Supplementary Material for:

"Risk and Preferences for Government Healthcare Spending:
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
Department of Political Economy,
King's College London
raluca.pahontu@kcl.ac.uk
www.ralucapahontu.com

This version: September 11, 2022

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

A   Research Design

Our main source of individual-level survey data is the BES Combined Wave 1-20 Internet Panel (the BES panel). In our main analyses, we use three waves (14, 15, 18) of the BES fielded between May 2018 and December 2019, as well as one wave (20) fielded during the pandemic in June 2020.

As our operationalisation of health risk (described below) comes from occupation-based measures, we subset the data to only working-age respondents (i.e. those between 16 and 65 years old). After subsetting, Wave 14 of the BES includes 17,706 individuals who provided a response other than "Don't know" to the question that underpins our main outcome variable ($taxSpendSelf$), 4,014 in wave 15, 1,863 in wave 18, and 20,947 in wave 20. The smaller number of responses in waves 15 and 18 reflect the fact that this question was posed to subsets of respondents in these waves. In total, we have 44,530 of individual-wave observations of our outcome variable from 33,364 individuals across the four waves. However, given the fixed-effect model we use in the analysis, the inferences we make are primarily based on within-person variation in $taxSpendSelf$ for those who appear in both pre-pandemic and during-pandemic waves. We have 9,375 individuals who appear in more than one wave in the data, of which 8,025 appear in both wave 20 and at least one other wave.[1]

As explained in the main text, we rely on information about respondent's ability to work from home, which was only asked to employed respondents. We therefore code those who are unemployed, retired or not in paid work as "I already worked from home" and full-time students or those who were furloughed as "Yes". Schools and universities largely transitioned to online teaching throughout the survey period and those who were furloughed were paid but did not attend their jobs.

B   Parallel Trends

As noted in the main text, an important identifying assumption for our DiD design is that, absent the treatment event, the treated and control groups would have exhibited parallel trends in the dependent variable. While we presented some evidence consistent with this for the pre-treatment period for the $taxSpendSelf$ variable, in this supplementary section we present further evidence to support the parallel trends assumption.

One shortcoming of the $taxSpendSelf$ variable is that it only appears in a subset of the BES panel waves and in some cases only to a subset of respondents. The question was posed to the full sample in wave 14 ($N = 31,063$) and wave 20 ($N = 31,468$) but to smaller subsets in waves 15 ($N = 6,880$) and 18 ($3,146$). Although this does not affect the estimation of our core empirical specifications below, it limits the time-series with which we are able to investigate the pre-pandemic dynamics in

---

[1]Note that for the analysis of the $redistSelf$ variable – described below – we have a larger number of observations, and a larger number of individuals who appear in multiple waves in the data. $redistSelf$ is measured in 7 BES waves (14-20), and we have more than 18,000 respondents providing a non-"Don't know" response to this question in each wave. For this variable, there are 15,935 individuals who provide a response in both the crisis-affected wave 20 and at least one other pre-pandemic wave.

this outcome for high- and low-risk respondents, something that is important for establishing the credibility of our results.

We therefore supplement our core analysis by using a second dependent variable, $redistSelf_{i,t}$, to evaluate evidence for differential pre-treatment trends for different groups of respondents. This variable asks respondents whether the "Government should try to make incomes equal" or whether the "Government should be less concerned about equal incomes". While this question clearly taps more generic attitudes about government spending, it is nonetheless useful because it is available for a larger number of survey waves. We again treat this variable as interval-level, code it such that higher values indicate more 'left-wing' positions, drop 'Don't know' responses, and scale the variable to have mean zero, standard deviation one. The resulting variables – $taxSpendSelf$ and $redistSelf$ – are reasonably highly correlated across survey waves. The correlations between the two in the BES panel are 0.28 for wave 14, 0.25 for wave 15, 0.38 for wave 18, and 0.31 for wave 20.

