THE UNIVERSITY OF WARWICK

Risk, Recession, and Declining Popular Demand for the Welfare State

Lucy Barnes & Timothy Hicks

WORKING PAPER SERIES

Centre for Competitive Advantage in the Global Economy

Department of Economics

CAGE
Competitive Advantage
in the Global Economy

E·S·R·C
ECONOMIC
& SOCIAL
RESEARCH
COUNCIL

Risk, Recession, and Declining Popular Demand for the Welfare State

Lucy Barnes

School of Politics & International Relations,

University of Kent

l.c.barnes@kent.ac.uk

https://kent.academia.edu/LucyBarnes/

Timothy Hicks

School of Public Policy,

University College London

t.hicks@ucl.ac.uk

http://tim.hicks.me.uk/

February 20, 2015*

Risk, Recession, and Declining Popular Demand for the Welfare State

Abstract

How do individual preferences over welfare spending respond to economic hard times? In this paper we reconcile two prominent, opposing expectations: that recessions lead to a 'hunkering down' such that individuals become less favorable to taxation and expenditure; and that downturns, being associated with increases in risk, should lead to increased demand for government expenditure. We present a simple formal model rooted in the risk/insurance literature, to demonstrate that these two intuitions both capture important effects. While the labor market risk literature correctly predicts that *individual*-level insecurity increases support for welfare expenditure, our model shows that it also predicts that poor *macroeconomic* performance has the opposite effect. The government budget constraint links the two levels, and we demonstrate that concern about budget balance is the mechanism driving declines in support for tax-and-spend. We test our argument using British individual-level panel data from before and during the Great Recession, and use Eurobarometer data from 32 countries to probe our deficit-based mechanism. The evidence is supportive of our claims on both counts.

1 Introduction

What effects should we expect recessions and economic crises to have on the attitudes of individuals towards the welfare state? In the aftermath of the Great Recession, this question is of obvious salience. Advanced industrialized democracies have seen falls in GDP and rises in unemployment as dramatic as they were unexpected. In the UK, for example, the magnitude of the economic shock is startling: unemployment increased from 5.3% to 8.0% between 2007 and 2011, while GDP dropped by 4.7% during 2008 and barely recovered in 2009. Even by September 2013, GDP was 2.7% lower than its 2008 peak.[1] Similar stories can and have been told about other countries (Pontusson and Raess, 2012). Meanwhile, government policy has the potential to offset some of the negative consequences of these economic shocks (Rueda, 2012). But government efforts to offset negative economic effects rely on public support for such policies. Do voters support government intervention in the economy in hard times?

It is helpful to quickly turn to some data that broadly highlights a core empirical point that stands as a response to this question — although our broader theoretical response is more nuanced. Aggregating individual-level panel data from before and after the onset of the financial crisis in the UK and the USA,[2] a basic empirical fact of the Great Recession emerges. As figure 1 shows, both countries saw significant drops in the level of support for tax-and-spend policies once the financial crisis hit, with the UK only recovering less than half of this decline by 2010.

According to an important school of thought, these patterns should not be at all surprising. Conceiving of government expenditure as a luxury good, declining national income implies a marked shift away from support for taxation and government expenditure, which effectively become too expensive. Indeed, in time-series analyses of general liberal–conservative policy mood in the USA, a healthier economy has been shown to increase aggregate support for leftist policies (De Neve, 2014; Durr, 1993), with similar findings in comparative context (Stevenson,
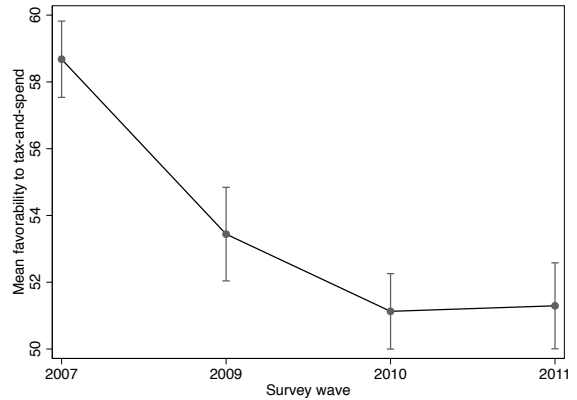
---

[1]Unemployment and GDP data are from Eurostat. Growth estimates and analysis from `http://niesr.ac.uk/press/september-2013-gdp-estimates-11587`.

[2]See section 3 (for the UK) and the accompanying supplementary materials (for the USA) for full details on how the measures in these figures were produced.
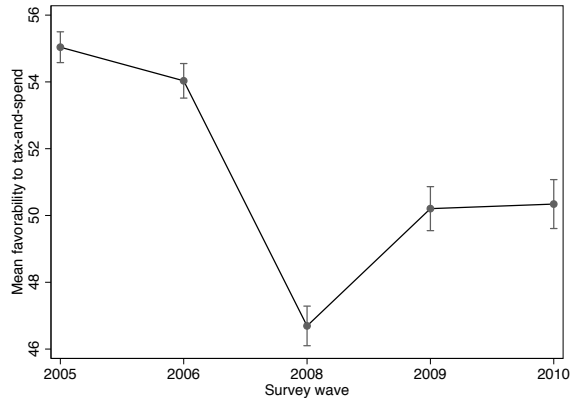
(a) UK

(b) USA

Figure 1: Average support for government taxation and expenditure (with 90% confidence intervals) before and during the Great Recession. Figures derived from individual-level panel data sets for the two countries. See section 3 (for the UK) and the accompanying supplementary materials (for the USA) for full details on how the measures in these figures were produced.

2001). At the individual level, those who have optimistic expectations about the state of the economy are more likely to focus on unemployment as an important policy problem, with the interpretation that solutions to unemployment require the kind of altruistic sacrifice that is more bearable in better economic times (Alt, 1979). Economic bad times lead to 'hunkering down' (see also Edsall, 2012). In keeping with this, recent studies of the electoral effects of bad times suggest that it is right-wing parties and ideologies that tend to benefit (Gamble, 2009; Kayser and Grafström, 2014; Lindvall, 2014).

However, a conflicting thread of literature has found contrary evidence in country-level public opinion time series.[3] For example, Erikson et al. (2002, p. 233) find that an increasing unemployment rate leads to a more liberal "policy mood" in the US. Similarly, Soroka and Wlezien (2005) found that higher unemployment led to greater aggregate support for government spending on education in the UK — although not for health. Notably, empirical findings of this sort are consistent with, but not really rooted in, the prominent and rising comparative political economy literature that conceives of government expenditure in terms of risk and insurance. Seeking

---

[3]Completing the set of possible results, studying aggregate opinion trends over a 40 year period for the US, Kenworthy and Owens (2011) find public opinion to be rather stable with respect to the incidence of recessions.

to explain patterns of welfare state preferences across countries, scholars in this tradition have emphasised the importance of individual risks — especially relating to unemployment — and the capacity of welfare states to insure these risks (Blekesaune and Quadagno, 2003; Cusack et al., 2006; Iversen and Cusack, 2000; Iversen and Soskice, 2001; Moene and Wallerstein, 2001, 2003; Rehm, 2009, 2011; Rehm et al., 2012). From this body of work, it would seem clear that recession and economic crisis should have a strong positive effect on attitudes towards welfare state expenditure and provision, coming, as they do, with large increases in economic risk for broad sections of populations. Indeed, this inference can be seen amongst prominent participants in this area. For example, summarising the risk-based literature, Ansell (2014, p. 384) indicates that, "aggregate labor market risk is considered to underpin support for national social policies", and a similar conclusion is arrived at by Marx (2014, p. 136).

In this theoretical and empirical context the contribution of this paper is to provide a reconciliation of these two apparently conflicting literatures. To this end, we show that a standard formalization of the risk/insurance theoretical framework yields predictions of increasing support for welfare expenditure as economic risk rises at the individual level, but also of decreasing support for the same as the economy deteriorates in the aggregate. While our model and its implications should be familiar to those most heavily involved in the development of the risk/insurance literature, our aim in presenting it is to correct an emerging tendency in the broader political economy literature to engage in what amounts to a fallacy of composition — to assume that the insights of the micro-level risk/insurance literature translate naturally to the aggregate level. The formalization also provides a theoretical bridge between the risk/insurance and luxury good literatures that have been so influential. We use this bridge to offer a second core contribution in this paper — to show that the government budget constraint is an important mechanism through which the macroeconomy influences attitudes towards welfare expenditure. This feature of the model has testable implications. Thus, we are able to produce a set of hypotheses that are both consistent with the risk/insurance literature and that also advance the broader literature in a strongly micro-founded way.

We demonstrate the value of the theoretical model by jointly testing the micro- and macro-level hypotheses using individual-level panel data from the British Election Study (BES), with observations from before and after the onset of the financial crisis (2005–2010). In adopting this panel data approach, we go beyond most of the existing risk/insurance literature, which has not tended to exploit within-individual, through-time variation. One notable exception to this is the recent work on the Financial Crisis in the USA by Margalit (2013). While his empirical aims are distinct from our own,[4] we show in the accompanying supplementary material that our core findings for the UK panel are replicated in the panel for the USA. Consequently, we feel more secure in the claim that our theoretical inferences are not simply specific to the UK.

The use of individual-level panels for particular countries is particularly well-suited to the testing of the micro-level hypotheses, but those panels provide far less information with which to test the macro-level hypotheses. While we, nonetheless, find strongly supportive evidence for the theory even in this empirically challenging context, we reinforce our macro-level findings with the use of Eurobarometer data. This allows us to expand our empirical range and macroeconomic N across 32 countries, and indeed down to the individual-level perceptions of these 32 macroeconomies — increasing our degrees of freedom, accordingly. It also allows us to demonstrate that the government budget constraint mechanism implied by our theory also has good empirical support.

