures may better predict preferences than objective ones (Melcher 2021), and this might be because individuals are not aware of their objective unemployment risks (such as occupational unemployment rates). This is unlikely to be a limitation for the use of *OccProximityRisk*, as this measure captures an occupation's proximity to others, which is a fact known to any individual (unlike, for example, occupational unemployment rates).

⁸We calculate the death rate for each occupation as the ratio of the number of deaths from COVID-19 between March and December 2020 to the average number of deaths from all causes for the period 2015–19. We use the average number of total

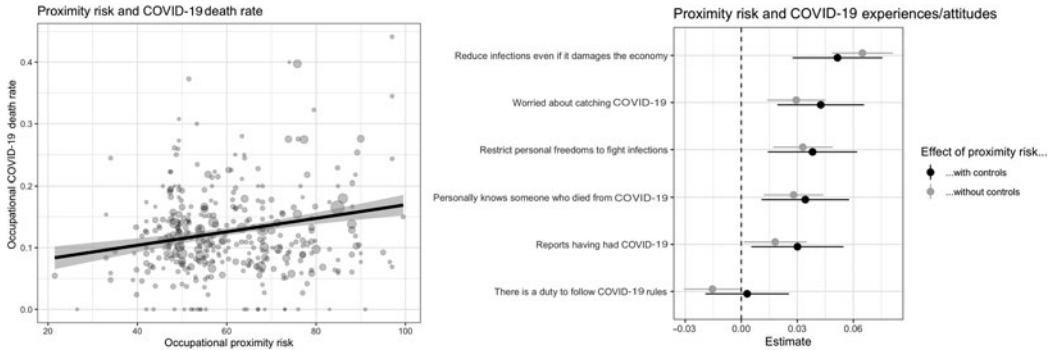


Fig. 1. Left-panel: association between occupational proximity-risk (x-axis) and occupational COVID-19 death rate (y-axis). Right-panel: linear association between occupational proximity-based risk and self-reported COVID-19 experiences and attitudes.

Note: Black points represent estimates from a model that controls for income, education, housing status, age and region; grey points represent estimates from bivariate regressions.

The right-hand panel shows estimates from a series of linear regression models, in which we regress BES variables that measure respondents’ attitudes to, and experiences of, the COVID-19 pandemic on the occupational proximity-based risk of respondents. In these regressions, we standardize all outcomes, as well as the occupational proximity-based risk variable, to have a mean of 0 and a standard deviation of 1. Grey points represent estimates from bivariate regressions, while black points represent estimates from regressions that also control for education, household income, housing status, age and region.⁹

Figure 1 demonstrates that respondents with higher levels of proximity-based risk express more concern about their own health risks and are more supportive of government actions to reduce the spread of COVID-19 infections than are respondents with lower levels of proximity-based risk. Importantly, respondents’ *perceived* health risks appear to correlate systematically with their *objective* health risks: BES respondents with higher levels of *OccProximityRisk* report being more worried about catching Covid-19, are more likely to report having had COVID-19 and are more likely to know someone who has died from COVID-19. These respondents are also more supportive of restricting personal freedoms in order to reduce infections and more likely to support measures to reduce infections even if doing so threatens to damage the economy. Of the variables considered here, it is only with respect to the statement that ‘It is every citizen’s duty to follow the coronavirus rules’ that we find insignificant differences between higher-risk and lower-risk respondents. Taken together, then, these results provide reassuring validation for our measurement strategy.

Models

We aim to identify the effect of *OccProximityRisk_i* on healthcare spending preferences by comparing the change in attitudes of those in high-risk occupations to the change in attitudes of those in low-risk occupations before and after the onset of the pandemic. Defining

deaths per occupation over a five-year period in order to reduce the degree to which our results are sensitive to idiosyncrasies in the number of deaths in a given year, which may be small for some occupations. This choice is consistent with previous work (see: <https://www.ons.gov.uk/peoplepopulationandcommunity/birthsdeathsandmarriages/deaths/bulletins/deathsregisteredweeklyinenglandandwalesprovisional/latest>).

⁹We use an eight-category measure of education, a nine-category measure of housing status, a seventeen-category measure of income and a fourteen-category measure of region, and we include a quadratic specification for age.