Figure B.1 plots the by-wave average $redistSelf$ value for each of the three categories of the $OccProximityRisk$ variable. It shows very clear evidence of pre-treatment parallel trends. For completeness, figure B.2 replicates the main analysis for the $redistSelf$ variable, as well. Consistent with the previous analysis, we also find no evidence that increased health risks during the pandemic led to changes in policy preferences amongst the UK public. The left-hand panel depicts a precisely estimated effect of zero and we again find (centre panel) that this effect does not differ significantly for those working from home and those who kept attending their places of work. Similarly, the right-hand panel again demonstrates that whether specified as a continuous or categorical measure, those in higher proximity-risk occupations were no more likely to change their economic policy preferences than those in low-risk occupations as a result of the pandemic.
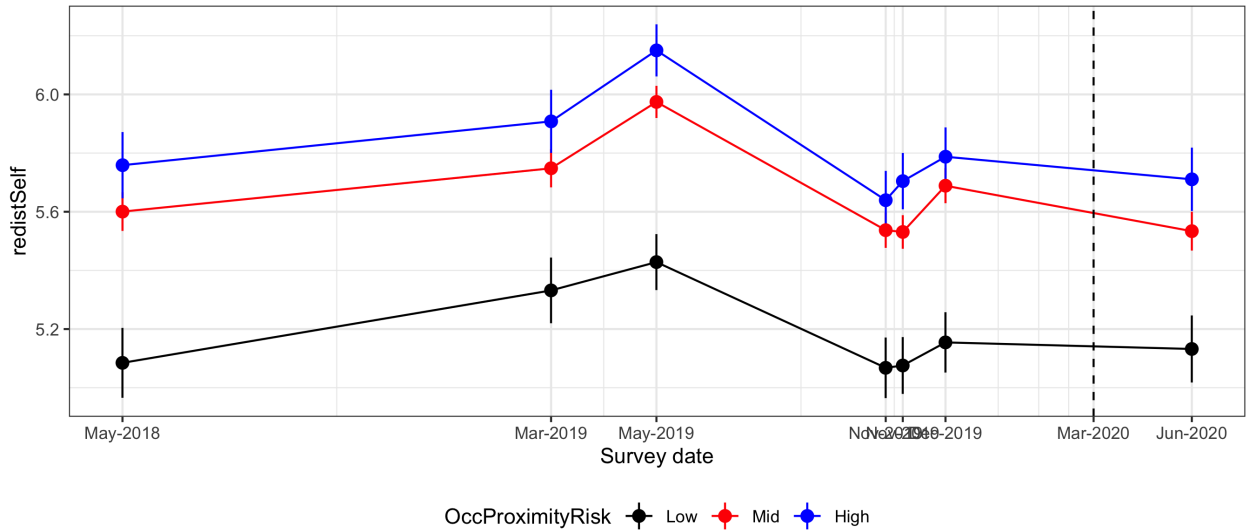

Figure B.1: Parallel trends for $redistSelf$. Points represent the average response to the $redistSelf$ variable in each survey wave for respondents in high, mid, and low categories of $OccProximityRisk$. The dashed vertical line indicates the beginning of the first COVID lockdown in the UK.

2
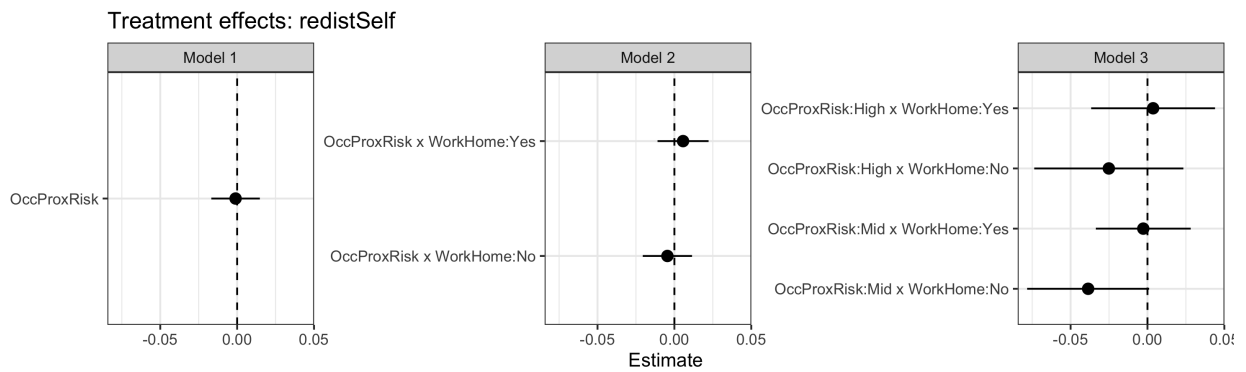
Figure B.2: The figure shows estimated treatment effects from two-way fixed-effect models where the outcome variables is $redistSelf$. Model 1 presents results from equation 1, which includes only the continuous proximity-risk treatment (plus controls for individual-level unemployment and the occupational unemployment rate measured at the 1-digit SOC level). Model 2 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.