2   Theory

For the past fifteen years, risk as a driver of social policy preferences has been one of the primary concerns of political economy. The core intuition of this body of literature is that welfare state institutions for the redistribution of income have an important insurance component. Even those

---

[4]He examines individual attitudes during the Great Recession using a panel of respondents and finds strong evidence that: personal misfortune (in the form of unemployment) makes individuals more supportive of welfare spending; this effect is larger for Republicans than Democrats; the effect is transient.

who may be net contributors to government social policy *ex post* may be better off subscribing to redistribution *ex ante* given the role of social policy in guaranteeing standards of living in the event of negative shocks. In the wake of the financial crisis, we use the theoretical framework developed in the context of individual-level risks to help us understand the dynamics of policy preferences in the context of major economic dislocation. However, we also integrate the macro-level parameters that have drawn more attention in the literature on policy mood. In doing so, we show how the risk literature to date looks in slightly the wrong place in attempting to understand change through time, and simultaneously provide more rigorous micro-foundations for the expectation of 'hunkering down' in response to crisis.

The literature on risks and redistribution centres on labor market risks, based on the dominance of labor income in determining most individuals' levels of material well-being. The biggest component of this labor market risk for most individuals is the risk of unemployment, a dimension directly implicated by the jump in unemployment rates in many countries in the aftermath of the 2008 financial crisis. Unemployment risks relate to redistributive policies because redistribution raises the income floor for all, and compresses earnings downward at the top, thus providing insurance by smoothing income across good and bad states. These bad, low-income states will often be transient, so these policies insure against that risk (Varian, 1980). The key insight of this theoretical literature is that insurance against the risk of unemployment may motivate support for redistributive policies, in addition to any pure redistribution sought to compensate low relative income.

From an empirical perspective, Phillip Rehm has led the documentation of a link between the risk of unemployment at the individual level, and the demand for redistribution (Cusack et al., 2006; Rehm, 2009, 2011). In these treatments, the risk of unemployment dictates the expected benefits of compulsory social insurance. Straightforwardly, variation in risk across individuals shapes variation in preferences. The greater the risk of unemployment, the greater the expected receipts from insurance, and the lower the expected contributions, thus the more support for government redistribution. In this context, the increase in risks associated with bad economic

times seems to imply that demand for redistribution should increase, and the failure of the Great Recession to instigate sustained increases in redistributive effort is puzzling.[5]

2.1   A Formal Model

In this section, we set out a generic and tractable version of model of tax-and-spend preferences in the presence of labor market risk. To reiterate, our goal is not to produce a new model, but rather to show that there are under-emphasised insights from an existing class of models. We present the core aspects of the model in its most accessible form, relegating mathematical derivations to the accompanying supplementary material. This allows us to be explicit about the many similarities of our modeling approach with the existing unemployment risk literature.[6] It also provides the space to draw connections with the wider literature.

The basic setup is as follows. The population is of size 1 individuals, of whom $\bar{\theta}$ are employed (at any time) and $1 - \bar{\theta}$ are unemployed. Employed individuals, indexed by i, earn an income y_i, while the unemployed have earned income of zero. Transfers from the government are flat rate, b. We set the model up as if individuals were infinitely long-lived and do not discount the future, so that their expected utility weights their utility when employed by the employment share and, when unemployed, by the unemployment share. Allowing for varying probabilities of employment at the individual level (θ_i), a general expected utility function is given by

$$V_i = \theta_i \cdot u((1 - \tau) \cdot l(\tau) \cdot y_i + b) + (1 - \theta_i) \cdot u(b) \ , \tag{1}$$

---

[5]A second, and prominent, theoretical thread relates welfare preferences to risks that arise from skill specificity (e.g. Estevez-Abe et al., 2001; Gingrich and Ansell, 2012; Iversen and Soskice, 2001). These theories were oriented towards explaining the divergence of national systems and national preferences. However, the skills demands of national economies and individuals' skills are not immediately and directly affected by short run — even if catastrophic — changes such as the Great Recession. As such, we do not emphasise this line of theory in our analysis.

[6]Regarding similarities, the basic setup shares important features with Iversen and Soskice (2001) and Rehm (2011). Precursor models that focus on aggregate unemployment, but exclude the individual-specific unemployment risk that has become so prominent in the literature, can be seen in Atkinson (1990) and Di Tella and MacCulloch (2002). A model that emphasises voting outcomes can be see in Wright (1986).

where $u(\cdot)$ is a function mapping income to utility, τ denotes a proportional tax rate, y_i is the wage rate, and $l(\tau)$ is the labor supply as a function of the tax rate. Following Iversen and Soskice (2001), we specify this latter function to capture the idea of deadweight losses from taxation:[7]

$$l(\tau) = \frac{1}{1 + \tau} \ . \tag{2}$$

Substituting equation 2 into 1, we have

$$V_i = \theta_i \cdot u \left(\frac{1 - \tau}{1 + \tau} \cdot y_i + b \right) + (1 - \theta_i) \cdot u(b) \ . \tag{3}$$

Following canonical models of redistribution and risk (Meltzer and Richard, 1981; Rehm, 2011), we assume a balanced budget for the government such that the sum of tax revenues is equal to the sum of transfer payments. This yields a government budget constraint of:

$$b = \frac{\tau \cdot \hat{y} \cdot \bar{\theta}}{1 + \tau} \ , \tag{4}$$

where $\hat{y} \equiv \bar{y}/\bar{\theta}$, and so denotes the mean income of those who are employed. In words, the transfer (b) equals the tax rate, multiplied by the average income of those employed and the fraction of the population that is employed. The average income of those employed is simply the average wage of those employed, multiplied by the amount of labor that they supply.

As equation 4 makes clear, it is because of the balanced budget assumption that the size of b depends on aggregate income and employment levels in the economy. While many of the accounts of 'hunkering down' discussed above invoke a vaguer, psychological mechanism (e.g. Durr, 1993; Gamble, 2009), this in fact provides direct material reasons for decreasing demands for redistribution in bad economic times. The same rate of taxation applied to smaller average

---

[7]Note that this specification of the deadweight losses incorporates all of the losses associated with taxation along the intensive margin for labour supply, with individual work effort (in work) affected, but abstracting from any effects on the extensive margin and thus on the aggregate unemployment rate $\bar{\theta}$. While a simplification, this is sufficient to capture the deadweight costs of taxation, and simplifies the analysis.

income, or fewer employed people earning that income, generates a less attractive level of benefits, while the cost to the individual is the same.

Since we are interested in the role that the changing magnitude of risk plays in preferences for what amounts to redistributive social insurance, the utility function $u(\cdot)$ must exhibit some degree of risk aversion. Log utility provides this risk aversion, in a moderate degree that is not inconsistent with empirical estimates (e.g. Chetty, 2006), and in a mathematically tractable form that others have found similarly useful (e.g. Alt and Iversen, 2014).

The preferred tax rate can then be recovered by solving the first order condition for τ, yielding:[8]

$$\tau^* = \frac{y_i \cdot (1 - \theta)}{y_i \cdot (1 + \theta) - \bar{\theta} \cdot \hat{y}} . \tag{5}$$

The important results for our purposes are the comparative statics and these follow rather easily:[9]

$$\frac{\partial \tau_i^*}{\partial \bar{\theta}} = \frac{\hat{y} \cdot (y_i \cdot (1 - \theta))}{(y_i \cdot (1 + \theta) - \bar{\theta} \cdot \hat{y})^2} > 0 \tag{6}$$

$$\frac{\partial \tau_i^*}{\partial \theta_i} = \frac{y_i \cdot (\bar{\theta} \cdot \hat{y} - 2y_i)}{(y_i \cdot (1 + \theta) - \bar{\theta} \cdot \hat{y})^2} < 0 \tag{7}$$

$$\frac{\partial \tau_i^*}{\partial \hat{y}} = \frac{\bar{\theta} \cdot (y_i \cdot (1 - \theta))}{(y_i \cdot (1 + \theta) - \bar{\theta} \cdot \hat{y})^2} > 0 \tag{8}$$

$$\frac{\partial \tau_i^*}{\partial y_i} = \frac{(\theta - 1) \cdot \bar{\theta} \cdot \hat{y}}{(y_i \cdot (1 + \theta) - \bar{\theta} \cdot \hat{y})^2} < 0 \tag{9}$$

These results help to clarify a number of mechanisms. First, the individual-level parameters' effects are similar to the expectations that emerge from the existing literature on redistribution. Higher income y_i always leads to lower preferred rates of taxation, and the same is true of more secure employment. This replicates exactly the logic of Rehm (2011) and Cusack et al. (2006),

---

[8]This expression constitutes the interior solution, the condition for which is that $y > \dfrac{\bar{\theta} \cdot \hat{y}}{2\theta}$. See the supplementary material for more details.

[9]Again, see the supplementary material for details on the constraints that apply here in order to sign these expressions.

and the intuition that (expected) net contributors to any scheme of social insurance will prefer its levels be lower than (expected) net beneficiaries.

However, the macro-level parameters operate entirely differently. As aggregate employment and income increase, the preferred tax rate increases. The intuition can be understood quite clearly. At the societal level, the question is how much insurance to purchase. A change in the aggregate risk of unemployment does not change the 'price' of this insurance as it remains actuarially fair by construction. Thus a change in risks does not lead to any particular substitution effect into, or out of, insurance. However, higher employment rates do have an income effect: they increase the resources available for the purchase of insurance. So, for any individual (with unchanged θ_i), a marginal increase in τ increases benefits by more than costs. The same is true for income with a logic that is commonly recognised (Meltzer and Richard, 1981).

