

Conjoint Experiments on Political Support for Governmental Spending Profiles

Ann-Kristin Kölln

KU Leuven

annkristin.kolln@kuleuven.be

Christopher Wlezien

University of Texas at Austin

Wlezien@austin.utexas.edu

Paper prepared for presentation at EPOP 2016, University of Kent, 9-11 September.

Abstract

Governmental budgets and the priorities they reflect are the subject of recurring political debate in modern democracies. Existing survey research has already shown that citizens hold distinct preferences on governmental spending when they indicate a preference for “more” or “less” spending in a particular domain. Responses to these relative spending questions reveal only limited information about people’s preferences, however. As the questions are relative, people’s answers obviously tell us little about how much spending they want in the different areas. Perhaps more importantly, given the difficulty of directly measuring absolute preferences, responses are unconstrained and do not explicitly take into account trade-offs people make between domains. We begin to redress these issues using conjoint experiments in a survey of Swedish citizens. Specifically, we investigate citizens’ latent preferences for precise governmental spending levels across four domains – defence, labour market, healthcare, education – as well as the potential relationships between these preferences. The results strongly suggest that respondents hold distinct preferences on particular levels of spending, and these tendencies correspond with general left-right partisan dispositions. In addition, we find clear connections between healthcare and defence spending preferences that appear to be about the total spending: Support for combinations of defence and healthcare spending increase as the total spending exceeds current spending levels. Our results also indicate that absolute spending preferences are strongly associated with relative preferences and that they are still picking up on a different aspect of citizens’ spending attitudes. These initial findings do not consider the effects of imposing spending constraints, tax considerations or explicit trade-off – something we will address in a next step.

Literature and contribution

Governmental spending is at the heart of much of politics. The question of how much money should be spent on economic, social, and security policy is the subject of party manifestoes, election campaigns, and public policy-making after Election Day. Previous research already provides insights into citizens' relative preferences for spending (e.g., Wlezien 1995; Soroka and Wlezien 2010). This work relies on survey questions asking citizens about the extent to which they want “more” or “less” spending in particular domains. However, real-world governmental spending decisions require choosing specific amounts of spending and involve multiple domains at once. Although responses to the relative spending questions will reflect isolated, general preferences, they tell us little about what specifically people want in the different areas and the trade-offs people would make between domains. This research attempts to redress this imbalance.

In the project we use conjoint experiments to reveal citizens' underlying preferences for particular governmental spending profiles (see e.g., Hainmueller et al. 2014). Conjoint experiments allow mimicking the governmental spending decision more closely than established relative survey questions because respondents are asked to rate entire spending profiles as opposed to single domains. The survey design incorporates the idea of a multidimensional choice and the trade-offs involved when spending more on one domain as opposed to another. Administered to a sample of 5,000 Swedish citizens we estimate a total of 20 causal effects across four spending domains – education, healthcare, defence, and labour market – simultaneously. As far as we can tell, this is the first such study in any country, so provides a critical exemplar and basis of comparison for studies in other countries. The results show the trade-offs citizens make and whether these are dominated by concerns about particular spending domains. Along the way, we also assess which spending profiles generate the highest levels of support and whether and how this variation is connected to important individual-level characteristics such as party preferences or relative spending preferences.

Previous research already provides important insights into citizen preferences for governmental spending from a variety of angles: their preferred change in governmental spending (e.g., Wlezien 1995); their preferences in trade-offs between spending, deficits and

taxes (e.g., Hansen 1998); their preferences in trade-offs between individual spending domains (e.g., Busemeyer et al. 2016) the coherence of mass spending preferences (e.g., Jacoby 1994); the micro-level determinants of spending preferences (e.g., Eismeier 1982; Rudolph and Evans 2005); the effect of issue framing on spending preferences (e.g., Jacoby 2000); or the causes and consequences of specific welfare spending preferences in various contexts (e.g., Gingrich and Ansell 2012; Rehm 2011). These studies all capture relative preferences for spending, which means that research has been limited so far to assessments of the extent to which citizens want more spending in a specific domain. The underlying assumption is that citizens could realistically only have relative and not absolute spending preferences. As a result, we do not know yet just how much spending citizens want – or will accept – in different domains. In addition, previous studies were unable to assess to what extent citizens are willing to make cuts in one domain to the benefit of another.

Design and data

The use of conjoint experiments can unveil latent preferences of respondents (see e.g. Franchino and Zucchini 2015; Hainmueller and Hopkins 2015; Hainmueller et al. 2014). We conducted a web-based conjoint experiment in Sweden to assess the structure, interdependence and covariates of citizens' preferences for governmental spending budgets with absolute spending levels.

Survey respondents were asked to compare two alternative spending profiles that were displayed on the screen, where each has randomly varied spending levels across a set of domains. We focus on four spending levels of would-be spending profiles that are all constructed around the proposed 2016 budget of the Swedish government, including education (60 billion, 65 billion, 70 billion, 75 billion and 80 billion), healthcare (60 billion, 65 billion, 70 billion, 75 billion and 80 billion), defence (40 billion, 45 billion, 50 billion, 55 billion, 60 billion), and labour market (70 billion, 75 billion, 80 billion, 85 billion, 90 billion).¹ In all domains the middle category represents the actual spending level of the budget. The close match between real-world spending and the experimental design yields the

¹ For more information on the 2016 budget of the Swedish government see <http://www.government.se/articles/2015/11/central-government-budget-for-2016-in-figures/> [last accessed 8 August 2016].

benefit of being realistic, closer to the choices the government and Parliament faces when proposing and voting on the annual budget. Table 1 summarises the spending levels per domain included in our experiment.