C Regression tables

	taxSpendSelf		
	(1)	(2)	(3)
OccProximityRisk	−0.021	−0.021	
	(0.015)	(0.015)	
OccProximityRisk: Mid			−0.050
			(0.037)
OccProximityRisk: High			−0.010
			(0.046)
Unemployed	0.002	0.016	0.013
	(0.061)	(0.061)	(0.061)
Occ. unemployment rate	1.665	2.569	2.730
	(1.834)	(1.897)	(1.897)
OccProximityRisk * workHome		−0.0005	
		(0.008)	
OccProximityRisk: Mid * workHome			0.057*
			(0.032)
OccProximityRisk: High * workHome			−0.074
			(0.049)
Individual FEs	Yes	Yes	Yes
Wave FEs	Yes	Yes	Yes
Observations	31,865	22,619	22,619
Note:			*p<0.1; **p<0.05; ***p<0.01

Table C.1: Treatment effect estimates for $taxSpendSelf_{i,t}$, with occupational unemployment rates measured at the 1-digit level. Model 1 corresponds to equation 1, which includes the (unconditional) continuous treatment effect for $OccProximityRisk$. Model 2 includes the interaction between the continuous measure for $OccProximityRisk$ and $workHome$. Model 3 includes the interaction between the categorical measure for $OccProximityRisk$ and $workHome$.

	redistSelf		
	(1)	(2)	(3)
OccProximityRisk	−0.001	−0.005	
	(0.008)	(0.008)	
OccProximityRisk: Mid			−0.039*
			(0.020)
OccProximityRisk: High			−0.025
			(0.025)
Unemployed	0.039*	0.046*	0.046*
	(0.021)	(0.025)	(0.025)
Occ. unemployment rate	0.368	0.390	0.384
	(0.648)	(0.752)	(0.754)
OccProximityRisk * workHome		0.010**	
		(0.004)	
OccProximityRisk: Mid * workHome			0.036**
			(0.018)
OccProximityRisk: High * workHome			0.029
			(0.026)
Individual FEs	Yes	Yes	Yes
Wave FEs	Yes	Yes	Yes
Observations	110,557	64,562	64,562
Note:			*p<0.1; **p<0.05; ***p<0.01

Table C.2: Treatment effect estimates for $redistSelf_{i,t}$, with occupational unemployment rates measured at the 1-digit level. Model 1 corresponds to equation 1, which includes the (unconditional) continuous treatment effect for $OccProximityRisk$. Model 2 includes the interaction between the continuous measure for $OccProximityRisk$ and $workHome$. Model 3 includes the interaction between the categorical measure for $OccProximityRisk$ and $workHome$.

	taxSpendSelf		
	(1)	(2)	(3)
OccProximityRisk	−0.020	−0.020	
	(0.015)	(0.015)	
OccProximityRisk: Mid			−0.052
			(0.037)
OccProximityRisk: High			−0.005
			(0.046)
Unemployed	0.0003	0.015	0.012
	(0.062)	(0.062)	(0.062)
Occ. unemployment rate	0.318	0.408	0.433
	(0.565)	(0.589)	(0.591)
OccProximityRisk * workHome		−0.0002	
		(0.008)	
OccProximityRisk: Mid * workHome			0.060*
			(0.032)
OccProximityRisk: High * workHome			−0.076
			(0.049)
Individual FEs	Yes	Yes	Yes
Wave FEs	Yes	Yes	Yes
Observations	31,600	22,486	22,486
Note:	*p<0.1; **p<0.05; ***p<0.01		

