


ARTICLE

Trends in Public Support for Welfare Spending: How the Economy Matters

Christopher Wlezien¹ and Stuart Soroka^{2*} 

¹University of Texas at Austin, USA and ²University of Michigan, USA

*Corresponding author. Email: ssoroka@umich.edu

(Received 4 January 2018; revised 19 November 2018; accepted 1 February 2019; first published online 27 August 2019)

Abstract

As spending on welfare in the United States has increased over time, preferences for more spending have remained fairly stationary. Given that previous research shows that the public adjusts its welfare spending preference thermostatically in response to welfare spending, the over-time pattern of preferences implies that something must be producing an increase in public support, but what? We address this question, focusing on individuals' demographics and a set of aggregate economic variables, both macroeconomic and distributional. Results reveal that individual-level factors matter little to the temporal variation and aggregate economics matter a lot: there are pro-cyclical and counter-cyclical elements in spending preferences and a dampening effect of income inequality over time. The combination of these variables accounts for the underlying trend in welfare spending preferences in the US, and the method used to reveal these dynamics can be used to analyze preference evolution in other spending domains and countries.

Keywords: aggregate economic variables; budgetary policy; public opinion; spending preference; welfare policy

The work that follows begins with two relatively straightforward observations. First, welfare spending in the United States has increased substantially over the past forty years.¹ Second, preferences for change in welfare spending have trended little, if at all. Although opinion has drifted upward and downward, mean support has remained fairly stationary, a pattern that is evident in other spending domains in the US and other countries (Soroka and Wlezien 2010). Consider [Figure 1](#), which shows over-time percentages of responses to a question capturing relative preferences: 'Do you think we should spend more, spend less, or spend the same amount on... welfare?'² The top panel shows trends across all response categories. It makes clear that a plurality, and often a majority, of Americans expresses support for spending less on welfare.³ It also makes clear that, even as there is real variation in welfare support over time, there is little

¹The trend in spending is apparent both in current and constant dollars, whether considered in total or on a per capita basis.

²Data are drawn from the General Social Surveys (GSS), conducted first annually and then bi-annually from 1972 to the present by the National Opinion Research Centre (NORC) at the University of Chicago. To be clear, data are missing in 14 years between 1973 and 2016: 1979, 1981, 1992, 1995, 1997, 1999, 2001, 2003, 2005, 2007, 2009, 2011, 2013, and 2015. We rely here (and below) on the primary representative sample in each year, ignoring minority oversamples. We do not apply survey weights, which are an issue since the GSS began employing a two-stage subsampling design for nonresponse in 2004, as we have to rely on the subsample that receives the 'welfare' spending question – or the 'assistance to the poor' variant, also examined below – and our regression models control for a range of demographics. For further information on GSS methodology, see www3.norc.org/gss.

³This is not the case in all spending domains, of course, though a consistent expressed preference to reduce spending is evident only for welfare and foreign aid. That said, we cannot easily interpret the neutral point, as discussed further below.

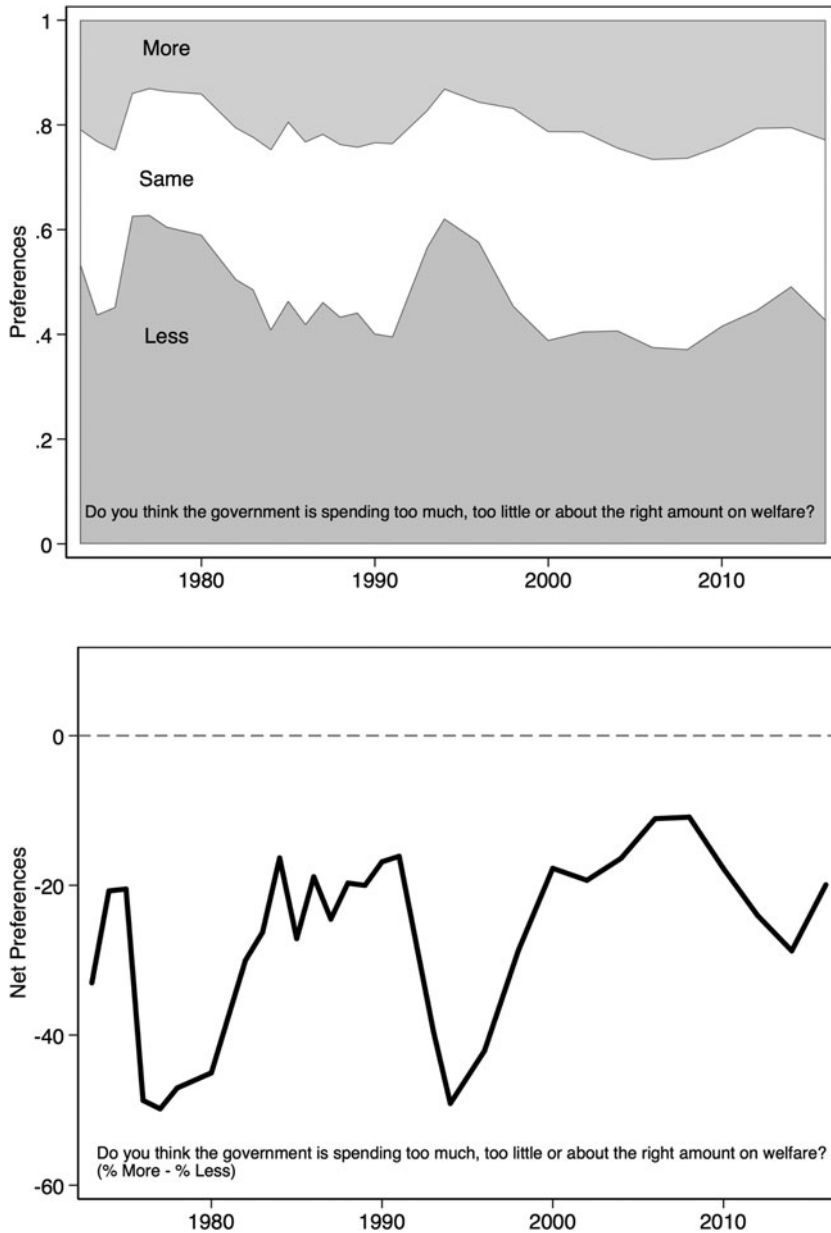


Figure 1. Relative preferences for welfare spending

systematic over-time trend. This may be most clear in the bottom panel of the figure, which shows a measure of ‘net support’ – the proportion of respondents saying they prefer ‘more’ minus the proportion saying they want ‘less’. Technically, there is a very slight positive trend (see n. 4) even as spending has increased over time.

The absence of a systematic downward trend in preferences over the forty-year time period is noteworthy because research shows that welfare spending preferences respond to policy change, increasing in the wake of spending decreases and decreasing after spending increases. What accounts for stationary preferences even as spending has increased substantially over

time?⁴ Answering this question is one primary goal of this article. There are hints in previous research but no clear answer. There appears to be an underlying positive trend in preferences for welfare spending over time, which research documents but does not explain (Soroka and Wlezien 2010). There are possible suspects, of course, and some of these have been included in previous research, such as economic conditions. As we will see, findings in that research are not in complete agreement, and so it is not clear how the economy would produce a positive trend. Recent research on growing income inequality would seem to predict a negative trend over time (Kelly and Enns 2010). There are other suspects, namely, the changing composition of the public, but this has not been the subject of previous research. This article incorporates all of these factors in a single analysis.

The objectives of what follows are fourfold. First, we want to document that underlying preferences for welfare spending have changed over time. Second, we want to provide an accounting of that change, focusing on both the evolving composition of the public and aggregate economic conditions. Third, to make this assessment, we want – and need – to isolate the effects of these factors, which we do by estimating a thermostatic model at the individual level that includes welfare spending, thus allowing us to estimate the effects of various demographic variables and macroeconomic indicators. Fourth, based on this analysis, we then want to quantify changes in the public's underlying preferred levels of welfare spending over time, and the specific contributions each variable makes to this change, both the apparent trend and the variation around it.