Note that on a strict reading of the literature of individual risks and income, and preferences for redistribution, none of this is surprising, as the key variable is *relative* income or risk. However, as we have noted above, this strict interpretation is not always borne in mind in theoretical or empirical applications. Decomposing this idea into its constituent parts — the individual and the average — allows us to be clear about the moving parts of these arguments. Just as individual income and employment security reduce support for redistribution, at the macro-level, greater income and employment increase it. Thus, a model constructed to express the risk/insurance mechanism at the individual level also provides a reconciliation of macro-level intuitions about the procyclical nature of support for redistribution and the political economy literature on income and risk. It also provides an explanation of why left parties and egalitarian policymaking may have struggled since the crisis. While all incumbents may see assessments of their competence suffer during bad economic times, governments of the left will be penalised on policy grounds as well — a point made more fully by Kayser and Grafström (2014).

In sum, the model provides a number of observational implications to take to data. Specifically, the comparative statics implied by the theory can be directly tested. We should be able to see that, all else equal, aggregate level employment rates and incomes increase support for redis-

tributive social insurance, while at the individual level higher incomes and higher employment probabilities decrease support.

3   A Unified Empirical Test Using British Panel Data

In this section, we present empirical evidence from an individual-level panel data set for the UK — a case that is of substantive interest and importance. As a major financial centre, with financial markets highly integrated to the USA and the rest of the world, Britain had high exposure to the shock of the financial crisis. In short, the case clearly affords us a good degree of macro- and micro-level variation in our key economic variables that enables us to estimate our parameters of interest. As noted above, we also demonstrate in the supplementary material that the basic findings are present using panel data from the USA, and so are not idiosyncratic to the UK. Indeed, this counters the possible concern that the results from the British case may be biased due to leftist policies being 'contaminated' by their association with the Labour government that presided over the crisis in the UK — undermining support for redistribution. That similar results are found for the American case, which had a right-Republican incumbent government that was widely blamed for the crisis, is reassuring.

The use of individual level panel data is particularly appropriate given our theoretical interests. First, given that our argument focuses on the offsetting influences of changes over time in, respectively, macro-variables and individual-specific economic parameters, micro-level panel data allow us to stay as close as possible to the theory. Second, in the context of investigating macro-parameters, the panel data represent a step forward in more easily avoiding ecological fallacies. For example, the impact of macro-level unemployment can be isolated from the effect of individual unemployment. Finally, in the context of investigating the impact on preferences of individual-level characteristics, the panel data represent the best possible (observational) guard against omitted variable bias. Drawing inferences on the basis of within-individual variation allows us to minimize the impact of unobserved heterogeneity in our units affecting both the

independent variables of interest and the outcome.

One possible objection to our use of panel data can be based on a claim that individual perceptions of labor market risks are very stable through time — amounting to a kind of permanent risk hypothesis. This would be because individuals are held to form their understandings of the working of the economy, and their place in it, early in their career. However, using such a model as an objection to our use of panel data requires that one assume that not even the Great Recession was a large enough shock to provide new and relevant information with which to update these beliefs about individual labor market risk. We suggest that this is rather a strong assumption to make.

Our basic empirical strategy is to estimate models of the following form:

$$\tau_{i,t}^* = \alpha_i + \beta_1 \theta_{i,t} + \beta_2 \bar{\theta}_t + \beta_3 y_{i,t} + \beta_4 \hat{y}_{i,t} + \beta_Z Z + \epsilon_{i,t} , \tag{10}$$

where we have extended the indexing so that t denotes time. The (sometimes empty) vector of control variables is denoted by Z, with β_Z a vector of accompanying parameter estimates, and $\epsilon_{i,t}$ is the stochastic error term.[10] The individual-specific intercepts, α_i are estimated, variously, as random or fixed effects. Our interest, of course, is in the parameters β_1–β_4, whose signs are predicted by our comparative static results.

The data we use come from the British Election Study (BES) panel. The timing of the two elections which bracket the panel's years, 2005 and 2010, is near-perfect for our analysis. In 2005, few signs of the impending collapse were evident but, by 2010, the economic crisis was a dominant economic and political issue. Furthermore, the BES contains a question on whether the government should increase taxation and spending on public services, or decrease them both, in each of its five survey years during the period we study[11]

---

[10]As discussed below, for the UK models, we assume an AR1 error process such that $\epsilon_{i,t} = \rho \epsilon_{i,t-1} + u_{i,t}$.

[11]The years surveyed are 2005, 2006, 2008, 2009, and 2010). Data for the election years are taken from the pre-campaign wave of questioning, as these contain the broadest set of demographic variables, and the tax-spend preference question of interest. The BES panel dataset

In order to assess changes in preferences, we study responses to the following question on taxation and government expenditure:

> Using the 0 to 10 scale below, where the end marked 0 means that the government should *cut taxes a lot and spend much less on health and social services*, and the end marked 10 means that government should *raise taxes a lot and spend much more on health and social services*, where would you place yourself on this scale?

Asked in each wave, this question has been widely (and for a long time) used to tap support for government intervention in the British economy (Alt, 1979). Note that there is an inherent balanced-budget aspect to the phrasing of this question. The expenditure described in the question is also relatively universalistic; referring, as it does, to the health service. For both these reasons, the question accords rather well with our theoretical formulation.[12] However, our theoretical model treats government expenditures (b) and private income (y_i) identically in terms of their contribution to utility. In contrast, the operationalisation we are using here refers to an in-kind benefit from the government. However, universal cash benefits are essentially non-existent in modern advanced welfare states.[13] We can consider the slippage in the operationalisation in one of two ways. First, we might think that the non-cash nature of b in reality can be captured by the deadweight losses built into the budget constraint: the reason for the penalty to the tax rate is that b buys less utility than a comparable amount of cash. Or we can note that since money is essentially fungible, any provision of services by the government acts to release cash elsewhere in consumers' budgets such that b does act like cash in the utility function. The most obvious example in this sense is private health insurance, which, where public health insurance is not adequately provided, takes up a large fraction of household budgets.

is available at `http://www.bes2009-10.org/panel-data0510.php`, and the working dataset and code available at [this journal's website].

[12]We rescale the variable so that it ranges from 0 to 100 so that our results tables require fewer leading zeroes.

[13]Cash transfers such as child benefits, or pensions, are nominally universal but obviously conditional on non-income criteria which nevertheless are directly implicated in systematically divergent 'risks': neither retirement nor having children can be seen as an insurable risk in the same way as unemployment, at least in advanced industrial contexts.

In either case, the empirical irrelevance of universal cash transfers, plus our theoretical focus, points us to operationalise preferences focusing on the universal rather than the cash dimension of redistribution.

We measure $y_{i,t}$ using self-reported income that falls within set bands, ascribing the mid-point of the respective band to each respondent.[14] This provides a good measure of individual-level income. However, income is measured only in the pre-election surveys in 2005 and 2010. This is problematic for us, as variation in our measures over time is critical to for the estimation of aggregate-level effects (of $\bar{\theta}_t$ and $\hat{y}_t$), which are of particular theoretical interest. To solve this problem, we augment the data with measures of income for the intervening years using additional information from the British Household Panel Survey (2014). Using covariates shared across the two surveys, we estimate a model of income in the BHPS data which is then used to predict income where it is missing from the BES. The full methodology is provided in the accompanying supplementary material.[15]

In keeping with the conceptual definition of $\hat{y}_t$, we calculate average income on the basis of $y_{i,t}$. That is: $\hat{y}_t = 1/N_t \sum_{i=1}^{N_t} y_{i,t}$. We do not differentiate between different types of income (for example, from employment, pension or benefits) since all such incomes are taxed similarly in the UK. This points to a further reason that benefits in kind be used to operationalise b, as these transfers are effectively untaxed, as they are in the theoretical model.

Following Rehm (2009), we measure $\theta_{i,t}$ using gender-specific occupational unemployment rates. Respondents indicated their occupational groups in the 2005 survey and we match this to unemployment data for seven different occupations, for each gender, at the quarter during which each survey interview took place from the Office for National Statistics (2013).[16] In order

---

[14]The bands are: £0; £0–£5k; £5k–£10k; £10k–£15k; £15k–£20k; £20k–£25k; £25k–£30k; £35k–£40k; £40k–£45k; £45k–£50k; £50k–£60k; £60k–£70k; £70k or more. We top-code those in the highest band to £100k. Finally, we divide by 1000 so as to avoid the unnecessary display of leading-zeroes in the results table, below.

[15]For those concerned about the use of such a measure of income, the results for the USA in the supplementary material should provide reassurance. In that sample, income is measured directly in each survey wave.

[16]The occupational groups that we matched were: professional or higher technical; manager

to match our conceptual variable with our empirical measure, we then calculate $\theta_{i,t} = (100 - OccUnemp_{i,t})/100.$[17] Naturally enough, our measure of $\bar{\theta}_t$ is simply the national unemployment rate.

