Risk and Health Policy Preferences: Evidence from the UK COVID-19 Crisis

Jack Blumenau
School of Public Policy,
University College London
j.blumenau@ucl.ac.uk
www.jackblumenau.com

Timothy Hicks
School of Public Policy,
University College London
t.hicks@ucl.ac.uk
tim.hicks.me.uk

Raluca L. Pahontu
European Institute,
London School of Economics & Political Science
r.pahontu@lse.ac.uk
www.ralucapahontu.com

This draft: December 8, 2021

Risk and Health Policy Preferences:

Evidence from the UK COVID-19 Crisis

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 – operationalised by occupation-based measures of proximity to other people – became more supportive of increased government healthcare spending during the crisis. Using panel data which tracks UK individuals before and after the outbreak of the pandemic, we implement a fixed-effect design which 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 death rates, there is no evidence that increased health risks during COVID-19 shifted attitudes on government spending on healthcare, nor broader attitudes relating to redistribution. Our findings are consistent with recent research demonstrating the limited effects of the pandemic on political attitudes.

1 Introduction

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, of morgues running out of space for bodies. What, if any, were the consequences of these events for mass preferences regarding the funding of health systems?

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 (e.g. Ares, Bürgisser, and Häusermann 2021; Blumenau, Hicks, Jacobs, et al. 2021) and trust in governing elites (e.g. Bol et al. 2021; Esaisson, Sohlberg, and Ghersetti 2021). As such, we also contribute to a broader literature that has sought to understand the 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 have shown lasting impacts on economic inequality, public spending, and voting behaviour (e.g. Gingerich and Vogler 2021; Grantz et al. 2016).

Our conceptualization of the COVID-19 pandemic as entailing a massive shock to health *risks* allows us to additionally 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 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 that risk (Rehm 2009). And 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 healthcare spending. We study the UK, one of the countries hardest hit by the pandemic, using a generalised difference-in-difference design in which we follow a panel of survey respondents in the period before and during 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 health care preferences (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 report being more worried about catching the coronavirus. This provides reassuring evidence that our null finding on preferences over healthcare spending are unlikely to be driven by people failing to notice the changing health risks they face. 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.

2 From Health Risk to Preferences Over Health Spending

Why might we expect that heightened health risk should be met with increased demand for public healthcare expenditure? While there are many possible mechanisms that could connect features of the COVID-19 pandemic to attitudes over health spending,[1] our theoretical interest in this paper 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% as of 2018) is channelled through the public sector (Office for National Statistics 2020, Section 5), and only around 11% 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 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. Staffing shortages were feared given the viral threat to healthcare workers. As a consequence, a reasonable inference for British residents to make is that the NHS urgently needed a lot more resources in order to cope with these demands. Increasing public health expenditure provides an obvious way to ensure the (future) availability of those additional resources. Moreover, the risk–insurance logic should operate

[1] E.g. 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.

fairly clearly, here. People who faced greater risks COVID-19 infection are also those for whom the capacity constraints of the healthcare system were likely to cause most concern, as they increase the risk of leaving them with inadequate care should they become infected with COVID-19. As such, those at high risk of infection should have 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 channeled into medical research that could yield improved medical treatments or preventions (e.g. vaccines) for the increased health risks *directly* caused by the SARS-CoV-2 virus.[2] 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 health expenditure might also be helpful there.

Our goal in this paper 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.

3 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, as further described in Section A. We operationalise health risks using measures of the physical proximity between people in different occupations. As we describe in more detail below, this key treatment variable is matched to respondents in the BES on the basis of Standard Occupational Cat-

[2] N.b. the period for which we have panel survey data is entirely before the availability of approved COVID-19 vaccines.

egory (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. We implement the analyses as described in our pre-analysis plan (Blumenau, Hicks, and Pahontu 2021) below.

3.1 Measuring Public Health Expenditure Preferences

Our main outcome variable comes from a question that asks respondents to place themselves on an 11-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". 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 'Don't know' responses, and then scale this variable to have mean zero, and standard deviation one. Our estimated effects can therefore be interpreted in standard deviations of the outcome and where higher values constitute more 'left-wing' positions.

3.2 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 which involve more socially-distant interactions.

To operationalise 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 U.S 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 which 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.[3]

In Section B we provide a series of validation checks for our measure of occupational health risk. First, we show that our measure has plausible face-validity in that high-risk occupations are those that involve close interactions with others (such as medical professionals and entertainers), while low-risk occupations involve little face-to-fact contact with others (such as farmers and artists). Second, we show that the occupational COVID death rate is positively correlated with the $OccProximityRisk$ measure that we use. Third, and crucially, we show that respondents' *perceived* heath risks correlate with their *objective* health risks: BES respondents with higher levels of $OccProximityRisk$ report being more worried about catching coronavirus. These respondents are also somewhat more likely to report being more supportive of government actions to reduce the spread of COVID infections. Taken together, these results provide reassuring validation for our measurement strategy.