Table C.3: Treatment effect estimates for $taxSpendSelf_{i,t}$, with occupational unemployment rates measured at the 3-digit level. Model 1 corresponds to equation 1, which includes the (unconditional) continuous treatment effect for $OccProximityRisk$. Model 2 includes the interaction between the continuous measure for $OccProximityRisk$ and $workHome$. Model 3 includes the interaction between the categorical measure for $OccProximityRisk$ and $workHome$.

	redistSelf		
	(1)	(2)	(3)
OccProximityRisk	−0.001	−0.005	
	(0.008)	(0.008)	
OccProximityRisk: Mid			−0.041**
			(0.020)
OccProximityRisk: High			−0.024
			(0.025)
Unemployed	0.044**	0.052**	0.052**
	(0.021)	(0.025)	(0.025)
Occ. unemployment rate	−0.090	−0.255	−0.281
	(0.226)	(0.262)	(0.264)
OccProximityRisk * workHome		0.010**	
		(0.004)	
OccProximityRisk: Mid * workHome			0.037**
			(0.018)
OccProximityRisk: High * workHome			0.028
			(0.026)
Individual FEs	Yes	Yes	Yes
Wave FEs	Yes	Yes	Yes
Observations	110,181	64,369	64,369
Note:			*p<0.1; **p<0.05; ***p<0.01

Table C.4: Treatment effect estimates for $redistSelf_{i,t}$, with occupational unemployment rates measured at the 3-digit level. Model 1 corresponds to equation 1, which includes the (unconditional) continuous treatment effect for $OccProximityRisk$. Model 2 includes the interaction between the continuous measure for $OccProximityRisk$ and $workHome$. Model 3 includes the interaction between the categorical measure for $OccProximityRisk$ and $workHome$.

D   Correlates of Proximity Risk

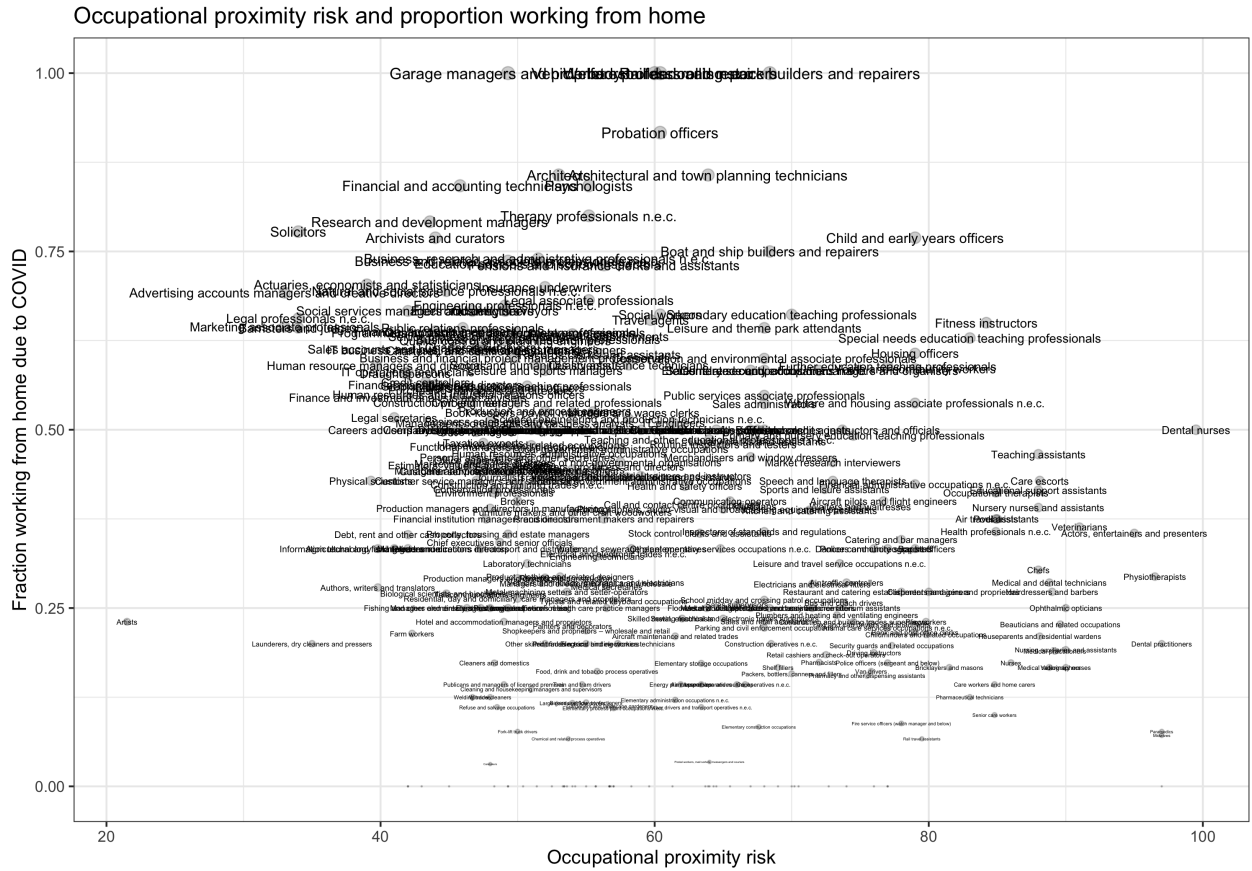## D.1   Occupational proximity risk and proportion working from home