Figure 2 shows occupational unemployment rates in the UK between 2005 and 2010.[18] It shows a clear increase in economic risk beginning at the start of 2008: for all occupations there is an uptick in unemployment at this time. Of course, this is consistent with the trajectory of the headline rate of unemployment (not shown) which went from a relatively stable level around five per cent between 2005 and 2008, and then increased to almost eight per cent by the end of 2010. In terms of our theoretical parameters, this is a substantial decrease in $\bar{\theta}$.
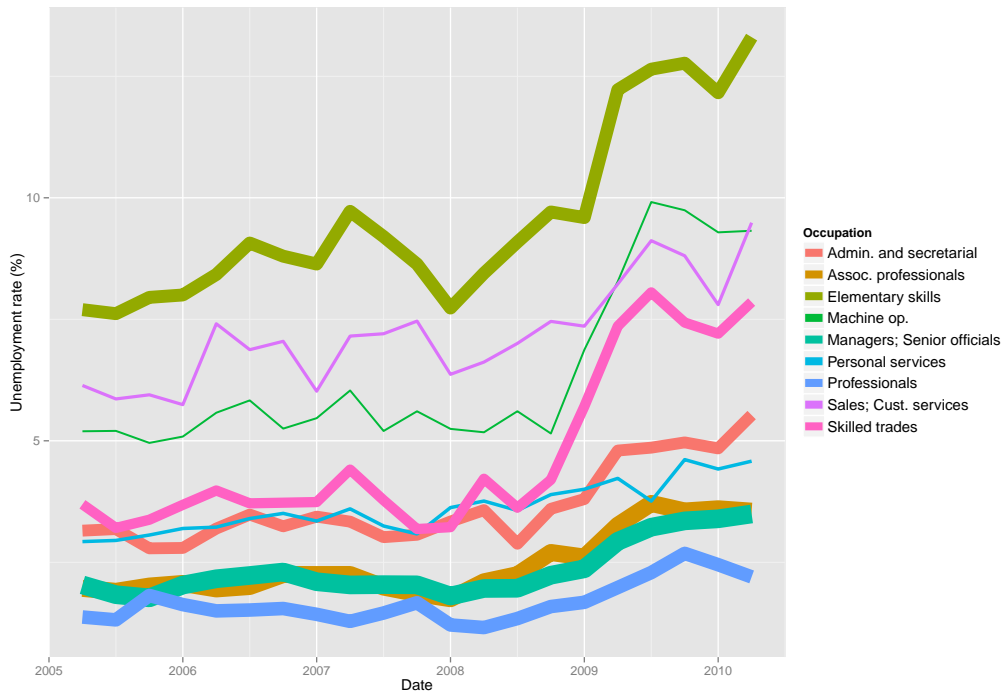

Figure 2: Unemployment rate by occupation (male and female) for the period 2005–2010. Line widths are weighted by the numbers of people within each occupational group.

---

or senior administrator; clerical; sales or services; foreman or supervisor; skilled manual; semi-skilled or unskilled manual.

[17]We use the assumption that respondents do not change occupation during the five-year period. Practically, this is because the three non-election-year surveys did not ask about respondents' occupation.

[18]To save space, we plot these rates for both genders, combined — but use gender-specific rates in the estimated models.

Our empirical aim is to estimate the parameters for each of $y_{i,t}$, $\hat{y}_t$, $\theta_{i,t}$, and $\bar{\theta}_t$, and then compare their signs to those derived in the theoretical comparative statics. To this end, we present the results from a range of model specifications with the goal of demonstrating the robustness of the inferences to different modelling choices. We estimate models that handle individual-level heterogeneity using either fixed or random individual-specific intercepts, as well as controlling for party-id. The BES data do not provide good information on individual-level unemployment status, so we neither control for it nor limit the sample on this basis. We also present a specification that includes a broad selection of demographic control variables.[19] For the fixed individual-specific intercepts, we do not employ such controls as this 'within' estimator effectively removes all between-individual variation without them.[20]

Finally, we estimate models with an autoregressive error term (AR1), as standard tests for autocorrelation in the error term (Drukker, 2003; Wooldridge, 2002) indicate that our core model does indeed suffer from this problem.[21]

The results are presented in table 1. Model UK1 is the most sparse of our specifications, using only the core theoretically relevant predictors and individual-specific random effects. It exhibits quite precise estimates for the parameters on each of $\bar{\theta}_t$, $\hat{y}_t$, and $y_{i,t}$, with each signed in the appropriate direction according to our theory. The coefficient for $\theta_{i,t}$ is rather imprecisely estimated and is in fact incorrectly signed, however. Model UK2 introduces party-id controls, measured at each wave, with the results essentially the same except for a larger and more

---

[19]Specifically: year of birth as well as dummies for gender, education level, ethnicity, marital status, housing tenure, occupation, and employment in the public or private sectors. Data limitations mean that these are not time-varying, and so can only appear with individual-level random intercepts.

[20]There remains some within-individual variation in education level and marital status, but the results from including them in the fixed effects models are essentially unchanged, so we do not report them here.

[21]We prefer this modelling strategy to a lagged dependent variable approach for a two reasons. First, it is not clear to us that a LDV model is theoretically appropriate. We include individual-level intercepts, and we do not think that the data generating process is one in which respondents deliberately give answers with reference to the previous year's 'baseline'. Thus for our purposes, the autocorrelation indicates nuisance rather than substance. Second, the LDV costs us the first of our five periods of observations, and so diminishes the sample size rather notably.

	UK1			UK2			UK3			UK4		
	b	se	p	b	se	p	b	se	p	b	se	p
$\theta_{i,t}$	6.60	9.10	0.47	15.8	8.66	0.07	-62.6	20.70	0.00	-23.2	20.88	0.27
$\bar{\theta}_t$	96.6	11.86	0.00	49.8	12.73	0.00	108.8	18.19	0.00	31.3	21.00	0.14
$y_{i,t}$	-0.079	0.01	0.00	-0.072	0.01	0.00	-0.053	0.01	0.00	-0.0030	0.02	0.89
$\hat{y}_t$	0.71	0.08	0.00	0.55	0.09	0.00	0.48	0.09	0.00	0.38	0.08	0.00
Party-id	No			Yes			Yes			No		
Demog.	No			No			Yes			No		
R^2	.02			.13			.16					
N	15174			12854			11110			10264		
Mean n per i	3.09			2.84			2.85			2.57		
Min. n per i	1			1			1			1		
Max. n per i	5			5			5			4		
ρ	0.14			0.13			0.13			0.14		
i effects	re			re			re			fe		

Table 1: Determinants of pro-welfare-spending attitudes for a panel of UK respondents (2005–2010). First, second, and third columns give OLS estimates with individual random effects. Fourth column gives OLS estimates with individual fixed effects. $\theta_{i,t}$ measures occupational employment security; $\bar{\theta}_t$ measures aggregate employment; $y_{i,t}$ measures individual-specific income; $\hat{y}_{i,t}$ measures mean income. Full details of variables is provided in the text.

statistically significant positive coefficient on $\theta_{i,t}$. This also results in substantially improved model fit, although there remains a question of whether the preference for a political party could just as well be the result of changed views on government intervention as its cause.

Model UK3 further adds the demographic controls, and yields the best performing model in terms of our theory. As well as the other three parameters remaining precisely estimated and correctly signed, the parameter for $\theta_{i,t}$ is now negative and statistically significant. Unreported estimates from a specification that excludes occupational controls suggest that their inclusion in model UK3 is to a large extent, but not exclusively, responsible for the change in sign of the $\theta_{i,t}$ paramater across these first three models. We shall return to this point when discussing the further results from table 2.

Finally, model UK4 presents estimates using individual-specific fixed effects. This model fully excludes between-individual variation and so does not allow for the partial pooling of within- and between-individual variation. As such, it is, in a sense, the most demanding of the data.

The model confirms the inference for $\hat{y}_t$, but it does not replicate the inferences for the other variables, although each of those variables retains their respectively appropriate sign. Given the necessarily coarse nature of our measure for $y_{i,t}$ and the fact that we had to estimate values for it in three of the five waves, we suggest that the partial pooling offered by the random effects models offers greater empirical scope to estimate the relevant effects.

Before proceeding, it is helpful to visualize the results from our model. We use model UK3 to produce predicted values for tax-and-spend preferences. Figure 3 focuses on the effects of aggregate unemployment (capured by $\bar{\theta}_t$) as compared to individual unemployment risk (captured by $\theta_{i,t}$). It shows how predicted values from the model vary across the survey waves for individuals with values of $\theta_{i,t}$ corresponding to the 25th and 75th percentile, while all predictors except $\bar{\theta}_t$ are held constant. The figure makes clear that there are consequential differences between individuals according to their level of occupational employment risk, such that variation across the interquartile range of θ yields approximately a 3-point change in tax-spend preferences — or about 14% of a standard deviation in the dependent variable. Meanwhile, the figure also shows that changes in $\bar{\theta}_t$ through the sample also accounted for around the same percentage of the variation. Meanwhile, figure 4 provides the analagous depiction for the results relating to y and $\hat{y}$. In keeping with general understandings of recessions, and compared to figure 3, the figure indicates that the stronger negative effect on tax-spend preferences was a result of increases in aggregate unemployment rather than falls in the income of those who remained employed. It also shows that an increase in $\hat{y}$ during the election year had a countervailing effect to that for aggregate unemployment.