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-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 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.[4]

3.3 Models

We aim to identify the effect of $OccProximityRisk_i$ on health 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 $OccProximityRisk_{i,t} = OccProximityRisk_i \times Pandemic_t$, where $Pandemic_t$ is a

[3]See "Which occupations have the highest potential exposure to the coronavirus (COVID-19)?".
[4]We provide further detail on our coding choices in Section 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:

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

γ represents the effect of proximity-risk on health spending attitudes and is our main quantity of interest. 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, by construction, it eliminates the possibility of confounding that stems from systematic differences between high- and low-risk individuals that are constant 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 below.

Consistent with our theoretical discussion, we expect that those with higher infection risk will become more supportive of government spending on health and social care during the pandemic period relative to those respondents with lower infection risk (i.e. $\gamma > 0$). 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 home-workers 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 preferences. We address this issue by estimating further (pre-registered) models which 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 above.

Although our fixed-effect 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 which correlate with health risk. As the equations discussed above 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 (e.g. Rehm 2011) and it is plausible that labour market risk 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.[5]

Second, we also adjust for *realised* 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 realised 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 realisation 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 *realised* COVID infection is necessarily post-treatment to the *risk* of infection. As such, including a variable which measures whether an individual contracted COVID into our regression results will bias the estimates of our risk variable.

[5] We measure occupational unemployment rates at the 1-digit SOC code level. In tables D.3 and D.4 we replicate the analysis while controlling for unemployment rates measured at the 3-digit SOC code level. The results are substantively identical.

4 Results

We present the results of all three of the models for the *taxSpendSelf* outcome in Figure 1.[6] The left-hand panel of the figure 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 zero. 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.

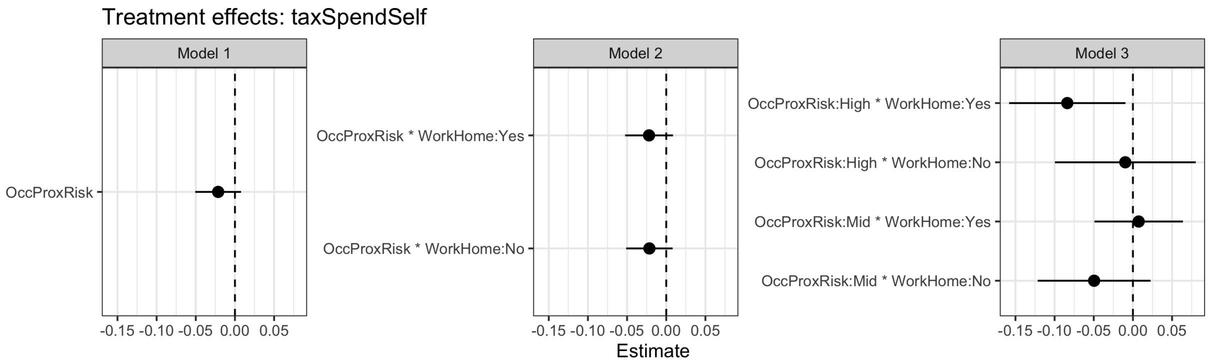

Figure 1: The figure shows estimated treatment effects from two-way fixed-effect models where the outcome variables is *taxSpendSelf*. Model 1 presents results from equation 1, which includes only the continuous proximity-risk treatment (plus controls for individual-level unemployment and the occupational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

In the centre panel of the figure, 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-risk do not affect attitudes towards government spending on healthcare.

Finally, the right-hand panel presents results from our categorical risk measure and again we 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,

[6] See appendix Tables D.1 and D.3 for full results.

we estimate a small *negative* effect for those in the highest-risk occupations and who worked from home during the pandemic.