We accomplish this through analyses of individual-level preferences for welfare using the data illustrated in Figure 1. Our analyses reveal a steady increase in support for welfare spending over the past forty+ years driven mostly by economic growth. Specifically, we find that the economy has powerful, pro-cyclical effects on preferences for redistribution. That upward trend in preferences is *partly* offset by growing income inequality, which dampens support for redistribution, as demonstrated in previous research (Kelly and Enns 2010). There also are counter-cyclical economic effects owing specifically to unemployment, which produce short-term ebbs and flows in support for welfare around the upward trend driven by long-term economic growth. These findings effectively integrate – and refine – findings of previous research on the sources of public preferences for welfare in the US and elsewhere (such as Durr 1993; Erikson and Tedin 2015; Funk 2000; Kelly and Enns 2010; Kreps 1990; Margalit 2013; Naumann, Buss, and Bahr 2015; Stevenson 2001). In so doing, they make clear that macroeconomics are particularly impactful and in ways that previous research has not clearly demonstrated.

Our analysis also highlights that the data in Figure 1 do not tell us how much welfare spending the public wants. The question asks for preferences *relative* to current levels of spending, after all, and this complicates a comparison of policy support over time. To be clear: although in Figure 1 it appears that the American public is roughly as supportive of welfare spending now as in 1990 and the mid-1970s, this really is not the case, as spending levels increased dramatically over this time period. We know that saying ‘about right’ in 2016 is not the same thing as saying ‘about right’ in 1976. We thus can infer that the public has supported higher levels of welfare spending over time. This is critical not just for those interested in trends in welfare attitudes, but also for work that uses measures of relative preferences to indicate policy support (such as Bartels 1991, 2015; Bartle, Delleliane-Avellaneda, and Stimson, 2011; Erikson, MacKuen, and Stimson, 2002; Gilens 2012; Gilens and Page 2014; Monroe 1979, 1998; Soroka and Wlezien 2010; Stimson 1991, 2004; Weissberg 1978; Wlezien 2017).

⁴Note that the Augmented Dickey Fuller (ADF) test of Net Support (from the second frame of Figure 1) rejects the null of nonstationarity ($\rho - 1 = -0.57$, $p = 0.03$) and the KPSS test, using Newey-West automatic bandwidth selection implemented in Stata, does not reject the null of trend stationarity ($p > 0.10$). The ADF provides positive evidence of a trend, namely, a slight, statistically significant increase over time, as noted in the text, though the estimate is sensitive to the starting point of the series, particularly before/after the transition quarter between fiscal years 1976 and 1977 specified in the Congressional Budget Act. To conduct the ADF tests, we ignore the numerous (14) missing observations (see footnote 2), which, while imperfect, allows analysis of 27 observations instead of 12.

Analysis (below) of the data in [Figure 1](#) reveals a clear over-time increase in public support for redistribution in the US. This is not to say that Americans are universally supportive of redistribution, of course; it also does not mean that support for redistribution is highest when it is most needed. Understanding the nature of public preferences for welfare spending requires a careful consideration of opinion data alongside aggregate economic trends. And that consideration points to both anti- and pro-redistributive tendencies in the American public.

In sum, results contribute to both what we know about preferences for welfare spending, and how we understand public preferences for policy more generally. While we focus on welfare spending in the US, the approach may be applied fairly directly to other spending domains in the US and other countries, particularly those in which there is an evident trend in underlying preferences (see Soroka and Wlezien 2010). Indeed, even as economic variables may not prove as useful in all spending (and non-spending) domains, the general approach still can offer insight into the dynamics of the public's preferences. We return to this issue in the concluding section.

Conceptualizing preferences for (welfare) policy

Survey organizations typically measure respondents' *relative* preferences for spending, their preferences for spending change. Specifically, they ask about whether the government is spending 'too little' or 'too much'. The item thus asks respondents to compare their absolute preferences and the policy status quo. The relationship between relative preferences and policy has been set out in some detail in past work on thermostatic public responsiveness (such as Wlezien 1995, 2004). In theory, the public's relative preferences (R) represent the difference between their ideal policy (P^*) and actual policy (P), such that

$$R = P^* - P \quad (1)$$

By implication, responses change when the ideal policy changes; they also change when policy changes.⁵ That is, responses are 'thermostatic'.

A growing body of work finds evidence of thermostatic responsiveness, at least for some issues (Bartle, Dellepiane Avellaneda, and McGann, 2018; Eichenberg and Stoll 2003; Enns and Kellstedt 2008; Erikson, MacKuen, and Stimson, 2002; Jennings 2009; Soroka and Wlezien 2010; Ura and Ellis 2012; Wlezien 1995; Wlezien and Soroka 2012). In some domains there is little to no evidence of a thermostatic relationship between opinion and policy, and even where evidence exists, the effect size is not the same. There is evidence of such an effect in the US welfare spending domain, and confirmation that the public responds specifically to spending in that domain, not to the general flow of spending across the range of social programs (Wlezien 1995, 2004). As shown in the Appendix, updating that analysis through 2016 confirms these patterns.⁶ In other words, welfare spending preferences are substantially unique, not the mere reflection of a single overarching public 'mood' in spending, and specific responsiveness to welfare spending – not broader social spending – in part accounts for that domain-specificity. This makes more meaningful our investigation of preferences in the domain.

Now, equation 1 highlights a point made in the introduction of this article: relative preferences (R) and absolute preferences (P^*) are not the same things. Ignoring the difference, that is, treating relative preferences as though they are absolute, clearly would be problematic. This is not to say

⁵These relationships are not linked to change specifically – they are true for levels as well. So we might just as easily write that responses differ when the ideal policy differs; they also differ when policy differs.

⁶The Appendix contains two sets of related results. First, in Appendix section A, factor analysis of aggregate spending preferences between 1972 and 2016 indicates that, while welfare preferences correlate positively with those in other social domains, they also exhibit substantial (46 per cent) unique variance. Second, in Appendix section D, regression analysis of welfare spending preferences that includes spending in a set of social domains shows that the public responds only to welfare spending, not to the trend in the set of related areas.

that people have absolute preferences for policy, that is, the levels they prefer. Almost no one can say exactly how much they think should be spent on welfare. An ‘ideal’ level – or range of levels – that would satisfy them presumably does exist, however. The idea is very much in line with Stimson’s (1991) ‘zone of acquiescence’, where a range of policy choices would have the public saying that policy is ‘just right’.⁷ In short, the argument there, and here, is that in most policy domains absolute preferences are difficult, and maybe even impossible, for respondents to express; *not* that there is no such thing as an absolute preference.

How has past work dealt with absolute preferences? One way to capture these is with proxies, such as economic measures. Work taking this approach typically does not actually acknowledge the connection between the measures and absolute preferences – researchers simply include macroeconomic measures as controls. To the extent these variables predict relative preferences controlling for policy, one can infer that they at least partly capture underlying absolute preferences. (See above for references to past work on this issue.)

Thus there is work that deals with absolute preferences, at least implicitly; that is, the connection between survey measures and absolute preferences typically is not directly considered. Our objective here is to explore relationships between economic measures and relative preferences, as a way of both revealing and explaining absolute preferences. Doing so allows us to speak both to (a) methodological issues in capturing absolute policy preferences and (b) the sources of preferences for welfare spending, a salient and important domain, and one for which we have reliable measures of relative preferences, proxies for absolute preferences and spending itself.⁸

Modeling preferences for welfare spending

Our analysis of welfare spending preferences builds on the thermostatic model depicted in Equation 1. Because we cannot directly capture the public’s ideal policy (P^*), and because the variables are measured using different metrics, it is necessary to rewrite the equation as follows:

$$R_t = \alpha + \beta_1 U_t + \beta_2 P_t + e_t, \quad (2)$$

where R is the public’s relative preference for policy, P is current levels of policy and U is a set of additional exogenous variables affecting R . To be clear: Equation 2 differs from Equation 1 only slightly – R is a function of P but we rely on proxies (U) for P^* .