We demonstrate the robustness of our inferences by adopting an approach suggested by Bell and Jones (2015), in which the random effects approach is augmented by estimating separate between-individual and within-individual effects from the explanatory variables. This is accomplished by creating by-individual means of the relevant explanatory variables, and then by-individual mean-deviations of those variables. The by-individual means are obviously constant for each individual, and so yield parameter estimates that are restricted to between-individual ef-
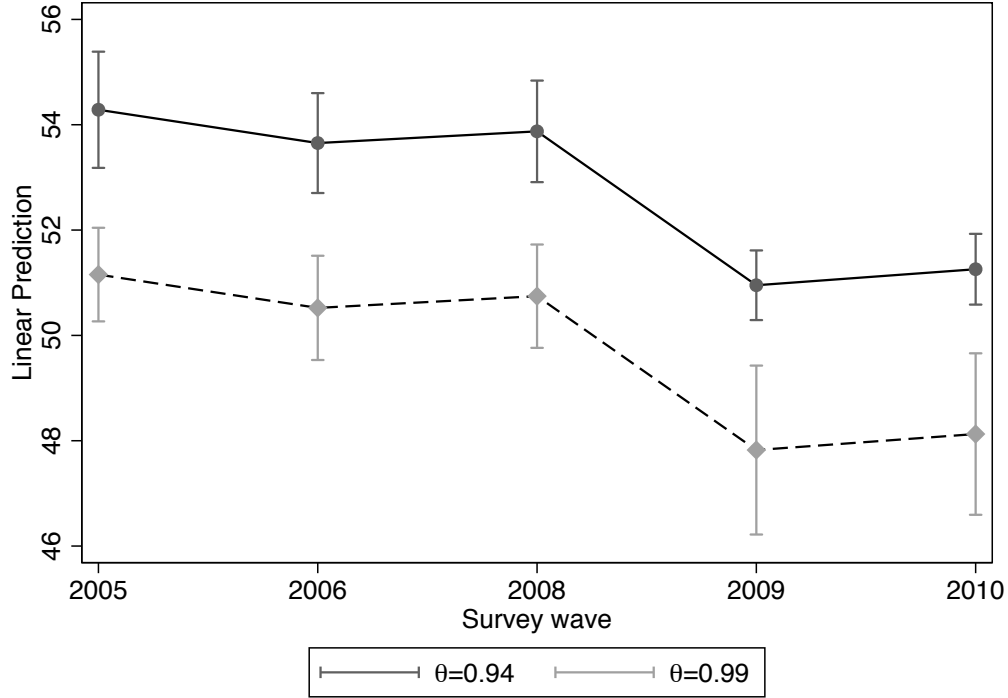
18

Figure 3: Predicted values for tax-spend preferences from model UK3, based on: holding y, $\hat{y}$, and party-id at their overall sample means, averaging over observed values of other control variables, and allowing $\bar{\theta}$ to take its by-wave mean value. Vertical bars depict 90% confidence intervals.

fects. Meanwhile, the by-individual mean-deviations are limited to variation within individuals, and so yield estimates that correspond to these effects. An important strength of the approach is that it removes the constraint that the between- and within-individual effects for each variable are the same, allowing us to directly assess which type of variation is driving our results.

Table 2 presents re-estimates of the three random effects specifications already discussed. Parameter names with the w and b superscripts denote within- and between-individual effects, respectively. As $\hat{y}_t$ and $\bar{\theta}_t$ are constant within time periods, it makes no sense to distinguish the two effects and so they enter once and untransformed. Most importantly, we continue to see similar inferences for the positive signs on each of these two aggregate measures.

The parameter estimate for $\theta_{i,t}^w$ is negative, as predicted, and quite accurately so when demographic controls are included. In contrast, the between-effect is positive, except in UK7 where it
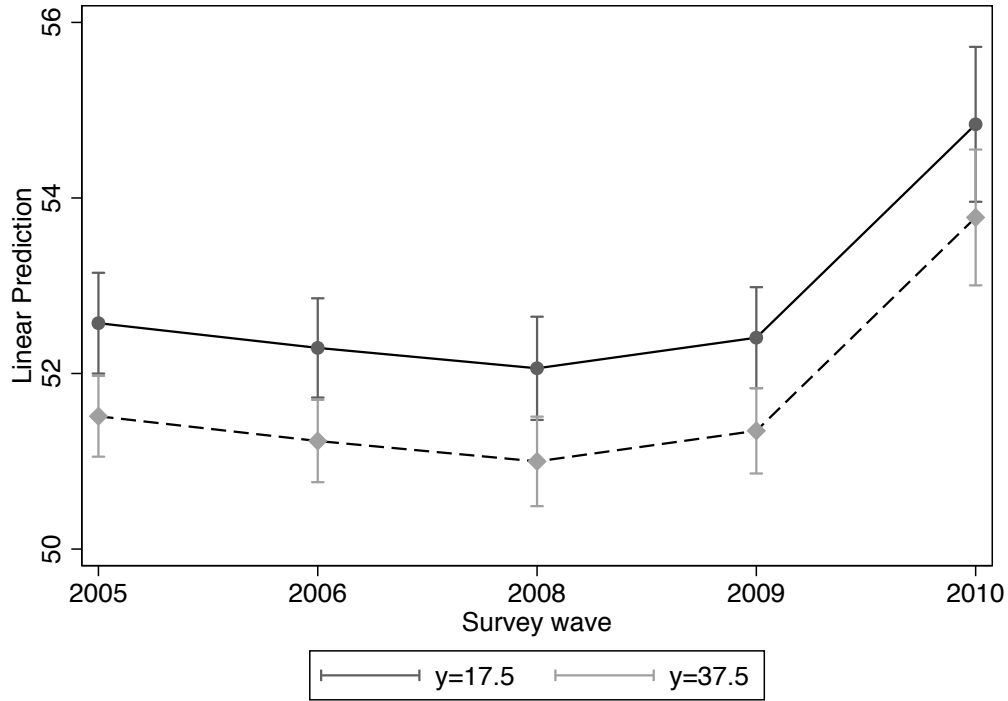
Figure 4: Predicted values for tax-spend preferences from model UK3, based on: holding θ, $\bar{\theta}$, and party-id at their overall sample means, averaging over observed values of other control variables, and allowing $\hat{y}$ to take its by-wave mean value. Vertical bars depict 90% confidence intervals.

is extremely imprecisely estimated. For a substantive interpretation of these estimates we must consider what the variables mean, as well as the other covariates included in the model. Conditional on the overall unemployment rate, individuals whose employment prospects are stronger than usual *for them*, prefer less intervention. This replicates the canonical risk models: greater risk, greater demand for redistribution. In contrast (subject to the same condition) individuals whose employment prospects are good *compared to others* prefer more intervention. One plausible explanation for this relates to heterogeneous risk preferences, and selection effects into occupational groupings. Since occupations are constant within individuals, facing a higher occupational unemployment rate *conditional on being in an occupation* leads to demand for more welfare expenditure. On the other hand, those who are more risk averse select into occupations that tend to have lower labor market risk associated with them, and so drive a countervailing

	UK5			UK6			UK7		
	b	se	p	b	se	p	b	se	p
$\theta_{i,t}^w$	-41.5	18.26	0.02	-48.3	19.52	0.01	-76.0	22.30	0.00
θ_i^b	27.2	10.39	0.01	36.2	9.79	0.00	-7.48	41.13	0.86
$\bar{\theta}_t$	128.6	15.64	0.00	92.2	17.05	0.00	116.1	18.62	0.00
$y_{i,t}^w$	-0.018	0.02	0.39	-0.018	0.02	0.41	-0.018	0.02	0.43
y_i^b	-0.11	0.01	0.00	-0.096	0.01	0.00	-0.070	0.02	0.00
$\hat{y}_t$	0.64	0.08	0.00	0.48	0.09	0.00	0.44	0.09	0.00
Party-id	No			Yes			Yes		
Demog.	No			No			Yes		
R^2	.02			.13			.16		
N	15174			12854			11110		
Mean n per i	3.09			2.84			2.85		
Min. n per i	1			1			1		
Max. n per i	5			5			5		
ρ	0.14			0.13			0.13		
i effects	re			re			re		

Table 2: Determinants of pro-welfare-spending attitudes for a panel of UK respondents (2005–2010). All models estimated with individual random effects. Within (w) and between (b) individual effects are denoted by superscripts. $\theta_{i,t}$ measures occupational employment security; $\bar{\theta}_t$ measures aggregate employment; $y_{i,t}$ measures individual-specific income; $\hat{y}_{i,t}$ measures mean income. Full details of variables is provided in the text.

between-individuals effect.

Meanwhile, the results also illustrate the interplay of within- and between-individual variation that drives the varying results across the initial specifications of table 1. However, it seems reasonable to conclude that the data limitations of $y_{i,t}$ make it difficult to estimate within-individual effects with much accuracy.

In sum, evidence suggests the following conclusions. First, aggregate-level good times — high employment rates and high average incomes — bolster support for redistributive intervention. This is in keeping with our theoretical explanation, but notably in contrast to common extrapolations from cross-individual comparisons to the macro-level. At the individual level, higher income is associated with less support for intervention. The results for expected employment chances are less clear-cut, but consistent with the argument that higher risk invites greater redistributive support once we have accounted for unobserved differences between occu-

pational groups either by occupational dummies (UK3), individual fixed effects (UK4), or using only within-individual variation (UK5–UK7). Overall, then, we find rather good support for the predictions that stem from the comparative static results of our model.

4   The Budget Constraint as the Macro-Mechanism

If the results from the BES confirm the traditional micro-logic of higher risk and lower income predicting greater support for redistribution, so the macro-results indicate that the Great Recession did lead to a 'hunkering down' (Alt, 1979; Anderson and Hecht, 2013).[22] While both consistent with our model and, perhaps, unsurprising in light of the stylised facts described in section 1, these macro-level results have two notable shortcomings. First, they are, of necessity, based on a relatively small number of observations at the macro-level — only five points in time. Although the macro-N *was* big enough to yield statistically significant results, there remains a concern about the reliability of the macro-level findings from such a sample. Moreover, the empirical results presented so far do not conclusively demonstrate the mechanism that explains this result, namely the need for budget balance.

Our goal in this section is to overcome both of these difficulties — while demonstrating further empirical support for our theory in an additional data set. Finding a large number of macro-level observations with unconfounded variation, to address the first problem, is not trivial. Cross-national variation in the unemployment rate and income levels would yield only the usual 20 to 40 countries.[23] Further, with a macro cross-sectional approach the effect of national differences in redistributive institutions and policies, and other economic policies on attitudes are difficult to separate out from those stemming from the severity of the crisis at the macro-level. Indeed, there are good reasons to think that redistributive policy and welfare state

---

[22]As noted, our supplementary material demonstrates that the macro-result for the UK is also present for the USA.