$OccProximityRisk_{i,t} = OccProximityRisk_i \times Pandemic_t$, where $Pandemic_t$ is a dummy equal to 1 for the survey wave during the COVID-19 pandemic and 0 otherwise, we estimate linear models of the following form:

$$taxSpendSelf_{i,t} = \gamma \cdot OccProximityRisk_{i,t} + \alpha_i + \delta_t + \sum_j^{j=1} \beta_j \cdot X_{i,t}^j + \epsilon_{i,t}. \quad (1)$$

where γ represents the effect of proximity-based risk on healthcare spending attitudes and is our main quantity of interest, while α_i and δ_t capture individual and survey-wave fixed effects. The individual fixed effects imply that the variation in our outcomes used to identify γ comes solely from within-respondent changes in the outcome over time. The key virtue of the fixed-effect design is that it eliminates the possibility of confounding that stems from systematic differences between high- and low-risk individuals that are constant over time. For instance, one might imagine that our measure of occupational proximity-based risk might correlate with respondents' income or social class, and that these features are likely to be systematically associated with preferences over government healthcare spending. If respondents from different positions in the income distribution, or from different social classes, differ in support for healthcare spending, then this could confound any inferences we might make about the effects of proximity-based risk in a cross-sectional analysis. However, the inclusion of unit-level fixed effects rules out confounding of this sort, not only for social-class or income-based differences, but for *any* covariate that does not vary within individuals over time. Similarly, the time fixed effects, δ_t , account for common shocks in each survey wave that contribute to tax/spend attitudes. Finally, we include a set of time-varying control variables, X , which we discuss in the following.

Consistent with our theoretical discussion, we expect that those with higher infection risk became more supportive of government spending on healthcare during the pandemic period relative to those respondents with lower infection risk (that is, $\gamma > 0$). However, the pandemic forced many people to work from home, something that could considerably alter the occupation-based health risks that they were subject to after the onset of the crisis. If homeworkers are no longer proximate to other humans (outside of their household), our estimate of γ in Equation 1 will likely represent an underestimate of health risk on spending preferences. We address this issue by estimating further (pre-registered) models that include an interaction between $OccProximityRisk_{i,t}$ and the $workHome_i$ dummy variable. We expect the effects of occupational risk to be smaller for those working from home.

Finally, the marginal effect of $OccProximityRisk_i$ may be non-linear, making the specification in Equation 1 inappropriate. We therefore report additional analyses, also pre-registered, which adopt the approach proposed by Hainmueller, Mummolo and Xu (2019) and estimate separate treatment effects for three equally sized groups of $OccProximityRisk$ (*Low*, *Middle*, *High*). Again, we interact this binned measure with the $workHome_{it}$ variable described earlier.

Although our fixed-effects approach adjusts for any confounding that is constant within individuals, we might still worry about other changes in other variables that occur within individuals and that correlate with health risk. As the equations discussed earlier indicate, we control for a set of time-varying variables in all specifications ($X_{i,t}$).

First, an influential literature shows that labour market risks are an important determinant of mass attitudes towards government expenditure (see, for example, Rehm 2011) and it is plausible that labour market risks and health risks may be correlated. Following Rehm (2009), we therefore adjust for quarterly, occupation-by-gender unemployment rates from the UK Labour Force Survey that correspond to the field dates of the BES panel waves, which we merge with the survey data.¹⁰

¹⁰We measure occupational unemployment rates at the one-digit SOC code level. In Table C.3 in the Online Appendix, we replicate the analysis while controlling for unemployment rates measured at the three-digit SOC code level. The results are substantively identical.




Fig. 2. Estimated treatment effects from two-way fixed-effect models where the outcome variable is *taxSpendSelf*. Notes: Model 1 presents results from Equation 1, which includes only the continuous proximity-based risk treatment (plus controls for individual-level unemployment and the occupational unemployment rate measured at the one-digit SOC level). Model 2 additionally includes an interaction between proximity-based risk and a dummy for whether a respondent reports working from home during the pandemic. Model 3 interacts the categorical version of the proximity-based risk measure with the work-from-home dummy.

Secondly, we also adjust for *realized* labour market risk by including an individual-level time-varying measure of whether someone is in work or unemployed. Margalit (2013) has shown that this realized labour market risk can have important effects on attitudes towards government expenditure (on unemployment benefits).

Finally, we note one other control variable that we could, but do not, use: individual-level reports of suspected contraction of COVID-19 by the respondents themselves or those close to them. It may appear that this could be an important variable to include in order to assess whether the possible effects of $OccProximityRisk_i$ are driven by risk itself or by the realization of a negative health outcome for an individual. However, while such a variable is available in the BES, our (pre-registered) concern with specifications of this sort is that *realized* COVID-19 infection is necessarily post-treatment to the *risk* of infection. As such, including a variable that measures whether an individual contracted COVID-19 into our regression results will bias the estimates of our risk variable.

Results

We present the results of all three of the models for the *taxSpendSelf* outcome in Figure 2.¹¹ The left-hand panel of Figure 2 shows the estimated treatment effect for the continuous measure of $OccProximityRisk$ from Equation 1. Contrary to expectations, the point estimate is negatively signed, but it is very small in magnitude and statistically indistinguishable from 0. The analysis is also sufficiently well powered that the estimated confidence intervals rule out even small effect sizes. In short, we find no evidence that – across all respondents – increased health risks during COVID-19 led to attitudinal change on government health spending.