The thermostatic model initially was conceived to apply longitudinally, where relative preferences respond dynamically to policy and other factors over time (t). But the model can apply across space as well, that is, we can introduce subscript j to indicate units:

$$R_{jt} = a_j + \beta_1 U_{jt} + \beta_2 P_{jt} + e_{jt}, \quad (3)$$

either at particular points in time or across time (Wlezién and Soroka 2012). The model can be adapted to fit individual-level analyses as well, as follows:

$$R_{it} = a_i + \beta_1 U_{it} + \beta_2 P_{it} + e_{it}, \quad (4)$$

where i denotes survey respondents. Just as we can look at variation in R across space (j), we can look at variation in R across individuals (i), based on varying levels of policy (P) and other factors

⁷The tendency relates to the ‘Goldilocks principle’, which is fairly widely recognized in the social sciences, for example, Goldilocks pricing. And note that we would amend his characterization somewhat to reflect the likelihood that people would prefer policy closer to the middle of the zone than at its edges. Also see Key’s (1961) characterization of opinion ‘dikes’.

⁸As noted earlier, welfare spending is a domain for which preferences are specific to the area, that is, they are not driven solely by the global forces of policy mood (Wlezién 1995; 2004). This is of obvious importance when modeling absolute preferences for spending in the particular area. Also see footnote 16 and the associated discussion in the Appendix, section D.

(U), including various individual-level characteristics, at time t . This is the micro-level variant of the more common macro-level thermostat model, and is the basis for the analysis that follows.

In our investigation, the dependent variable, R_i , represents individuals' responses to the relative preference question described above. It is a trichotomous variable tapping support for less, the same and more spending, and varies across both individuals and years. These are the same GSS data used to produce Figure 1 above, although in this case we assess stacked individual-level data, that is, a series of repeated cross-sectional surveys, in which the respondents are different from one survey to the next. This allows us to account for the contribution of individual-level characteristics to the trend in spending preferences over time. It is necessary to account for spending, which we measure using per capita national welfare appropriations in 2009 USD, for which reliable data are available since fiscal year 1976.⁹ The measure clearly varies over time but is the same for all individuals in a given year.

We have argued that, because it shows *relative* preferences, Figure 1 obscures what must be a large over-time increase in absolute preferences for welfare spending. Accounting for context thus is critical. We focus on the economic climate here, using various measures (for U) based on previous research. We rely first on measures used in previous analyses of the public's policy preferences – unemployment and real per capita GDP (in thousands of 2009 USD).¹⁰ Based on some previous research, we would expect increasing unemployment to lead to greater demand for spending (Erikson, MacKuen, and Stimson, 2002). This expectation is highly intuitive, and has roots in basic self-interest, where the unemployed are expected to be more supportive of government welfare (Margalit 2013; Naumann, Buss, and Bahr 2015). There are reasons to expect a 'sociotropic' effect as well, where the employed themselves demand more spending when unemployment increases, owing to the effects of broader societal interest (Funk 2000; also see Erikson and Tedin 2015).¹¹ Based on other previous research, we would expect per capita GDP to increase demand for spending (Durr 1993; Stevenson 2001). This expectation also is intuitive, if less obvious, and relates to what in economics is referred to as an 'income effect', where demand for goods varies with discretionary income (Kreps 1990). Note that the expectations relating to the two economic variables contrast, in that the former is counter-cyclical and the latter is pro-cyclical. Of course, it may be that both are at work. We also include a measure of income inequality in the model, namely, the Gini index (see n. 10). This is less typically found in analyses of the public's policy preferences, but is especially relevant given our current focus on welfare spending. Following the logic of redistribution, we might expect that, as inequality increases, support for welfare spending also increases. Recent empirical work nevertheless suggests a very different link between income inequality and public preferences, whereby increasing inequality leads to support for less spending on welfare (Kelly and Enns 2010).¹² While counter-intuitive on its face, there is a self-interest logic for such a pattern, at least among segments of the

⁹The data are from table 5.1 of the *Historical Tables* published by the Office of Management and Budget (OMB), as of June 2018. In line with Wlezien (1995), we rely on 'income security' minus three subcategories: 'general retirement and disability insurance', 'federal employee retirement and disability', and 'unemployment compensation'. This leaves 'housing assistance', 'food and nutrition assistance' and 'other income security', the latter of which includes Supplemental Security Income and earned income credit but also child credit, temporary assistance for needy families, and Making Work Pay. There are various other accounts under other income security, but those listed above add up to approximately 90 per cent of 'other' income security spending. Logging the spending variable does not alter the direction or significance of any of the estimated coefficients. See Appendix Table A6.

¹⁰Unemployment is drawn from the FRED database (UNEMP) and real GDP is from the OMB's Historical Tables (table 10); the Gini index, introduced below, is from the US Census Bureau (table A-2 Selected Measures of Household Income Dispersion). All measures are the versions distributed as of June 2018. These measures are correlated to varying degrees over time. We thus have to keep in mind that any one coefficient estimate is conditional on the impact of other related and sometimes countervailing indicators.

¹¹The expectation also fits with the economic function of government relating to stabilization (Musgrave 1985).

¹²Including the Gini index may introduce issues of endogeneity, since the level of inequality will be driven in part by spending on welfare, which may reflect preferences for spending. It is not clear how this would account for the negative relationship Kelly and Enns (2010) found, and excluding the variable also makes little substantive difference to the estimates for

population who would be expected to pay for the additional spending. That the pattern appears to hold across income groups (see Kelly and Enns 2010) implies that the mechanism is not entirely or even mostly self-interest.

Finally, our models include a host of demographic variables, based largely on existing work examining individual-level variation in support for redistributive policy (especially Margalit 2013, alongside other research referenced above). Specifically, we include: age (in years),¹³ gender (binary, where 1 = female), education (categorical, with binary variables for high school and more than high school with less than high school as the residual category),¹⁴ work status (categorical, with binary variables for unemployed and other with employed as the residual category), family income (categorical for quartiles, with the first quartile as the residual category), race (categorical, with binary variables for Black and Hispanic with White as the residual category), and region (categorical; regions are shown in Table 1, and New England is the residual category). Our aim here is to explore exogenous sources of policy preferences, so we do not include partisanship in these models. That said, including party ID makes little difference to the coefficients for the other variables in the model – see Appendix Table A3.

Analyses of relative preferences

Table 1 shows results for the model of support for welfare spending. Despite having a three-point scale as our dependent variable – less, same, more – we present results using a simple GLS model. Doing so simplifies the interpretation of coefficients and also is helpful with some of the macro-level estimations that follow. We nevertheless include ordered logit results in Appendix Table A2, in which the pattern of results is very similar, confirming that the results presented here are not the product of using a linear model. Given the panel structure of our data, models include either yearly random effects (for GLS models) or clustered standard errors (in the case of ordered logit).¹⁵

To begin, notice in Table 1 that per capita welfare spending has the expected thermostatic effect on preferences for more spending, as the coefficient (-0.055) is significantly less than 0. Other things being equal, increases (decreases) in welfare spending lead to decreases (increases) in relative preferences for policy.¹⁶ The magnitude of the coefficient suggests the following: a \$191 increase in per capita welfare spending (which is a one-standard deviation change in spending in our sample) is associated with an average 0.11-point decrease in support for more spending. Put differently: the range in per capita spending over the time period investigated here – roughly \$640 (from \$400 to \$1,040) – accounts for a 0.35 shift in spending preferences over the time period, 89 per cent of the total range of spending preferences – 0.39 points – over the same period.

Other factors matter for relative spending preferences as well. Demographic variables are important, particularly age, race, ethnicity and income. Effects are in line with what past work has shown: for instance, age and income tend *ceteris paribus* to be negatively correlated with support for welfare spending, and non-Whites tend to be more supportive of welfare spending. We do not wish to interpret these coefficients in much detail here, since there is a good deal of previous work focusing on the individual-level correlates of welfare state support. Nevertheless, each

the other economic variables, as we will see. Given all of this, as well as the fact that the including the Gini Index allows us to account for the underlying trend in preferences, we retain the variable in the model.