Figure D.1: The figure shows, for each occupation, the occupational proximity-risk ($OccProximityRisk$) on the x-axis and the fraction of BES respondents (unweighted) who report working from home as a result of the COVID-19 pandemic.

D.2   Occupational proximity risk and occupational unemployment rates

One potential concern about the analysis we present in the paper is that occupational *health* risks might simply proxy for occupational *labour-market* risks, thus confounding our inferences. Although we control for occupational unemployment rates (separately for men and women) in the main analysis presented in the paper, in this section we also descriptively analyse these variables to assess whether they are, in fact, correlated.

In figure D.2, we depict (left panel) the relationship between our measure of proximity risk and the occupational employment rate before the crisis, and (right panel) the relationship between proximity risk and the *change* in the occupational employment rate from before (October to December 2019) to during (April to June 2020) the pandemic. Each point represents an occupational category from the ONS Standard Occupational Classification (SOC) Hierarchy, measured at the three digit level. We estimate the occupational unemployment rates using the Labour Force Survey and points are sized in proportion to the number of observations falling into each category in the relevant wave of the LFS. We also superimpose a weighted linear regression line onto both plots, where we use the
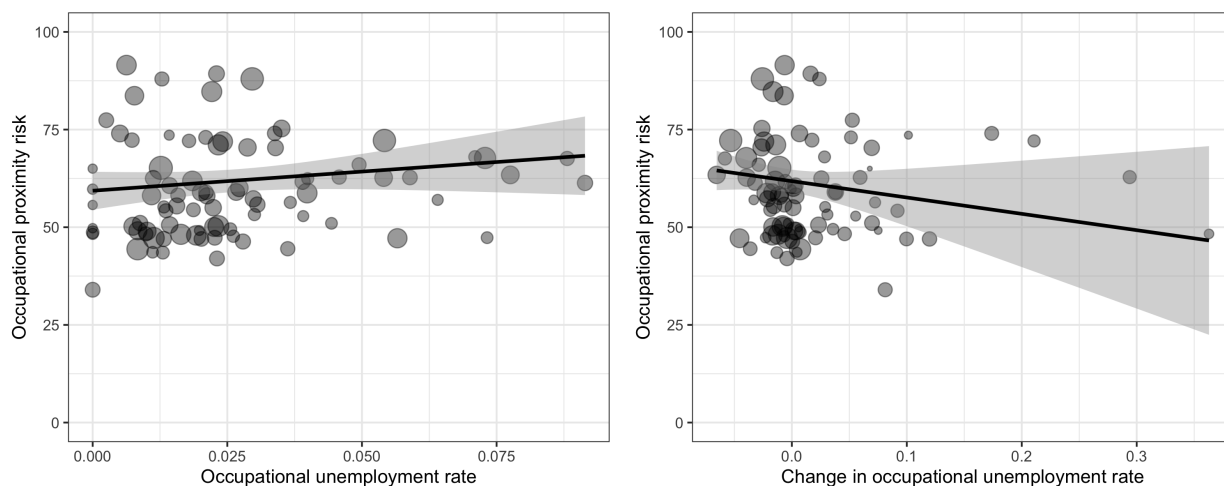
Figure D.2: **Left-panel**: Occupational unemployment rate in 4th quarter 2019 (x-axis) versus occupational proximity risk (y-axis). **Right-panel**: Change in occupational unemployment rate from 4th quarter 2019 to 2nd quarter 2020 (x-axis) versus occupational proximity risk (y-axis). Points sized in proportion to number of observations falling into occupational categories in the Labour Force Survey. Black line estimated from a weighted regression using the number of observations in each category as weights.

number of observations in each category as weights.