In the centre panel of Figure 2, we evaluate whether this precisely estimated null effect masks heterogeneity across respondents who do and do not work from home. This does not appear to be the case: regardless of whether respondents report working from home during the pandemic, higher levels of occupational proximity-based risk do not affect attitudes towards government spending on healthcare.

Finally, the right-hand panel presents results from our categorical risk measure, and we again find no evidence that those in higher-risk occupations became more in favour of increased government spending than those in low-risk occupations during the crisis. In fact, we estimate a small *negative* effect for those in the highest-risk occupations and who worked from home during the pandemic.¹²

¹¹For full results, see Tables C.1 and C.3 in the Online Appendix.

¹²In Section E in the Online Appendix, we move beyond an analysis of average treatment effects to investigate whether the pandemic was associated with any mean-preserving polarization or convergence in preferences towards government spending on healthcare. We find that there is no evidence of such shifts when averaging either across all respondents or for each of the groups defined by our occupational health risk variable.




Fig. 3. Parallel trends for *taxSpendSelf*.

Notes: Points represent the average response to the *taxSpendSelf* variable in each survey wave for respondents in the high, mid and low categories of *OccProximityRisk*. The dashed vertical line indicates the beginning of the first COVID-19 lockdown in the UK.

Two possible concerns regarding our inferences may arise at this point. First, one may object to our design on the basis that it was not just those with higher COVID-19 infection risk that faced increased healthcare risks: *anyone* who expected that they would need to make use of health services during the pandemic period would have had an incentive to see funding increase in order to avoid the capacity problems that were so clear. An obvious group for whom this kind of logic may operate would be those who are older, as they would tend to have more ongoing health issues that require medical attention. Due to us restricting the BES sample to those under 66 years old, we believe this issue is unlikely to be problematic for our inferences. In addition, if there were a much broader rise in concerns about NHS capacity, we should expect a much more widespread rise in support for health spending across the population. As can be seen in Figure 3, there is no evidence of this in our data.

The second possible concern regarding our inferences may arise from our DiD design requiring a common-trends assumption across our treatment and control groups. Figure 3 indicates that there is fairly good evidence of common trends in the pre-treatment survey waves when splitting respondents into the three groups defined by the categorical version of our *OccProximityRisk* variable. There is some evidence that the high- and medium-risk groups saw slightly larger increases in support for *taxSpendSelf* in the BES wave immediately prior to the pandemic, but in general, the trends of support are similar over time. Moreover, in the Online Appendix Section B, we show that a closely related dependent variable for which we have observations from many more waves also exhibits clear parallel trends in the pre-treatment period. In short, the analysis of pre-pandemic trends is reassuring, in that it suggests the behaviours of the low-risk group during the pandemic are likely to provide a suitable counterfactual for the higher-risk groups.

Conclusion

This article studies whether individuals responded to the risk of ill health prompted by the pandemic by increasing support for government healthcare spending. We find evidence that individuals more exposed to the shock were more likely to worry about their risk of ill health and had objectively higher deaths, but we find little evidence that this translated into a change in preferences on spending. This is consistent with recent work showing limited effects of the pandemic on attitudes (Ares, Bürgisser and Häusermann 2021; Blumenau et al. 2021; Lowande and

Rogowski 2021). However, our findings differ from those of studies of the United States, which have shown an increase in support for the expansion of government-provided healthcare coverage and spending among those most exposed to the pandemic (Fox et al. 2022; Rees-Jones et al. 2022). This very likely reflects a key difference between the universal healthcare system of the UK and the employer-sponsored insurance system of the United States, rather than an inherently conflicting set of findings. In the United States, the pandemic's associated recession increased the probability of people losing their jobs and, therefore, their health insurance. As such, the US findings are consistent with people wanting to increase the probability of government-funded health insurance being available to them if they need it, rather than increasing per-person-insured government health spending – which is what the UK data relate to.

Overall, the results in this article provide evidence against the operation of the risk–insurance logic of public spending attitudes – at least in this case. The interesting question is: why? On this, we can only speculate. A plausible explanation – especially given that the pandemic-period data that we use come from rather early in the COVID-19 era (June 2020) – is that respondents may have seen the health emergency as providing little in the way of a guide for how public policies should change. To the extent that the pandemic was seen as extraordinary and expected to be relatively short-lived, then it may have been that people did not change their views about the appropriate levels of expenditure on the health system.

Supplementary Material. Online appendices are available at: <https://doi.org/10.1017/S0007123422000539>

Data Availability Statement. Replication data for this article can be found in Harvard Dataverse at: <https://doi.org/10.7910/DVN/AU6FCP>.