¹³Note that we test alternative specifications for age. Details are included in Appendix section H.

¹⁴We also test alternative specifications for education; see Appendix section H.

¹⁵Some might find GLS estimation a conservative approach, though we note that it makes almost no difference to the results. Given that respondents in each year share a common (correlated) experience that can impact estimates, we report the GLS results in the body of the article.

¹⁶Note that we can test directly the possibility that welfare preferences are responding to welfare spending rather than general social spending by including the latter in our models. See Appendix Table A4 for the results of that analysis, which show that responsiveness in the welfare spending domain is specific, based on the more fully specified model, in column three of Table 1, described below. These findings confirm what has been demonstrated in previous research (Wlezien 1995; also see Wlezien 2004).

Table 1. Modeling relative preferences, across individuals and time

	Relative preferences for spending: B (s.e.)		
	Column 1	Column 2	Column 3
Spending (100s, per capita) _t	-0.055*** (0.007)	-0.068*** (0.007)	-0.075*** (0.005)
Time trend	0.011*** (0.001)	-0.026*** (0.006)	-0.004 (0.005)
Unemployment _t		0.062*** (0.011)	0.044*** (0.007)
GDP (1,000s, per capita) _t		0.069*** (0.011)	0.061*** (0.008)
Gini (*100) _t			-0.080*** (0.011)
Female	0.007 (0.009)	0.008 (0.009)	0.007 (0.009)
Age	-0.003*** (0.000)	-0.003*** (0.000)	-0.003*** (0.000)
Education (HS)	-0.091*** (0.013)	-0.091*** (0.013)	-0.090*** (0.013)
Education (HS +)	0.012 (0.015)	0.012 (0.015)	0.014 (0.015)
Work: Unemployed	0.201*** (0.026)	0.199*** (0.026)	0.198*** (0.026)
Work: Student/home/retired	0.113*** (0.012)	0.113*** (0.012)	0.112*** (0.012)
Family income: 2nd Quartile	-0.178*** (0.013)	-0.178*** (0.013)	-0.178*** (0.013)
Family income: 3rd Quartile	-0.237*** (0.014)	-0.238*** (0.014)	-0.238*** (0.014)
Family income: 4th Quartile	-0.258*** (0.015)	-0.258*** (0.015)	-0.260*** (0.015)
Race: Black	0.374*** (0.014)	0.374*** (0.014)	0.374*** (0.014)
Race: Hispanic/other	0.128*** (0.022)	0.127*** (0.022)	0.128*** (0.022)
Region: Mid Atlantic	-0.077** (0.024)	-0.079** (0.024)	-0.077** (0.024)
Region: EN Central	-0.054* (0.024)	-0.056* (0.024)	-0.055* (0.024)
Region: WN Central	-0.055* (0.027)	-0.056* (0.027)	-0.056* (0.027)
Region: S Atlantic	-0.093*** (0.024)	-0.095*** (0.024)	-0.094*** (0.024)
Region: ES Central	-0.104*** (0.028)	-0.105*** (0.028)	-0.104*** (0.028)
Region: WE Central	-0.101*** (0.026)	-0.102*** (0.026)	-0.100*** (0.026)
Region: Mountain	-0.055* (0.028)	-0.056* (0.028)	-0.056* (0.028)
Region: Pacific	-0.045 (0.025)	-0.047 (0.024)	-0.045 (0.024)
Constant	2.151*** (0.044)	-0.056 (0.354)	3.467*** (0.539)
N	26,290	26,290	26,290
N (groups)	27	27	27

Cells contain coefficients and standard errors from an GLS regression estimated with random effects for years. Ordered logit models are included in Appendix Table A2. *p < 0.05; **p < 0.01; ***p < 0.001.

of these micro-level factors does help account for differences across individuals and may help explain variation over time as well, to the extent that distributions of characteristics have changed over time. (See for instance Appendix Figure A1, which examines the racial and ethnic composition of the GSS over time.)

For our purposes, the central result in the first column of Table 1 is the statistically significant trend. The coefficient (0.011) reveals upward movement in absolute preferences (P^*) that is only implicit in Figure 1. It implies that the demand for welfare spending has increased in a fairly consistent way over time, in line with previous research (Wlezien 1995; Soroka and Wlezien 2010). The effect, while seemingly small, actually implies nearly half a point increase in the mean response over the 1976–2016 period, which is one quarter of the theoretical range of the variable, and is substantial in terms of spending, as the depictions that follow make clear. Note that we can translate this over-time effect of trend into spending amounts based on the coefficient for per capita spending (–0.055), which implies an increase in the underlying preferred level of about \$800 (2009) per capita, an amount that is larger than the mean level of spending during the period (roughly \$675).¹⁷ Having statistically documented this trend in spending preferences, let us see whether our economic variables account for it. The second column of Table 1 reports results that include unemployment and per capita GDP in the model. This allows us to begin to directly assess whether and to what extent macroeconomics account for the increasing demand for welfare spending over time.

¹⁷That is, the trend coefficient (0.011) implies an increase of 0.44 over the period (40 times 0.011) and the spending required to equalize this increase equals 0.44 divided by the coefficient of feedback (0.055) multiplied by \$100, as the spending units are in hundreds.

Notice first from the second column of [Table 1](#) that controlling for the two variables increases the magnitude of estimated thermostatic responsiveness by 25 per cent, from -0.055 to -0.068 . The impact of the two macroeconomic variables themselves actually is complex, however. Based on the positive, statistically significant coefficient (0.069), increases in per capita GDP are associated with increases in support for spending. This is a pro-cyclical effect, one that comports with most of the literature (see [Durr 1993](#); [Stevenson 2001](#); [Wlezien 1995](#)); that is, as the economy expands, public support for more government appears to increase.¹⁸ Based on the positive, statistically significant coefficient (0.062) increases in unemployment also are associated with increases in support for spending, a clear counter-cyclical effect, which fits with other research ([Erikson, MacKuen, and Stimson 2002](#)).¹⁹ Since the two variables tend to move together inversely over time, the effects partly contrast and to some degree cancel out. The impact of unemployment, while important, only dampens the positive effect of GDP, as the latter substantially outweighs that of the former.²⁰

Note also that the macroeconomic variables added in column 2 completely account for the positive trend we saw in column 1 of [Table 1](#).²¹ Indeed, in column 2 the linear trend variable has a negative, statistically significant effect (-0.026) on welfare spending preferences. This may surprise, but keep in mind that we have not included the Gini index, which has trended upward over time and has been shown in previous research to negatively influence support for redistribution, as noted above. When we add that variable to the equation, in the third column of [Table 1](#), it has the expected negative effect, indicating that as inequality has increased, net support for welfare spending has decreased, other things, especially the economy and spending itself, being equal.²² With the Gini index in the model, the estimated coefficient for the trend variable (-0.004) is not distinguishable from 0. That is, our model now accounts for much if not all of the underlying positive trend in the public's welfare spending preferences.²³

¹⁸There is an interesting related literature on economic conditions and postmaterialist values; see [Clarke and Dutt 1991](#); [Duch and Taylor 1993](#); [Inglehart and Abramson 1994](#).

¹⁹Note that the model does not contain a measure of prices, as the coefficient for the consumer price index (CPI) is not statistically significant ($b = 0.001$; $s.e. = 0.002$), and including it makes little difference to the other estimates included here and adds no explanatory power. Given this, and also its correlation with the other trending variables in the equation, we exclude the variable from the analyses reported in the text.