[23]Depending on the inclusion of Central and Eastern European countries along with the advanced industrial 'usual suspects'.

institutions actually explain variation in the magnitude of the crisis (Barnes and Wren, 2012; Iversen and Soskice, 2012).

We deal with the macro-N issue by using survey data across multiple countries that capture individual-level variation in perceptions of the national economic situation — in terms of both income and jobs. In doing so, we move from a macro-N of 5 to one that is conceptually far larger. Similarly, we advance our inferential claims about the importance of macroeconomic variables by showing that the mechanism implied by our theoretical model does indeed appear to be empirically supported. Thus, in this section we present an account of precisely how the macro-variables translate into lower support for redistribution in hard times.

From our model, it is clear that the level of national income and employment should affect individual preferences for redistribution via mechanical effects on the government budget constraint, and thereby on the burden to any individual of providing benefit levels at a given level of per recipient generosity. This can be seen clearly in equation 4, which documents the government budget constraint and provides the link between the macro variables and the components of individual utility in equation 3. In this way, the theoretical model makes it clear where macro level variables enter into any individual calculus about the level of redistribution. It shows that as economic circumstances deteriorate, any given level of τ becomes worse 'value for money' as it raises less revenue — which will also tend to be spread over a larger recipient population. Note that this mechanism is quite separate from any more behavioral or psychological effects of hard economic times — for example a narrowing in spheres of 'in group' identities, as sometimes implied by advocates of the 'hunkering down' logic (e.g. Gamble, 2009, p. 150).

Before proceeding to a discussion of data and estimation issues, it is helpful to clarify how we move from what we called the mechanical effects of the budget constraint to empirically relevant hypotheses. Consistent with standard models in this area, the theoretical model we outlined above builds-in a balanced budget assumption. The reason for doing this is to abstract away from theoretically distinct questions about whether and why intertemporal transfers, perhaps between generations, occur via government borrowing. While such considerations are surely of

value in the study of public finance more broadly, we wish to stay closer to the intuition and formal features of the existing risk/insurance literature in this area.

Notwithstanding this intention, the balanced budget assumption is clearly not empirically true: governments clearly do run budget deficits, and often large ones, in current circumstances. Rather than seeing this as theoretically problematic, we seek to take inferential advantage of this occurrence. Our claim is that the comparative statics of our model with respect to τ^* will also apply with respect to attitudes regarding whether or when governments should reduce the size of their budget deficit. For example, an individual with reason to demand higher levels of expenditure for insurance reasons should, other things equal, be more likely to tolerate such expenditure being deficit financed. Similarly, those experiencing a given level of labor market risk should, other things equal, be more concerned about the size of the budget deficit if they perceive the economy as a whole to be weaker. To reiterate, the benefit we see of this empirical extension is that it allows us to probe the causal mechanism that ties the individual- and macro-level arguments together. Overall, then, we have the same diverging expectations regarding the effects of the micro and macro variables on attitudes regarding the budget deficit as we did for attitudes regarding the level of tax-and-spend.

Empirically, we test these expectations using data from Eurobarometer 75.3 (fielded in 2011), which provides information on a battery of questions for around twenty thousand respondents in European Union and accession countries.[24] Respondents were asked, 'how would you judge the current situation in each of the following': 'the situation of the [nationality] economy' ($\hat{y}$); 'the employment situation in [our country]' ($\bar{\theta}$); 'the financial situation of your household' (y). and 'your personal job situation' (θ). As indicated parenthetically, we see a clear mapping from these questions to our conceptual variables of interest — especially for the macro-level variables.[25]

---

[24]Belgium, Denmark, Greece, Spain, Finland, France, Ireland, Italy, Luxembourg, The Netherlands, Austria, Portugal, Sweden, Germany, Great Britain & NI, Bulgaria, Cyprus (Republic), Czech Republic, Estonia, Hungary, Latvia, Lithuania, Malta, Poland, Romania, Slovakia, Slovenia, Turkey, Iceland, Croatia, Macedonia and Montenegro.

[25]The micro-level variables are less explicit than we would like regarding relative income position and risk of unemployment, but they are conceptually close enough that we believe it worth pursuing the analysis.

One possible drawback of these kinds of subjective measures, rather than the objective positions calculated in the British data, is that we may be capturing the impact of individuals' general personalities or outlooks rather than variation directly related to differences in economic situation. This is discussed at more length in the supplementary material, where we also present further empirical evidence to demonstrate that our theoretically relevant inferences are not driven by this alternative mechanism. Even without this extra evidence, we can say that this kind of problem would tend to work in the same direction for the individual and national level assessments. Thus, we are satisfied that the impact of general personalities or outlooks does not bias results in our favor because we predict that the effects of individual and national level assessments work in opposition to each other.

The Eurobarometer survey allows us to investigate the mechanism we hypothesise via a pair of questions on government budget balance.[26] The survey asks half of respondents whether 'means to reduce the deficit and debt... cannot be delayed', and half whether they are 'not a priority', and allows for four possible responses from 'Strongly agree' to 'Strongly disagree'.[27] We combine these two variables into a single dependent variable that captures concern about the government's budget balance as an ordered factor.[28]

Given the nature of our outcome variable, we use ordered logit models of the following form:

$$\text{logit}(b_{i,k}) = \alpha_j + \beta_1 \theta_i + \beta_2 \bar{\theta}_i + \beta_3 y_i + \beta_4 \hat{y}_i + \beta_Z Z + \epsilon_{i,t} . \tag{11}$$

As before, individuals are indexed by i and now countries are indexed by j. Thus, we include country dummies to capture all cross-national variation — most notably in the prevailing gov-

---

[26]Unfortunately, the EB 75.3 does not contain any variables which we could use to measure redistribution preferences themselves, thus we are unable to pursue the additional test of the full model to 'replicate' the British analysis in the cross-national cross-section.

[27]Graphical summaries of the key dependent and independent variables can be found in the supplementary material.

[28]Specifically, we reverse the scale of responses from those receiving the 'not a priority' version of the question, and then merge it with the other responses. In all models, we include a dummy variable for the question version to account for any level effects stemming purely from the different question wordings.

ernment budget balance.[29] As in the British case, Z is a vector of control variables, sometimes empty, and alternatively including demographic characteristics and other covariates. Specifically, we control for gender, age, the age at which the respondent completed their education, occupation (in a categorical variable which includes categories for the unemployed, students, retired and home-makers) as well as political ideology (in some specifications). Each model also includes a dummy variable indicating which version of the budget balance question the respondent saw.

The results are shown in table 3. The three specifications presented in the table show models of increasing complexity; the first includes only the substantive independent variables of interest, plus dummy variables indicating the respondent's country and the version of the outcome question they saw. Model EB2 adds the basic demographic controls as discussed above, and EB3 also incorporates political ideology. Since the substantive results are extremely consistent across these specifications, we focus on EB2. The negative and significant coefficients for both variables confirm that concern about the balance of the government budget is associated with the (perceived) levels of national income and employment exactly as predicted by the theoretical model.

Given the non-linear model and the measurement of our variables of interest, though, predicted probabilities provide a better gauge of the substantive effect size. To this end, figure 5 shows the effect of specific changes to each of the national economy variables. The figures show the change to the probability of each anti-deficit outcome resulting from the economic situation of the country, or national employment outlook, shifting from 'very bad' to 'rather bad'. In both cases, 'rather bad' is the modal and median response; while 'very bad' the second most common. The demographic variables are held at their average values, except for the other economic assessment variables which are all held at 'very bad'. Because the model estimated is an ordinal model with four levels, the predicted probabilities from holding covariates at any level is

---

[29]It is not inferentially interesting to us whether two otherwise identical individuals in different countries have differing views on deficit reduction purely because one country has a very much bigger deficit.

	EB1			EB2			EB3		
	b	se	p	b	se	p	b	se	p
$\bar{\theta}_i$	-0.07	0.02	0.01	-0.08	0.02	0.00	-0.06	0.02	0.03
$\hat{y}_i$	-0.11	0.02	0.00	-0.11	0.02	0.00	-0.09	0.02	0.00
θ_i	0.08	0.02	0.00	0.07	0.02	0.00	0.07	0.02	0.01
y_i	0.09	0.02	0.00	0.09	0.02	0.00	0.07	0.03	0.01
Demog.	No			Yes			Yes		
Ideology	No			No			Yes		
N	22347			21971			18180		
Intercepts:									
1 | 2	-1.73	0.09	0.00	-1.51	0.13	0.00	-1.57	0.14	0.00
2 | 3	-0.03	0.09	0.79	0.19	0.13	0.14	015	0.14	0.28
3 | 4	1.83	0.09	0.00	2.05	0.13	0.00	2.02	0.14	0.00
Country effects	fe			fe			fe		

Table 3: Determinants of concern about the government's budget balance for Eurobarometer respondents (32 EU and accession countries, 2011). All models contain a dummy variable indicating the version of the outcome question seen. $\bar{\theta}_i$ and $\hat{y}_i$ denote perceptions of, respectively, aggregate "employment situation" and "the situation of the [...] economy' (where higher is better). θ and $\hat{y}$ denote perceptions of, respectively, respondents' personal job situation and personal financial situation (where higher is better). See text for more details.

a vector of four probabilities, thus we obtain four first differences (one for each outcome level) from the comparison.