The figure demonstrates that occupational health and unemployment risks are broadly uncorrelated, at least when measured using these variables. Although there is a modest positive relationship between the unemployment rate and the proximity measure in the left-hand panel, and a modest negative relationship between the change in the unemployment rate and the proximity measure in the right-hand panel, neither of these is significantly different from zero in our weighted regressions (p = 0.192, left-panel; p = 0.208, right-panel). As a consequence, this analysis suggests that the health and labour-market risks faced by individuals during the crisis were indeed distinct from one another. It also – in combination with our regression analyses which control for the occupational unemployment rate (see tables C1– C4) – reinforces our belief that any effects of the proximity-risk measure we estimate are very unlikely to be confounded by occupational labour-market risks.

E Polarization/convergence effects

The analysis we present in the main paper suggests that heightened health risks during COVID had no impact on the *average* preferences for government spending on healthcare expressed by BES respondents. However, it remains possible that increased health risks could have had mean-preserving effects on the distribution of preferences by increasing support for expanded healthcare spending for some respondents while decreasing support for other respondents. In this section, we investigate whether the pandemic was associated with any such polarization or convergence in preferences, both averaging across all respondents and for different groups defined by our proximity-risk variable.
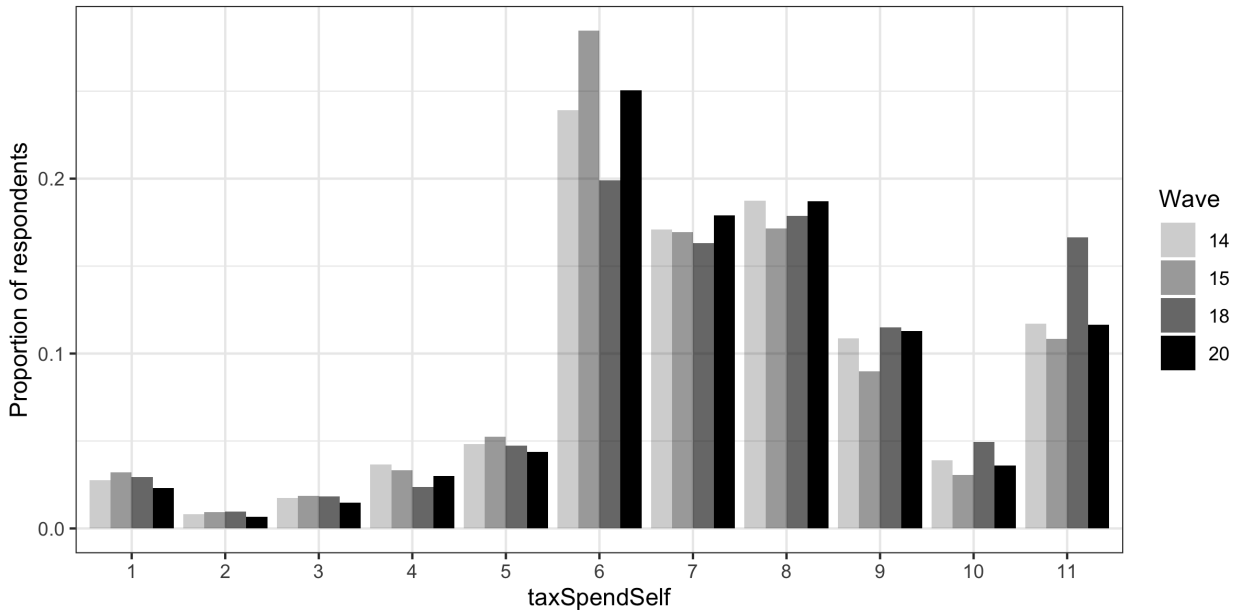

Figure E.1: Figure shows the (unweighted) distribution of the $taxSpendSelf$ variable for each wave in our data.