²⁰This is relatively easy to examine. Here is an example using the GLS coefficients for GDP and unemployment from the second model in [Table 1](#). A one-standard deviation increase in within-sample per capita GDP (7.245) is associated with an average 0.500 impact on *Net Support* (based on multiplying 7.245 by the coefficient for per capita GDP, 0.069). Over this time period, a one-standard deviation increase in per capita GDP is associated with an average -0.667 decrease in unemployment (based on multiplying 7.245 by -0.092 , the coefficient for per capita GDP from an OLS regression of unemployment on per capita GDP, using annual data). That 0.667 decrease in unemployment produces an average -0.041 shift in *Net Support* (-0.667 multiplied by the coefficient for unemployment in [Table 1](#), which is 0.062). Overall, then, the resulting shift in *Net Support* is $+0.459$ ($0.500 - 0.041$). The positive impact of GDP thus appears to overwhelm the negative impact of unemployment.

²¹Additional modeling suggests that it actually is GDP alone that accounts for the positive trend, which might help explain the trend found in other spending domains ([Soroka and Wlezien 2010](#)).

²²Note that when added to the model, the poverty rate has a positive but insignificant effect ($b = 0.011$, $s.e. = 0.010$) and the other estimates change trivially and insignificantly.

²³To further assess the significance of the trend estimate in the face of multicollinearity with per capita GDP, we decompose the latter into two components – one predicted by trend and the remaining, residual component – and then substitute these into the model. Re-estimating an otherwise identical regression to that in column 3 of [Table 1](#), though now excluding time, which is perfectly collinear with the predicted component of per capita GDP, produces a coefficient of 0.055 ($s.e. = 0.004$) for the predicted component and 0.062 ($s.e. = 0.008$) for the residual component, very similar estimates that are not significantly different statistically ($p = 0.47$). Also note that the null effect of trend is evident in annually aggregated preferences: the coefficients for GDP (0.060) and unemployment (0.046) are positive and statistically significant, the coefficient for Gini (-0.081) is negative and statistically significant, and the coefficient for trend (-0.003) is less than 0 and not statistically significant. Finally, note that the aggregate welfare spending preferences time series pretty clearly is stationary (see footnote 4), that spending and the Gini index appear to be integrated (with a trend), GDP seems trend stationary and unemployment stationary, and aggregations of certain respondent demographics also may be integrated. This implies that preferences may be the linear

What, in sum, do these results tell us about welfare spending preferences over the past forty years? First, there clearly has been an upward trend in spending preferences. Second, that trend reflects a combination of economic factors. Upward movement in per capita GDP has been the most important factor – increases in Americans' welfare spending preferences are in large part a pro-cyclical reaction to the expanding economy, which has trended substantially over time, alongside a somewhat weaker counter-cyclical reaction to unemployment. Were Americans responding to GDP (and unemployment) alone, however, support for welfare spending seemingly would be even greater. But, as we have seen, the positive trending impact of macroeconomics has been muted by growing income inequality, which also has trended upward over time and led people to be less supportive of redistribution, other things being equal. This seemingly perverse result is in line with work by Kelly and Enns (2010) on the self-reinforcing nature of income inequality preferences.²⁴ Taken together, the contrasting effects of trends in these two variables still is positive; that is, the effect of GDP overwhelms the effect of the Gini index. And, to be clear, the results imply that the effects of GDP and the Gini index owe not just to their upward trends over time but also to the variation around those lines.²⁵

Decomposing relative preferences

Just how much of the variance in support for spending is accounted for by spending or macroeconomics, income inequality and demographics? Answers to this question provide information about how and why the public's underlying preferences for welfare spending evolve over time. Table 2 shows the *R*-squareds for a stepwise series of models of relative preferences, following the specification in the final column of Table 1.

The first row shows results from a model of relative preferences that includes just demographic and regional variables. These variables account for rather little of the variation in relative preferences – about 23 per cent of the between-year variance, and 6.3 per cent of the within-year (individual-level) variance.²⁶ The explanatory power of demographics for between-year variance

combination of some integrated series – together with other stationary variables – much as is implied by the thermostatic model, in which case the equation would be balanced (see Enns and Wlezien 2017). Directly diagnosing the combination on the right-hand side is difficult given the numbers of time-series observations and potential variables, but see footnote 32. Importantly, estimating a first differences – and error correction – model of aggregate preferences produces coefficients for spending and the economic variables that are statistically significant and have the same signs as those from a regression in levels, and are similarly-sized. This implies that the results of our analysis are not dependent on our modeling assumptions.

²⁴Other variables might have produced the same effect, including trust in government, which trended downward over the period, though note that adding Hetherington and Rudolph's (2015) aggregate measure to the model produces a small and insignificant coefficient ($p = 0.24$) and only marginally (and insignificantly) changes the estimated effects of the Gini index and the other economic variables. There are strong correlations between the economic measures, however, and so there thus is reason for care in the interpretation of their estimated effects in Table 1. This is particularly true for the Gini coefficient and GDP, which trend upwards steadily over time. (The correlations between either of these two measures and a linear time trend are 0.97 and 0.98.) Excluding one macroeconomic variable thus may shift the coefficients for others, and so it is important to note that the finding that preferences are pro-cyclical (and effectively 'pro-inequality') does not change as variables are dropped or added to the model. The predictions of the model – critical to estimations in the sections that follow – also change only marginally with different combinations of macroeconomic variables. That said, the negative effect of the Gini evident in Table 1 holds only when other positively trending variables are in the model; indeed, the coefficient turns positive and is significantly different from 0 when both GDP and the linear trend variable are excluded. Importantly, the estimated effects of the economic variables seen in column 3 of Table 1 all are unchanged when they (and spending) are detrended – see Table A9 in the Online Appendix. Because detrended variables are used in that analysis, the net effect of trend in the different variables is directly evident from the estimate for the yearly counter variable, which is positive and statistically significant as a result.

²⁵This is clear since if *only* the upward trend mattered, we would not expect macroeconomics to matter independent of the trend variable; nor would we expect both per capita GDP and the Gini index to matter at the same time. Also see the results of analysis that detrends spending and the economic variables discussed in n. 24.

²⁶Including attitudinal variables does not take us much further in explaining the individual-level variance. For instance, we can estimate the models adding results from one welfare-related question asked consistently over the time period analyzed

Table 2. How much variation can we explain?

	Stepwise <i>R</i> -squared		
	Within	Between	Overall
Demographics and regions	0.063	0.229	0.065
+unemployment & GDP	0.063	0.269	0.072
+Gini	0.063	0.373	0.074
+Spending	0.063	0.923	0.088
(Spending w/o macroeconomics)	0.063	0.457	0.065

Based on random-effects GLS regressions. $N = 26,290$.

suggests that some of the over-time change in preferences is a product of demographic change. Given what we have seen in previous models, as the American population – or at least the survey sample – becomes more educated over time, expressed preferences for welfare spending should shift leftward. And, as noted above, changes in the racial and ethnic composition help account for the between-year variance (see Appendix Figure A1).²⁷

Subsequent rows in Table 2 show the proportion of variance explained by models that add macroeconomics, the Gini index and spending. There are no differences in the proportion of within-year variation explained here – after all, we are adding variables that only vary over time. But the proportion of between-year variance explained jumps dramatically when macroeconomic and spending variables are added to the model. It is important to note that *only together* do they have strong explanatory power. (See the drop in between-year variation explained when spending is included without macroeconomics, in the final row.) With both spending and macroeconomics included, however, the model accounts for just over 90 per cent of the between-year variance in relative preferences.²⁸ Indeed, it is not even clear that the variance that remains is real and not mere survey error.

Simulating absolute preferences

Clearly, spending, macroeconomics and income inequality all are important to relative preferences for welfare spending. This is exactly as we would expect if economics drive the public's underlying absolute preferences. The strength of the models in Table 1, at least where across-year variance is involved, also allows us to explore what aggregate absolute preferences might look like over time. That is, we are able to use the model of relative preferences in Table 1 to impute the underlying level of absolute preferences.

Our approach is as follows. We begin with the final model of relative preferences for spending in Table 1. We then use the coefficients from that model to calculate estimates of P^* by including various combinations of the predictors, as depicted in Figure 2. (Note that complete syntax for this, and all other analyses, is available through the Harvard Dataverse.) A first line shows the predicted value of P^* when we use just the coefficients for the demographic variables and their observed values over time. This results in a *slight* upward trend in the estimated P^* . The

here, 'Some people say that people get ahead by their own hard work; others say that lucky breaks or help from other people are more important ...' Including responses to this question increases the within-year variance explained in the 'Demographics and regions' model from 0.063 to 0.068. (Though note that we also lose nearly 40 per cent of the cases due to missing responses on this variable; when estimated for the same set of observations, the explained within-year variance goes from 0.063 to 0.066.)