The figure highlights the consistent negative impact on concern for budget balance: with the one-category improvement in the national economy, the probability of the highest level of concern about the budget declines by almost two percentage points, while the two categories which represent a relaxed attitude both increase (categories 1 and 2, by 0.8 and 1.4 percentage points, respectively). The impact of the assessment of employment is not quite as large, but we see the same consistent pattern. Further, although a percentage point difference here or there may not seem like a huge substantive change, it is of a comparable size to the effects of the individual-level economic assessments (which have been consistently documented, both in our prior analyses, and in the canonical literature on risk, income and redistribution (e.g. Rehm (2009)). Alternatively, a one-level difference in national income assessment is of a magnitude comparable to the shift in attitudes between small business proprietors (typically very conservative in terms of government
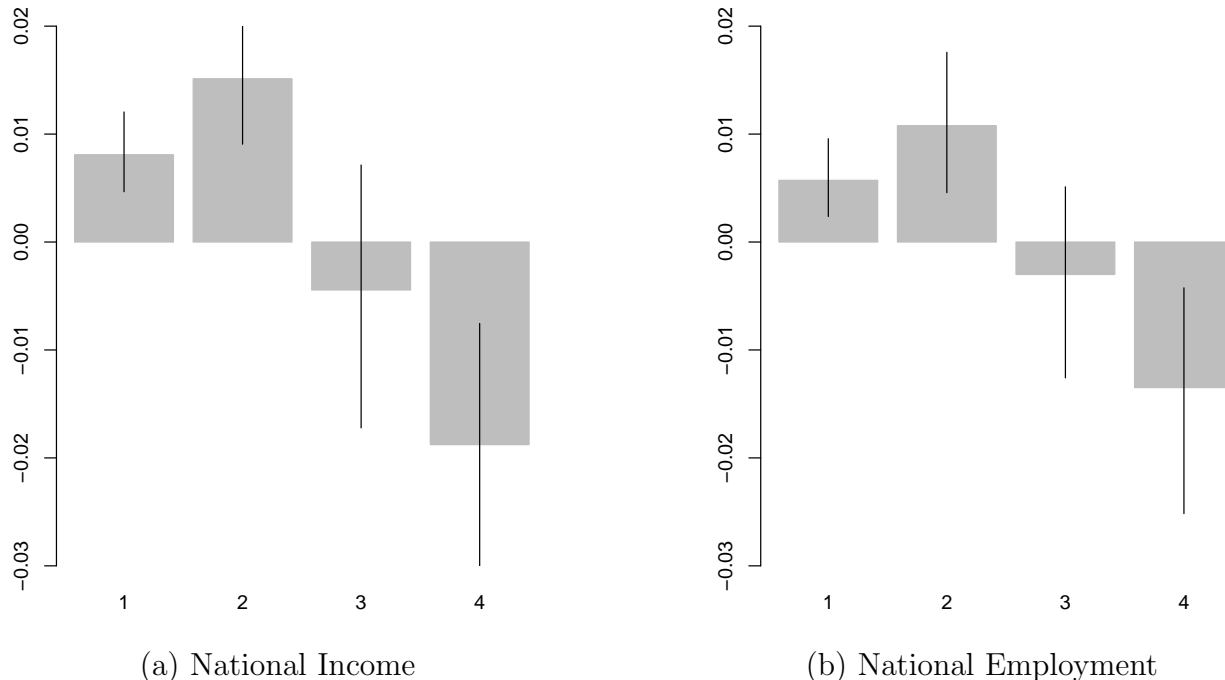
(a) National Income	(b) National Employment

Figure 5: First differences in emphasis placed on budget balance when assessments of the named variable change from 'very bad' to 'rather bad'. Grey bars indicate the estimated change in the probability of each response in prioritising budget balance, from 1 (low priority) to 4 (high priority). Black segments indicate 95% confidence intervals for these changes. Figures derived from model EB2.

fiscal policy) and office employees.[30] The analysis also confirms the effects of the individual-level variables are analagous to those in the analysis of the UK. That is to say, in contast to the macro-perceptions, those in a worse personal financial situation and those with worse job prospects are less concerned about the government budget deficit — consistent with them being more favorable towards government tax-and-spend.

One potential concern about our operationalization here relates to endogeneity. In the context of economic crisis, it might be that those who think imbalances in public finances are more problematic are more likely to assess the economic situation negatively. Or, both might be artefacts of some unobserved underlying set of attitudes. As far as the latter is concerned, there is only so far any observational study can go to assuage these omitted variable concerns.

---

[30]Note that the absence of individual income measurement prevents us from drawing the direct comparison between the size of the effect of national economic assessments and an additional dollar income amount, for example.

We discuss the possibility of an omitted 'doom and gloom' propensity in the supplementary material. However, the most plausible omitted variable is general economic ideological position, which might drive both attitudes towards budget balance, and towards the state of the economy. Specifically, right wing respondents are likely to be more hawkish about budget balance; perhaps they are also more likely to bemoan the state of the economy. However, the inclusion of a control for ideology does not appreciably affect our results. The reverse-causation link is harder to dismiss for assessments of the economy overall: concern that government redistribution plans are unsustainably underfunded budget-wise could easily translate into greater concern about the economy overall. However, for this logic to truly damage the empirical results here, we would need to be able to make the same case for concern about the national employment situation, which stands in the same relation to attitudes towards the budget. This is a much more difficult story to tell, as the health of revenues and that of employment rates are more obviously linked in a short-run trade off than they are likely to be complementary problems or priorities.

Thus, the evidence from the Eurobarometer data indicates systematic effects of (perceptions of) macro-economic variables on attitudes towards the budget. The impact of a bad national economy — in general terms, and in terms of employment — is to increase concern about the balance between expenditure and revenues. Worse assessments of the government budget by individuals perceiving greater national economic stress thus provide a key mechanism by which economic downturns may erode public support for redistribution. It follows that a public turn away from redistribution is explicable in exactly the same kinds of model as one predicting the opposite impact for individuals based on their individual, relative position. Furthermore, this logic does not require that the calculations made by voters in hard times change fundamentally. An identical calculus about risks and returns provides different preferences over redistribution depending on the macroeconomic situation.

5   Conclusions

The economic changes wrought by recession do appear consequential for policy preferences. But where the existing cross-sectional contrasts indicate that bad *positions* in terms of risk and income increase support for redistributive policy, bad *times* do not have the same effects. Our evidence from panel data that span the period before and during the Great Recession shows that higher average income levels increase support for intervention; as do lower levels of unemployment risk overall.

The theoretical model that explains these results owes a large debt to previous analyses of risk and redistribution. What we contribute, however, is explicit attention to the component parts of relative income and relative risk. This allows a clearer explanation of the different mechanisms at work, as individuals stand to gain via pure insurance as long as they are risk averse; gain from redistribution if they are relatively poor; and gain from social insurance if they are relatively likely to be unlucky. The first of these mechanisms implies that good economic times will promote intervention, while the latter two implicate low relative standing in increasing demand. In this light it is clear why recessions may dampen support for redistributive programs. As a country gets poorer together, they can afford less (pure) insurance, and by definition the relative changes must average out at the aggregate level. There may be changes in the distribution of risk, and indeed income, as a result of economic crisis. These are likely to be very interesting to scholars of redistribution preferences, but in terms of examining the consequences of economic crisis, changes to the variance and skew are — quite literally — second order effects: the changes to the levels are the most obvious effects of recession.

The empirical evidence on our main hypotheses is confirmed in two individual-level panel datasets, a luxury that is relatively uncommon in the study of political economy policy preferences. These panel data allow us to consider aggregate changes in a principled manner, accounting for individual-level heterogeneity in terms of labor market risk and income that is of direct theoretical importance. Neither of our datasets are perfect for our purposes. As is usual

with survey data, each have drawbacks in the precise questions asked or the scope and timing of the survey waves. However, we believe that the weaknesses in the data are complementary: the BES allows for greater confidence in the aggregate level results, given its longer panel and fortuitous timing, bracketing the sharp changes associated with the financial crisis. The Margalit (2013) data — which we present in the supplementary material — provide better information on individual level income, and so assuage fears that measurement error drives the British results.

The general consistency of the results across two country contexts is reassuring. However, the UK and the USA are relatively similar political economies. Thus, one clear avenue for further research is to investigate whether the same risk and income dynamics apply under different institutional arrangements. For example, one may suspect that more coordinated market economies (Hall and Soskice, 2001), where the preservation of specific skills for specialist manufacturers is integral to the system, may exhibit softer government budget constraints during downturns — reducing the procyclicality of popular attitudes towards welfare expenditure.

Moreover, we provide further evidence that the micro and macro variables are related as the model indicates through an examination of attitudes towards the budget constraint. Bad economic times are indeed perceived as putting pressure on the government budget for redistributive policy. This is an important contribution given the focus of most political economy accounts of public expenditure on the 'softness' of the government budget constraint. Beneficiaries not internalizing the tax costs of expenditure is held to increase public debt (e.g. Alesina and Tabellini, 1990; Bawn and Rosenbluth, 2006; Roubini and Sachs, 1989). While literature of this sort tends to be elite/policy-maker focused, the logic should, in principle, apply to individuals, too — at least in models such as Bawn and Rosenbluth (2006), which explicitly model electoral accountability. It is voters' preferences for particular projects that drive spending. Yet our results indicate that the aggregate budget constraint does have real consequences for individual preferences: the reduction in support for redistribution in hard times results from a recognition that resources are more limited. This poses a puzzle as to how and why budget constraints matter for individual-level attitudes, and how this can be reconciled with the somewhat inconsistent

31

accounts of elite behaviour. This is an interesting avenue for future research.