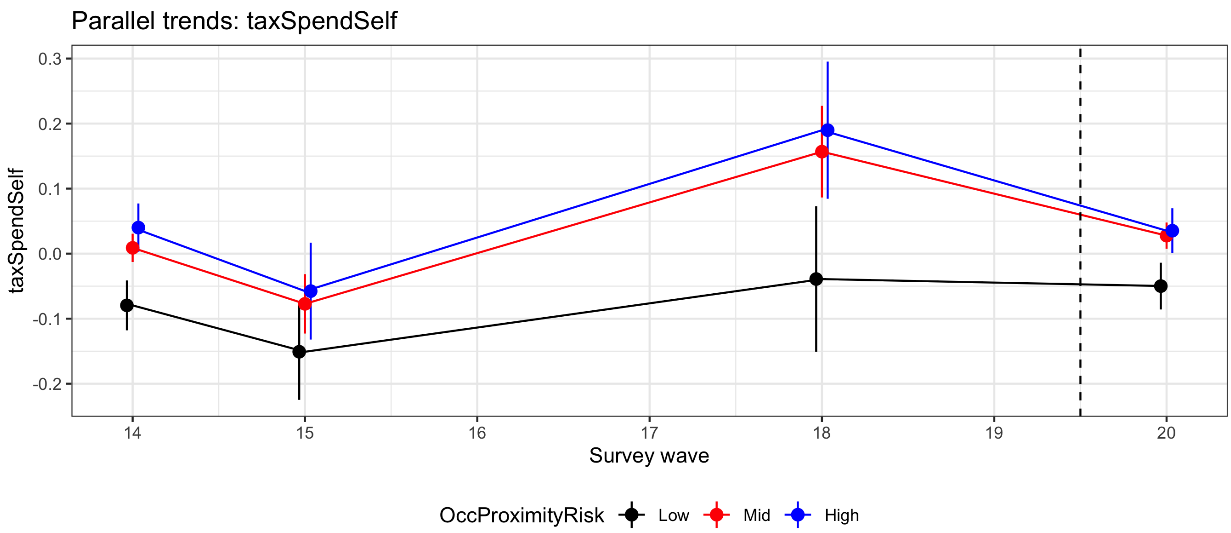

Figure 2: Parallel trends for $taxSpendSelf$.

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 2, there is no evidence of this in our data – actually, quite the opposite.

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 2 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 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 Section C, 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.

5 Conclusion

This paper studies whether individuals respond to the risk of ill health prompted by the pandemic by increasing support for 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 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, Hicks, Jacobs, et al. 2021; Lowande and Rogowski 2021).

Overall, the results in this paper 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.

References

Ares, Macarena, Reto Bürgisser, and Silja Häusermann (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(sup1): 41–55.

Blumenau, Jack, Timothy Hicks, Alan M. Jacobs, J. Scott Matthews, and Tom O'Grady (2021), "Testing Negative: The Non-Consequences of COVID-19 on Mass Ideology", Presented at the Aerican Political Science Association Annual Meeting, September 2021.

Blumenau, Jack, Timothy Hicks, and Raluca L. Pahontu (Jan. 6, 2021), *Risk and Health Policy Preferences: Evidence from the UK Covid19 Crisis*, Pre-registered analysis plan, URL: https://osf.io/dfe8t (visited on 11/23/2021).

Bol, Damien, Marco Giani, Andre Blais, and Peter John Loewen (2021), "The Effect of COVID-19 Lockdowns on Political Support: Some Good News for Democracy?", *European Journal of Political Research* 60(2): 497–505.

Esaisson, Peter, Jacob Sohlberg, and Marina Ghersetti (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.

Gingerich, Daniel W. and Jan P. Vogler (2021), "Pandemics and Political Development", *World Politics* 73(3): 393–440.

Grantz, Kyra H. 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, Jens, Jonathan Mummolo, and Yiqing Xu (2019), "How Much Should We Trust Estimates from Multiplicative Interaction Models? Simple Tools to Improve Empirical Practice", *Political Analysis* 27(2): 163–192.

Iversen, Torben and David Soskice (2001), "An Asset Theory of Social Policy Preferences", *American Political Science Review* 95(4): 875–893.

Lowande, Kenneth and Jon C. Rogowski (2021), "Executive Power in Crisis", *American Political Science Review* 115(4): 1406–1423.

Margalit, Yotam (2013), "Explaining Social Policy Preferences: Evidence from the Great Recession", *American Political Science Review* 107(1): 80–103.

Office for National Statistics (Apr. 28, 2020), *Healthcare expenditure, UK Health Accounts: 2018*, URL: https://www.ons.gov.uk/peoplepopulationandcommunity/healthandsocialcare/healthcaresystem/bulletins/ukhealthaccounts/2018 (visited on 11/18/2021).

Pahontu, Raluca L., Gerda Hooijer, and David Rueda (2020), "Insuring Against Hunger? Long Term Political Consequences of Exposure to the Dutch Famine".

Rehm, Philipp (2009), "Risks and Redistribution: An Individual-Level Analysis", *Comparative Political Studies* 42(7): 855–881.

— (2011), "Social Policy by Popular Demand", *World Politics* 63(2): 271–299.

The King's Fund (2014), *The UK private healthmarket*, URL: https://www.kingsfund.org.uk/sites/default/files/media/commission-appendix-uk-private-health-market.pdf (visited on 11/18/2021).