²⁷Note that the impact of demographics on over-time change is also consistent with work on cohort effects in redistributive preferences. See Giuiliano and Spilimbergo (2014), Neundorf and Soroka (2018).

²⁸Note that the fact that spending and macroeconomics (and trend) explain more over-time variation than demographics is not surprising given how much these variables change over time. What strikes us here is that the aggregate-level variables account for *so much* of the variation.

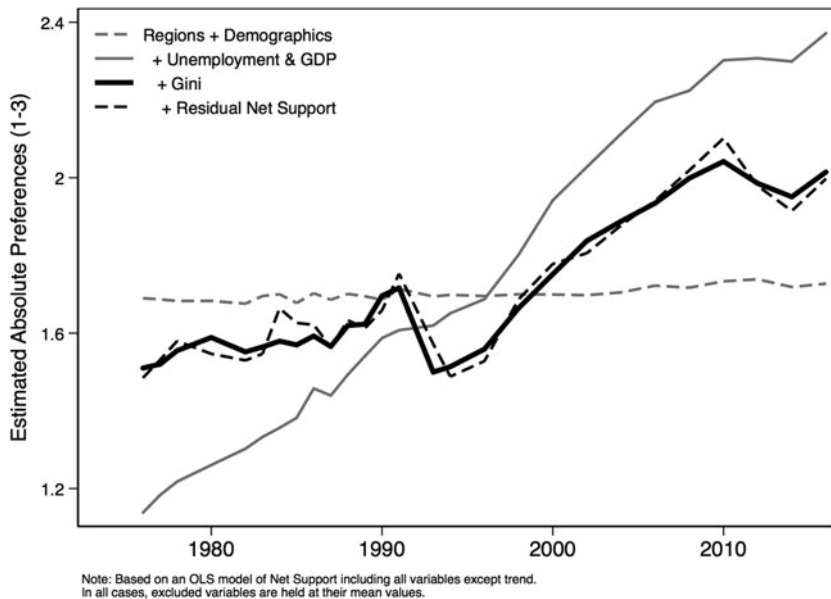


Figure 2. Estimated absolute preferences for welfare spending

real difference comes when we add the economic variables. Including per capita GDP and unemployment produces estimates for P^* that trend steadily upwards. Adding the Gini index dampens that upward trend.²⁹ The importance of macroeconomics and inequality in capturing over-time variance in P^* is made particularly clear when we estimate P^* when we also include the residuals from the model of R , that portion of welfare spending preferences for which our model does *not* account. There clearly is some additional variation there, particularly in the mid-1980s to mid-1990s, but the overall trend is not fundamentally different, usually indistinguishable. Indeed, the correlation between the two aggregate-level series is a 0.96. Most of the over-time variation in P^* clearly is accounted for by our economic variables.

Figure 3 shows the estimated P^* series (without the residual component) alongside the survey measure of R . Here we illustrate both relative and absolute preferences together, and we do so to make clear how different the two series are. Relative preferences move up and down over time, in response to the combination of P^* and P . P^* , on the other hand, trends upwards. The two series are correlated at 0.62, which implies that only a little more than a third of their variance is common. It makes clear how using measures of R to capture P^* , besides being theoretically tenuous, would pose empirical problems, not just for our welfare spending question, but potentially for a wide range of survey questions tapping ‘relative’ support for policy.

Estimating the ‘preferred’ level of spending

We know that the public’s *underlying* absolute preferences for welfare spending change and have increased over time. But how much spending satisfies the public at each point in time? We cannot say for sure based on our analyses, as we do not have a direct measure of P^* , at least not in dollars. We seemingly can infer it from our results, however – we can estimate the amount of spending that would, holding other things constant, lead to the average respondent preferring the spending

²⁹We do not make predictions for years in which data are not available, though drawing lines through the available data in Figure 3 (and in subsequent figures as well) effectively interpolates values for missing cases.

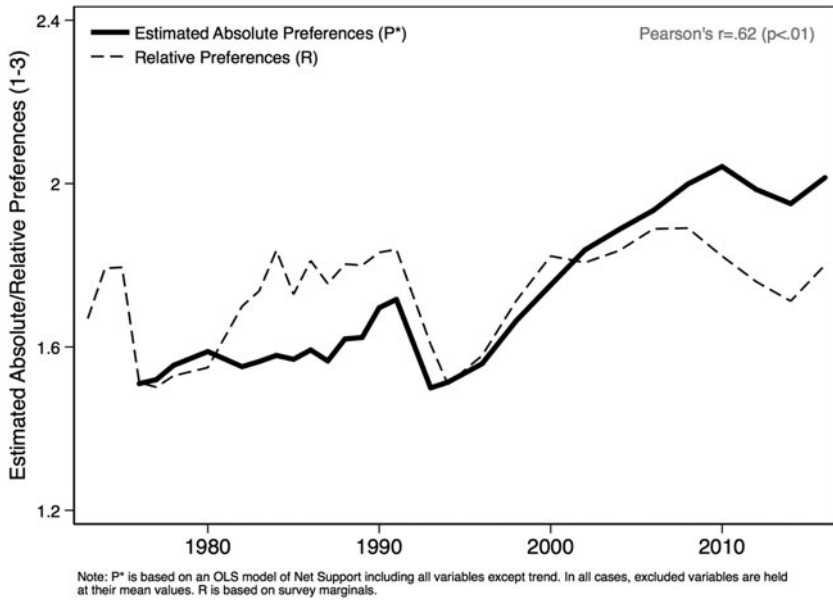


Figure 3. Relative and implied absolute preferences for welfare spending

status quo. This is the amount required to predict a mean spending preference of ‘2’, that is, ‘the same’ rather than ‘more’ (3) or ‘less’ (1), in each year. The approach is implied by the survey question itself, which asks about the congruence between the spending they want and the spending they get, and is the assumption underlying work on majoritarian ‘consistency’ (Brooks and Manza 2007; Gilens 2012; Monroe 1979; Petry 1999; Weissberg 1978) and provides the basis for Bartels’ (2015) assessment of the democratic deficit across countries.³⁰

The prediction of PS_1^* , the level of welfare spending that would produce a mean R_t of 2, is relatively straightforward.³¹ We begin with a simple time series model of R_t regressed on both spending (P_t) and the estimated absolute preferences (\hat{P}_t^*) series illustrated in Figure 3:

$$R_t = a + \beta_1 P_t + \beta_2 \hat{P}_t^* + e_t. \tag{5}$$

This is an aggregate level equation and is estimated using the annual values of net preferences (R , illustrated in the bottom panel of Figure 1), the annual values of P^* (illustrated in Figure 3) and annual spending figures (P).³² The coefficient for spending (β_1) indicates the shift in R that follows from the one-unit increase in spending (P); $1/\beta_1$ thus is the amount of spending required to

³⁰If people say ‘about right’, after all, they seemingly are happy with the status quo. By contrast, if they say ‘too little’ or ‘too much’, they are indicating that they prefer policy change. The assumption also is reflected in Ellis and Stimson’s (2012) characterization of ‘conflicted conservatives’ – those who think of themselves as conservatives but support liberal spending policies.

³¹This does assume that the model applies out of sample, as the mean value of R_t never does reach 2.0 in our data – see Figure 1.