Acknowledgements. We thank the editor, the anonymous reviewers and the audience at King's College London for useful comments and suggestions.

Financial Support. This work was supported by a Leverhulme Trust Research Fellowship (RF-2021-327, Blumenau).

Competing Interests. None.

References

- 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(sup 1), 41–55.
- Blumenau J, Hicks T and Pahontu RL (2021) Pre-analysis Plan: 'Risk and Health Policy Preferences: Evidence from the UK Covid19 Crisis'. Available from <https://osf.io/dfe8t> (accessed 23 November 2021).
- Blumenau J et al. (2021) Testing Negative: The Non-consequences of COVID-19 on Mass Ideology. Paper presented at the American Political Science Association Annual Meeting, September.
- Blumenau J, Hicks T and Pahontu RL (2022) "Replication Data for: Risk and Preferences for Government Healthcare Spending: Evidence from the UK COVID-19 Crisis", <https://doi.org/10.7910/DVN/AU6FCP>, Harvard Dataverse, V1, UNF:6:bg5pScXXIR50LVxqSuec1g== [fileUNF].
- Bol D et al. (2021) The effect of COVID-19 lockdowns on political support: some good news for democracy? *European Journal of Political Research* 60(2), 497–505.
- Esaïsson P, Sohlberg J and Ghersetti M (2021) How the coronavirus crisis affects citizen trust in institutions and in unknown others: evidence from 'the Swedish experiment'. *European Journal of Political Research* 60(3), 748–760.
- Forbes K (2016) *A Tale of Two Labour Markets: The UK and US*, technical report, Bank of England. Available from <https://www.bankofengland.co.uk/-/media/boe/files/speech/2016/a-tale-of-two-labour-markets-the-uk-and-us.pdf>
- Fox A et al. (2022) Health insurance loss during COVID-19 may increase support for universal health coverage. *Journal of Health Politics, Policy and Law* 47(1), 1–25.
- Gingerich DW and Vogler JP (2021) Pandemics and political development. *World Politics* 73(3), 393–440.
- Goos M and Manning A (2007) Lousy and lovely jobs: the rising polarization of work in Britain. *Review of Economics and Statistics* 89(1), 118–133.
- Grantz KH et al. (2016) Disparities in influenza mortality and transmission related to sociodemographic factors within Chicago in the pandemic of 1918. *Proceedings of the National Academy of Sciences* 113(48), 13839–13844.
- Hainmueller J, Mummolo J and Xu Y (2019) How much should we trust estimates from multiplicative interaction models? Simple tools to improve empirical practice. *Political Analysis* 27(2), 163–192.

- Iversen T and Soskice D** (2001) An asset theory of social policy preferences. *American Political Science Review* **95**(4), 875–893.
- Lowande K and Rogowski JC** (2021) Executive power in crisis. *American Political Science Review* **115**(4), 1406–1423.
- Margalit Y** (2013) Explaining social policy preferences: evidence from the Great Recession. *American Political Science Review* **107**(1), 80–103.
- Melcher CR** (2021) Economic self-interest and Americans' redistributive, class, and racial attitudes: the case of economic insecurity. *Political Behavior*. <https://doi.org/10.1007/s11109-021-09694-x>
- Office for National Statistics** (2020a) Healthcare expenditure, UK health accounts: 2018. Available from <https://www.ons.gov.uk/peoplepopulationandcommunity/healthandsocialcare/healthcaresystem/bulletins/ukhealthaccounts/2018> (accessed 18 November 2021).
- Office for National Statistics** (2020b) Which occupations have the highest exposure to the coronavirus (COVID-19)?. Available from <https://www.ons.gov.uk/employmentandlabourmarket/peopleinwork/employmentandemployeetypes/articles/whichoccupationshavethehighestpotentialtotheexposuretothecoronaviruscovid19/2020-05-11> (accessed 26 October 2022).
- Pahontu RL, Hooijer G and Rueda D** (2020) Insuring against hunger? Long term political consequences of exposure to the Dutch famine. Available from <https://ssrn.com/abstract=3895192>
- Rees-Jones A et al.** (2022) Experience of the COVID-19 pandemic and support for safety-net expansion. *Journal of Economic Behaviour and Organization* **200**, 1090–1104.
- Rehm P** (2009) Risks and redistribution: an individual-level analysis. *Comparative Political Studies* **42**(7), 855–881.
- Rehm P** (2011) Social policy by popular demand. *World Politics* **63**(2), 271–299.
- The King's Fund** (2014) The UK private health market. Available from <https://www.kingsfund.org.uk/sites/default/files/media/commission-appendix-uk-private-health-market.pdf> (accessed 18 November 2021).