Before concluding, it is worth highlighting a further interesting aspect of our findings. We have provided theory that accounts for opposing effects of micro- and macro-level unemployment and income variables, but the simple evolutions in tax-and-spend attitudes depicted in figure 1 make plain that it is the macro effect that has dominated. It is natural to ask why this would be the case. While it is possible to ask this question of our formal model, we are not keen to push too far in this direction as to do so requires making statements that become rather dependent on functional form assumptions. More fruitfully, however, figure 2 is suggestive of one possible explanation. It shows how the UK exhibited marked variation in how the recession affected unemployment risk across different occupations. In short, the figure indicates that a majority of the population saw quite a muted increase in personal unemployment risk, but a sizeable minority saw a much larger increase. According to our theory and evidence, this pattern of change would be associated with the majority seeing the effect of a small personal risk increase being outweighed by the larger aggregate unemployment effect (driven by the less fortunate occupations). If this interpretation is correct, then it may go further in explaining conflicting findings in the literature about the effects of non-effects of recessions on aggregate opinion. The character of recessions — in the sense of the distribution of their effects — is likely to be important.

In sum, this paper has drawn together the prominent and flourishing literature on labor market risks with that more focused on aggregate-level policy mood. It has done so in a way that provides theoretical microfoundations for arguments of the latter sort, whilst reconciling what previously had seemed to be rather contradictory empirical findings regarding the impact of recessions on attitudes towards redistribution and welfare expenditure. We believe this to be an advance over existing work not only theoretically, but also empirically. It also points the way towards further comparative work that offers the hope of improving our understanding of the political consequences of recessions; and so of one of the more profound socioeconomic phenemena of our time.

References

Alesina, Alberto and Guido Tabellini (1990), "A Positive Theory of Fiscal Deficits and Government Debt", *Review of Economic Studies* 57:3, 403–414.

Alt, James E. (1979), *The Politics of Economic Decline: Economic Management and Political Behaviour in Britain Since 1964*, (Cambridge, UK: Cambridge University Press)

Alt, James E. and Torben Iversen (2014), "Inequality, Labor Market Segmentation, and Preferences for Redistribution", Paper presented at the Annual Meeting of the American Political Science Association, Washington, DC, August 2014.

Anderson, Christopher J. and Jason D. Hecht (2013), "Crisis Of Confidence? The Dynamics of Economic Opinions During the Great Recession", in: *Mass Politics in Tough Times: Opinions, Votes, and Protest in the Great Recession*, ed. by Nancy Bermeo and Larry M. Bartels, (Oxford, UK: Oxford University Press) chap. 2, pp. 40–71.

Ansell, Ben (2014), "The Political Economy of Ownership: Housing Markets and the Welfare State", *American Political Science Review* 108:2, pp. 383–402.

Atkinson, A. B. (1990), "Income Maintenance for the Unemployed in Britain and the Response to High Unemployment", *Ethics* 100:3, pp. 569–585.

Barnes, Lucy and Anne Wren (2012), "The Liberal Model in the Crisis: Continuity and Change in Great Britain and Ireland", in: *Coping with Crisis: Government Reactions to the Great Recession*, ed. by Nancy Bermeo and Jonas Pontusson, (Russell Sage Foundation) chap. 10, pp. 287–324.

Bawn, Kathleen and Frances Rosenbluth (2006), "Short versus Long Coalitions: Electoral Accountability and the Size of the Public Sector", *American Journal of Political Science* 50:2, pp. 251–265.

Bell, Andrew and Kelvyn Jones (May 2015), "Explaining Fixed Effects: Random Effects Modeling of Time-Series Cross-Sectional and Panel Data", *Political Science Research & Methods* 3:1, pp. 133–153.

Blekesaune, Morten and Jill Quadagno (2003), "Public Attitudes toward Welfare State Policies: A Comparative Analysis of 24 Nations", *European Sociological Review* 19:5, pp. 415–427.

British Household Panel Survey (Jan. 2014), *British Household Panel Survey, Waves 1–18, 1991–2009*, University of Essex. Institute for Social and Economic Research.

Chetty, Raj (2006), "A New Method of Estimating Risk Aversion", *American Economic Review* 96:5, pp. 1821–1834.

Cusack, Thomas R., Torben Iversen, and Philipp Rehm (2006), "Risks at Work: the Demand and Supply Sides of Government Redistribution", *Oxford Review of Economic Policy* 22:3, pp. 365–389.

De Neve, Jan-Emmanuel (June 2014), "Ideological change and the economics of voting behavior in the US, 1920–2008", *Electoral Studies* 34:0, pp. 27–38.

Di Tella, Rafael and Robert J. MacCulloch (2002), "The Determination of Unemployment Benefits", *Journal of Labor Economics* 20:2, pp. 404–434.

Drukker, David M. (2003), "Testing for serial correlation in linear panel-data models", *Stata Journal* 3:2, pp. 168–177.

Durr, Robert H. (1993), "What Moves Policy Sentiment?", *American Political Science Review* 87:1, 158–170.

Edsall, Thomas Byrne (2012), *The Age of Austerity: How Scarcity will Remake American Politics*, (London, UK: Doubleday)

Erikson, Robert S., Michael B. Mackuen, and James A. Stimson (2002), *The Macro Polity*, (Cambridge, UK: Cambridge University Press)

Estevez-Abe, Margarita, Torben Iversen, and David Soskice (2001), "Social Protection and the Formation of Skills: A Reinterpretation of the Welfare State", in: *Varieties of Capitalism: The Institutional Foundations of Camparative Advantage*, ed. by Peter A. Hall & David Soskice, (Oxford, UK: Oxford University Press) chap. 4, pp. 145–183.

Gamble, Andrew (2009), *The Spectre at the Feast: Capitalist Crisis and the Politics of Recession*, (New York: Palgrave MacMillan)

Gingrich, Jane and Ben Ansell (2012), "Preferences in Context: Micro Preferences, Macro Contexts, and the Demand for Social Policy", *Comparative Political Studies* 45:12, pp. 1624–1654.

Hall, Peter A. and David Soskice, eds. (2001), *Varieties of Capitalism: The Institutional Foundations of Comparative Advantage*, (Oxford, UK: Oxford University Press)

Iversen, Torben and Thomas R. Cusack (2000), "The Causes of Welfare State Expansion: Deindustrialization or Globalization?", *World Politics* 52:3, pp. 313–349.

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

— (2012), "Modern Capitalism and the Advanced Nation State: Understanding the Causes of the Crisis", in: *Coping with Crisis: Government Reactions to the Great Recession*, ed. by Nancy Bermeo and Jonas Pontusson, (Russell Sage Foundation) chap. 10, pp. 287–324.

Kayser, Mark and Cassandra Grafström (2014), "The Luxury Goods Vote: Why the Left Loses from Asymmetric Electoral Accountability", Presented at the Hertie School of Governance, 9 October 2014.

Kenworthy, Lane and Lindsay A. Owens (2011), "The Surprisingly Weak Effect of Recession on Public Opinion", in: *The Great Recession*, ed. by David B. Grusky, Bruce Western, and Christopher Wimer, (New York, NY: Russell Sage Foundation) chap. 7, pp. 196–219.

Lindvall, Johannes (2014), "The Electoral Consequences of Two Great Crises", *European Journal of Political Research* 53:4, pp. 747–765.

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

Marx, Paul (2014), "Labour market risks and political preferences: The case of temporary employment", *European Journal of Political Research* 53:1, pp. 139–159.

Meltzer, Allan H. and Scott F. Richard (1981), "A Rational Theory of the Size of Government", *Journal of Political Economy* 89:5, pp. 914–927.

Moene, Karl Ove and Michael Wallerstein (2001), "Inequality, Social Insurance, and Redistribution", *American Political Science Review* 95:4, 859–874.

Moene, Karl Ove and Michael Wallerstein (2003), "Earnings Inequality and Welfare Spending: A Disaggregated Analysis", *World Politics* 55:4, pp. 485–516.

Office for National Statistics (2013), *Labour Market Statistics, November 2013*, URL: `http://www.ons.gov.uk/ons/rel/lms/labour-market-statistics/november-2013/index.html` (Retrieved 06/13/2014).

Pontusson, Jonas and Damian Raess (2012), "How (and Why) Is This Time Different? The Politics of Economic Crisis in Western Europe and the United States", *Annual Review of Political Science* 15:1, pp. 13–33.

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

— (2011), "Social Policy by Popular Demand", *World Politics* 63:02, pp. 271–299.

Rehm, Philipp, Jacob S. Hacker, and Mark Schlessinger (2012), "Insecure Alliances: Risk, Inequality, and Support for the Welfare State", *American Political Science Review* 106:2, pp. 386–406.

Roubini, Nouriel and Jeffrey D. Sachs (1989), "Political and economic determinants of budget deficits in the industrial democracies", *European Economic Review* 33:5, pp. 903–933.

Rueda, David (2012), "West European Welfare States in Times of Crisis", in: *Coping with Crisis: Government Reactions to the Great Recession*, ed. by Nancy Bermeo and Jonas Pontusson, (New York, NY: Russell Sage Foundation) chap. 12.

Soroka, Stuart N. and Christopher Wlezien (2005), "Opinion–Policy Dynamics: Public Preferences and Public Expenditure in the United Kingdom", *British Journal of Political Science* 35:4, pp. 665–689.

Stevenson, Randolph T. (2001), "The Economy and Policy Mood: A Fundamental Dynamic of Democratic Politics?", *American Journal of Political Science* 45:3, 620–633.

Varian, Hal R. (1980), "Redistributive taxation as social insurance", *Journal of Public Economics* 14:1, pp. 49–68.

Wooldridge, Jeffrey M. (2002), *Econometric Analysis of Cross Section and Panel Data*, (Cambridge, MA: MIT Press)

Wright, Randall (1986), "The redistributive roles of unemployment insurance and the dynamics of voting", *Journal of Public Economics* 31:3, pp. 377–399.