³²Specifically, the dependent variable is the mean preference in each year. The resulting equation is $R_t = 0.526 + -0.072P_t + 0.982\hat{P}_t^* + e_t$, where all coefficients are significant at $p < 0.001$; $N = 27$. Here, the time series of preferences is the linear combination of two seemingly integrated but jointly stationary series, which is supported by an ADF test of residuals from a regression relating the right-hand side variables ($\rho - 1 = -.42, p < .10$); also see footnote 23. Note that, rather than using aggregated preferences, we can estimate the coefficients using an individual-level model, as in Table 1, and doing so produces essentially the same results, as it should.

produce a one-unit increase in R .³³ In each year, then, $P\$\ast_{1t}$ – the estimated level of spending that would be required for predicted values of $R_t(\widehat{R}_t)$ to be equal to 2 – is:

$$P\$\ast_{1t} = P_t + \frac{(2 - \widehat{R}_t)}{\beta_1} . \quad (6)$$

The results are shown in [Figure 4](#), which plots the predicted $P\$\ast_1$ (the thick solid line) alongside actual spending levels (the thin dotted line). We ignore the thick dashed line for the time being and discuss it below. In the figure we can see that predicted $P\$\ast_1$ closely parallels spending across the time period, with a correlation of 0.75. The variable is consistently lower than the actual level of spending, however.³⁴ This implies that politicians – both Democrat and Republican – always provide more welfare spending than the public wants. The pattern is surprising, particularly given research that implies that politicians are more responsive to the rich (such as [Bartels 2008](#); [Gilens 2012](#); [Gilens and Page 2014](#); but see [Branham, Soroka, and Wlezien 2017](#); [Brunner, Ross, and Washington, 2013](#); [Ellis 2017](#); [Enns and Wlezien 2011](#); [Rigby and Wright 2013](#)). If we take the neutral point from the response categories seriously, it looks like the poor, who favor the status quo on balance, are better represented on welfare (also see [Wlezien 2017](#)).

We cannot take the scale of relative preferences at face value, however. Question wording introduces problems. To begin with, there is the classic problem relating to policy labels, as the difference between support for spending on ‘welfare’ and ‘the poor’ differ dramatically, mostly in the form of an intercept ([Rasinski 1989](#)). There also is an assumption that the measures of spending match the public’s conception of the particular areas, for example, that the programs in a measure of welfare spending actually are the same as the programs that the public thinks of as welfare spending. In addition to the labels, the spending question registers people’s unconstrained preferences, with no trade-offs between spending on different programs, between spending and taxes, or between spending and deficits. [Hansen \(1998\)](#) has shown that taking these into account produces very different spending preference distributions, and greater support for the status quo.

What then is the correct neutral point? Instead of basing our estimate of $P\$\ast$ on the middle response category in the survey question, we can instead scale based on the mean, or long-term equilibrium, value of R .³⁵ The assumption makes sense if policymakers on average represent the median respondent over time, perhaps providing more than is wanted at some points in time and less than is wanted at others, but in a way that nets out. Of course, the assumption may not be true. That is, it may be that policymakers consistently represent the rich, in which case welfare spending could end up consistently below what the average person wants (see [Bartels 2008](#); [Gilens 2012](#); also see [Soroka and Wlezien 2008](#)).

This is less of a problem than one might think, and for three reasons. First, it is not clear that policymakers only or even mostly represent the rich ([Branham, Soroka, and Wlezien 2017](#); [Enns and Wlezien 2011](#)). Second, the middle and the rich have very similar welfare spending preferences, and so even if policymakers do represent the rich, welfare policy ends up pretty much in line with what the average person wants ([Enns and Wlezien 2011](#); [Soroka and Wlezien 2008](#)). Third, research ([Wlezien 2017](#)) on presidential elections indicates that the opinion average is

³³This is a mean effect, which presumably varies over the time period of our analysis owing to the level of spending (and GDP). One way to take this into account is to divide $1/\beta_1$ by the mean of spending (\bar{P}) to produce a standardized measure, that is, the proportional shift in spending that would produce a unit shift in R . The main effect of standardizing in this way is to flatten estimated $P\$\ast$, shifting values up early in the series and down later in the series. Here, we focus on $1/\beta_1$ using the unstandardized estimate.

³⁴Looking annually (and considering just those 27 years for which estimates of $P\$\ast$ are available), $P\$\ast_1$ is 42 per cent of actual spending, on average; mean per capita spending is \$644 (2009) and mean $P\$\ast_1$ is \$269.

³⁵Technically, the mean response (1.73) is just slightly greater than the equilibrium level (1.71) because of the case that is dropped in the course of estimating the latter.




Figure 4. Predicted absolute preferences and actual welfare spending

the effective neutral point for the public, that is, incumbents are punished to the extent the public mood drifts away – either above or below – the mean.³⁶ There thus is good reason to suppose the average welfare spending preference over time is the neutral point or close to it.

To produce these alternative estimates of P^* , we substitute the mean welfare spending preference (1.73) for the middling survey response value (2) in Equation 6. The resulting P^* series is depicted using a dark dashed line in Figure 4. Here we can see that, while the original estimate suggested that the government consistently spends *more* on welfare than the public wants, our alternative measure suggests otherwise: government now appears to spend about the right amount.³⁷ Of course, this is by construction and, while there is reason to think the assumption underlying it is correct, we cannot be sure. It thus is not absolutely clear that we have captured the level of spending that the average person wants – or at least does not want to change – at each point in time. It nevertheless does appear that we have captured the change in those preferences over time, which is nicely reflected in the parallelism of the two series of estimates in Figure 4.

Conclusions

Motivated in part by the burgeoning literature on the relationship between public opinion and policy, this article has explored the relationship between measures of relative and absolute preferences for policy. The existing literature does not always recognize the difference between relative and absolute preferences, and we explicitly set out the difference here, building on work that develops a thermostatic model of preferences. While it is much easier to directly measure relative preferences than specific absolute preferences using opinion surveys, we still *can* get at the latter, and have done so using proxies in analyses of relative preferences. Doing so in the case of welfare spending makes clear that economic variables, which are often used (if only implicitly) as proxies for absolute preferences, can tell us a lot about absolute preferences. Indeed, it is in large part due to macroeconomic variables that we are able to identify change in absolute preferences over time –

³⁶ Adopting a neutral point different from the mean performs worse the further the point is from the mean.

³⁷ The mean of this estimate of P^*_2 is \$6,424 per capita (2009) USD, just slightly higher than P .

both pro-cyclically with the expansion of GDP and, to a lesser extent, counter-cyclically with the variation in unemployment. This produces an upward trend in spending preferences that has been muted by responses to increasing income inequality, tapped by the Gini index, though the net effect of the economic variables still is positive. These variables – along with spending levels themselves – almost completely account for over-time variation in welfare spending preferences.³⁸

Our results thus provide a glimpse at the absolute preferences that underlie the relative preferences that we observe directly, and on which much recent work on representation relies. In short, absolute preferences for welfare (a) have been trending upwards (pro-cyclically) over the past forty years and (b) are roughly in line with actual spending on welfare. We can effectively monetize *changes* in those preferences over time; we are less confident estimating the underlying preferred *levels* of spending, however. For this, we would need to identify the public's true neutral point, which is debatable; although we think there are good reasons (discussed above) to suppose that it is close to the mean preference across years, we cannot be absolutely sure.

The approach we have suggested here can apply beyond welfare, and beyond the US, of course. Towards that end, it seems unlikely that economic variables will be equally useful across all other spending domains. They might work in some areas, such as education and health, and other proxies might work in other domains, including survey-based attitudinal measures, the effects of which can be directly modeled, that is, using a thermostatic framework. Clearly, much work remains, and hopefully we have provided useful guideposts.

Acknowledgements. Earlier versions of the article at the Annual Meetings of the American Political Science Association, Chicago, 2013, and the Southern Political Science Association, New Orleans, 2018, and also at Columbia University and Texas A&M University. We are thankful to John Bullock, Peter Enns, Bob Erikson, Anthony Fowler, Andy Gelman, Sunshine Hillygus, Stephen Jessee, Kathleen Knight, Jeffrey Lax, Yotam Margalit, Justin Phillips, Bob Shapiro, Alexander Tahk, Joe Ura, Guy Whitten, and Greg Wolf for comments, and to Marc Hetherington and Thomas Rudolph for data.