First, figure E.1 shows the raw (unweighted) distribution of the $taxSpendSelf$ variable for each wave in our data. Pre-pandemic waves are shown in various shades of gray, and the pandemic wave is depicted in black. At least at this aggregate level where we are averaging across all respondents, there is no evidence that the pandemic-wave $taxSpendSelf$ responses are any more or less polarized than the pre-pandemic responses.

Second, in figure E.2 we present the standard deviation of the $taxSpendSelf$ variable (alongside bootstrapped confidence intervals) for low, middle and high proximity-risk respondents for each wave in our data. If the crisis had polarizing effects on the preferences of a given group, we would expect the standard deviation to increase in the pandemic wave relative to the pre-pandemic waves. If the crisis had lead to preference convergence, we would expect the standard deviation to decrease after the onset of the pandemic.

The figure reveals that there has been very little change in the variation of responses to this question over time for any of the groups we examine. For instance, the standard deviation varies between 2.05 and 2.18 for the low-risk group; 2.09 and 2.37 for the middle-risk group; and 2.13 and 2.29 for the high-risk group. Although the changes appear to be somewhat larger between the November 2019 and June 2020 waves than between other waves, they remain trivial in magnitude.
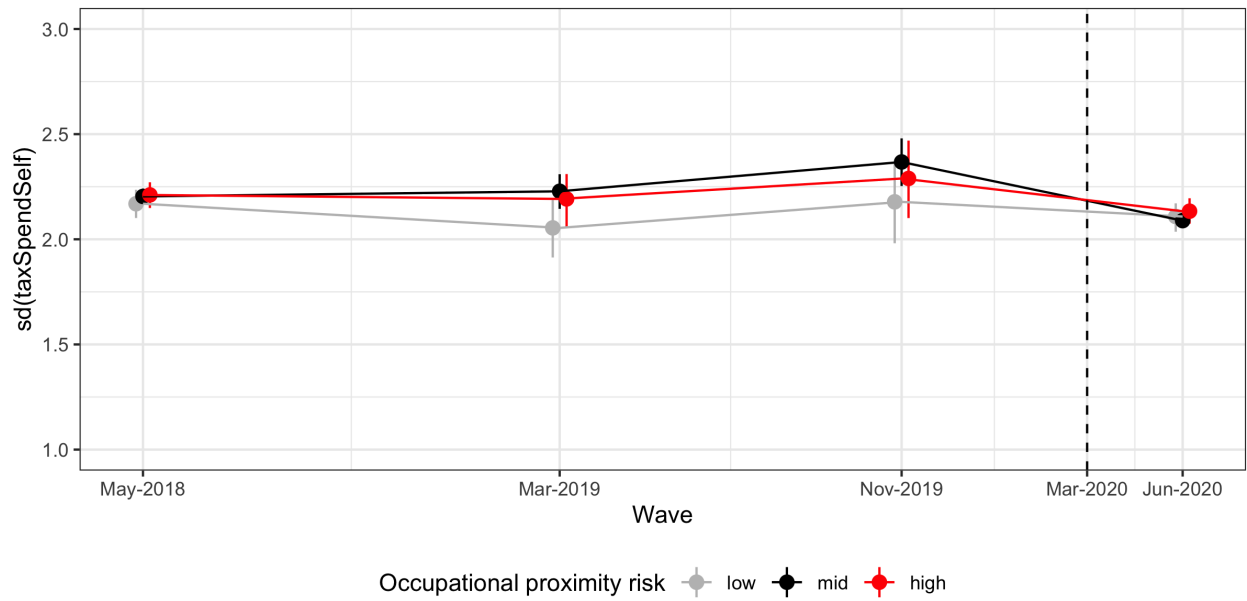
Figure E.2: Figure shows the standard deviation of the *taxSpendSelf* variable for low, middle, and high-risk respondents for each wave in our data.

Taken together, these results suggest that – in addition to having no effect on average attitudes towards government healthcare spending – the pandemic did not have either polarization- or convergence-inducing effects for the *taxSpendSelf* variable.