Supplementary material. Data replication files can be found at: <https://doi.org/10.7910/DVN/TYJOMG> and online appendices at: <https://doi.org/10.1017/S0007123419000103>

References

- Bartels LM** (1991) Constituency opinion and congressional policy making: the Reagan defense build-up. *American Political Science Review* **85** (2), 457–474.
- Bartels LM** (2008) *Unequal Democracy*. New York: Russell Sage.
- Bartels LM** (2015) The social welfare deficit: public opinion, policy responsiveness, and political inequality in affluent democracies. Unpublished manuscript.
- Bartle J, Dellepiane Avellaneda S and McGann A** (2018) Policy accommodation versus electoral turnover: policy representation in Britain, 1945–2015. *Journal of Public Policy*. Currently in First View.
- Bartle J, Dellepiane-Avellaneda S and Stimson J** (2011) The moving centre: preferences for government activity in Britain, 1950–2005. *British Journal of Political Science* **41** (2), 259–385.
- Branham JA, Soroka SN and Wlezien C** (2017) When do the rich win? *Political Science Quarterly* **132** (1), 43–62.
- Brooks C and Manza J** (2007) *Why Welfare States Persist: Public Opinion and the Future of Social Provision*. Chicago: University of Chicago Press.
- Brunner E, Ross SL and Washington E** (2013) Does less income mean less representation? *American Economic Journal: Economic Policy* **5** (2), 53–76.
- Burstein P** (2010) Public opinion, public policy, and democracy. In Leicht K and Jenkins C (eds), *Handbook of Politics*. New York: Springer.
- Clarke H and Dutt N** (1991) Measuring value change in western industrialized societies: the impact of unemployment. *American Political Science Review* **79**, 10–27.
- Duch R and Taylor M** (1993) Postmaterialism and the economic condition. *American Journal of Political Science* **37**, 747–779.
- Durr R** (1993) What moves policy sentiment? *American Political Science Review* **87** (1), 158–170.

³⁸Specifically, the variables account for over 90% of the variance in measured preferences, keeping in mind that some of the unexplained variance is survey error.

- Eichenberg RC and Stoll R** (2003) Representing defense: democratic control of the defense budget in the United States and Western Europe. *Journal of Conflict Resolution* 47 (4), 399–422.
- Ellis C** (2017) *Putting Inequality in Context*. Ann Arbor, MI: University of Michigan Press.
- Ellis C and Stimson JA** (2012) *Ideology in America*. Cambridge University Press.
- Enns P and Wlezien C**, eds (2011) *Who Gets Represented?* New York: Russell Sage Foundation.
- Enns P and Wlezien C** (2017) Understanding equation balance in time series regression. *The Political Methodologist* 24 (2), 2–13.
- Enns PK and Kellstedt PM** (2008) Policy mood and political sophistication: why everybody moves mood. *British Journal of Political Science* 38 (3), 433–454.
- Erikson RS and Tedin KL** (2015) *American Public Opinion*, 9th edn. New York: Routledge.
- Erikson RS, MacKuen MB and Stimson JA** (2002) *The Macro Polity*. Cambridge: Cambridge University Press.
- Funk CL** (2000) The dual influence of self-interest and societal interest in public opinion. *Political Research Quarterly* 53 (1), 37–62.
- Gilens M** (2012) *Affluence and Influence*. Princeton: Princeton University Press.
- Gilens M and Page B** (2014) Testing theories of American politics: elites, interest groups and average citizens. *Perspectives on Politics* 12, 564–581.
- Giuliano P and Spilimbergo A** (2014) Growing up in a recession. *Review of Economic Studies* 81 (2), 787–817.
- Hansen JM** (1998) Individuals, institutions, and public preferences over public finance. *American Political Science Review* 92 (3), 513–531.
- Hetherington M and Rudolph TJ** (2015) *Why Washington Won't Work*. Chicago: University of Chicago Press.
- Inglehart R and Abramson P** (1994) Economic security and value change. *American Political Science Review* 88, 336–354.
- Jennings W** (2009) The public thermostat, political responsiveness and error-correction: border control and asylum in Britain, 1994–2007. *British Journal of Political Science* 39, 847–870.
- Kelly NJ and Enns PK** (2010) Inequality and the dynamics of public opinion: the self-reinforcing link between economic inequality and mass preferences. *American Journal of Political Science* 54 (4), 855–870.
- Key VO** (1961) *Public Opinion and American Democracy*. New York: Knopf.
- Kreps DM** (1990) *A Course in Microeconomic Theory*. Princeton: Princeton University Press.
- Lax JR and Phillips JH** (2011) The democratic deficit in the states. *American Journal of Political Science* 56, 148–166.
- Margalit Y** (2013) Explaining social policy preferences: evidence from the great recession. *American Political Science Review* 107, 80–103.
- Miller WE et al.** (1999) *Policy Representation in Western Democracies*. Oxford: Oxford University Press.
- Monroe A** (1979) Consistency between constituency preferences and national policy decisions. *American Politics Quarterly* 12, 3–19.
- Monroe AD** (1998) Public opinion and public policy, 1980–1993. *Public Opinion Quarterly* 62 (1), 6–28.
- Musgrave RA** (1985) A brief history of fiscal doctrine. In Auerbach A and Feldstein M (eds), *Handbook of Public Economics, vol 1*. New York: North Holland, pp. 1–59.
- Naumann E, Buss C and Bahr J** (2015) How unemployment experience affects support for the welfare state: a real panel approach. *European Sociological Review* 32 (1), 81–92.
- Neundorf A and Soroka S** (2018) The origins of redistributive policy preferences: political socialization with and without a welfare state. *West European Politics* 41 (2), 400–427.
- Petry F** (1999) The Opinion-Policy Relationship in Canada. *The Journal of Politics* 61 (2), 540–550.
- Powell GB** (2011) Party polarization and the ideological congruence of governments. In Dalton R and Anderson C (eds), *Citizens, Context, and Choice*. Oxford: Oxford University Press.
- Rasinski K** (1989) The effect of question wording on support for government spending. *Public Opinion Quarterly* 58 (3), 388–394.
- Rigby E and Wright GC** (2013) Political parties and representation of the poor in the American states. *American Journal of Political Science* 57 (3), 552–565.
- Soroka SN and Wlezien C** (2008) On the limits to inequality in representation. *PS: Political Science and Politics* 41 (2), 319–327.
- Soroka SN and Wlezien C** (2010) *Degrees of Democracy*. New York: Cambridge University Press.
- Soroka S and Wlezien C** (2019) Replication Data for: Trends in Public Support for Welfare Spending: How the Economy Matters, <https://doi.org/10.7910/DVN/TYJOMG>, Harvard Dataverse, V1, UNF:6:0ATrQK5GG65pZ1S1kYAeEA==[fileUNF].
- Stevenson RT** (2001) The economy and policy mood: a fundamental dynamic of democratic politics? *American Journal of Political Science* 45 (3), 620–633.
- Stimson JA** (2004) *Tides of Consent: How Public Opinion Shapes American Politics*. New York: Cambridge University Press.
- Stimson J** (1991) *Public Opinion in America: Moods, Cycles, and Swings*. Boulder, CO: Westview Press.
- Ura J and Ellis C** (2012) Partisan moods: polarization and the dynamics of mass party preferences. *Journal of Politics* 74 (1), 277–291.
- Weissberg R** (1978) Collective vs. Dyadic representation in congress. *American Political Science Review* 72, 535–547.

- Wlezien C** (1995) The public as thermostat: dynamics of preferences for spending. *American Journal of Political Science* **39**, 981–1000.
- Wlezien C** (2004) Patterns of representation: dynamics of public preferences and policy. *Journal of Politics* **66**, 1–24.
- Wlezien C** (2017) Policy (mis)representation and the cost of ruling: US presidential elections in comparative perspective. *Comparative Political Studies* **50** (6), 711–738.
- Wlezien C and Soroka S** (2012) Political Institutions and the Opinion-Policy Link. *West European Politics* **35** (6), 1407–1432.