Table 1. Spending domains and levels.

Spending domain	Spending levels
	<i>billion SEK</i>
Education and academic research (utbildning och universitetsforskning)	60 65 70 75 80
Healthcare, medical care and social services (hälsovård, sjukvård och social omsorg)	60 65 70 75 80
Defence and contingency measures (försvar och samhällets krisberedskap)	40 45 50 55 60
Labour market and working life (arbetsmarknad och arbetsliv)	70 75 80 85 90

The introductory text to the experimental design reads as follows: “Each year the government needs to decide how to distribute money across the different issue domains. Please read the descriptions of the potential spending profiles carefully. For each domain the level of governmental funding is given. Then, please indicate which of the spending profiles you would personally prefer.” Respondents were asked to express their preference for one of the budget alternatives, as well as to rate their level of support per budget alternative on a scale from 1 (“should definitely not be adopted”) to 7 (“should definitely be adopted”). The English translation of the experimental design is as follows:

Each year the government needs to decide how to distribute money across the different issue domains.

Please read the descriptions of the potential spending profiles carefully. For each domain the level of governmental funding is given. Then, please indicate which of the spending profiles you would personally prefer.

Spending Domain	Spending Profile 1 <i>SEK in thousands</i>	Spending Profile 2 <i>SEK in thousands</i>
Education and academic research	70,000,000	75,000,000
Healthcare, medical care and social services	70,000,000	65,000,000
Defence and contingency measures	50,000,000	40,000,000
Labour market and working life	80,000,000	90,000,000

1. If you had to choose between them, which of these two spending profiles would you adopt, if you were in power?

- 1 – Spending Profile 1
- 2 – Spending Profile 2

2. On a scale from 1 to 7, where 1 indicates that the government should absolutely not adopt the spending profile and 7 indicates that the government should definitely adopt the spending profile, how would you rate Spending Profile 1?

<i>Absolutely not adopt</i>							<i>Definitely adopt</i>
1	2	3	4	5	6	7	

3. Using the same scale, how would you rate Spending Profile 2?

<i>Absolutely not adopt</i>							<i>Definitely adopt</i>
1	2	3	4	5	6	7	

In our experimental design individual spending levels are the experimental treatments, of which there are 20. Given the number of spending dimensions and spending levels, there are in total $5^4 = 625$ possible profiles. With a pairwise comparison of spending profiles per respondent the full enumeration of comparisons amounts to $\binom{625}{2} = 195,000$. Even when

surveying 195,000 individuals to obtain information on all possible combinations it would not suffice for the necessary statistical power. However, Hainmueller et al. (2014) show that a full enumeration of comparisons is not necessary and that instead a lot fewer respondents can be surveyed with sufficient statistical power as long as spending levels are randomly assigned to profiles.


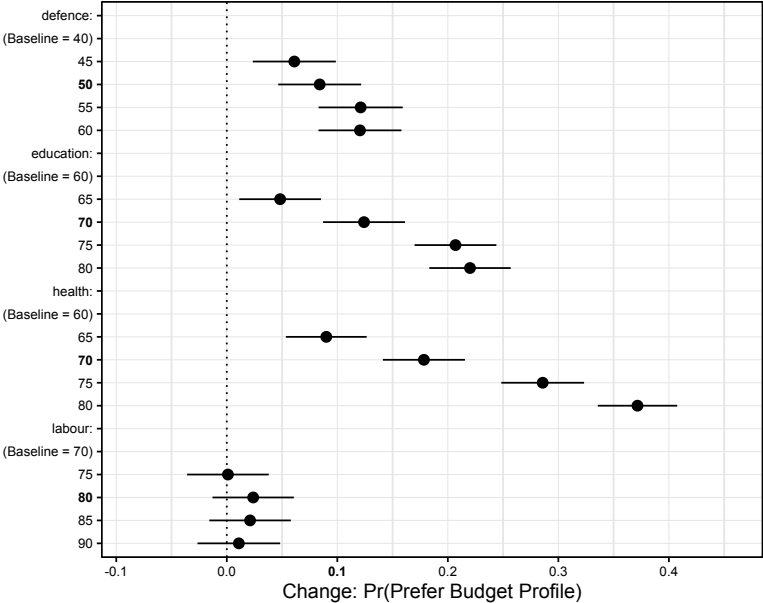
Our survey was administered to a sample of 5,000 Swedish respondents from a web-based survey conducted by the Laboratory of Opinion Research (LORE) at the University of Gothenburg in May/June 2016. Respondents are self-selected into the pool of respondents LORE has access to; the gross sample for our survey was pre-stratified according to age, gender and education. In those characteristics the gross sample is representative of Swedish residents aged 18-70. A total of 3,009 respondents participated in the entire experiment, which corresponds to a response rate of 60.2 per cent. All respondents were presented with one choice task, i.e., a comparison of two spending profiles. It means that our data cover more than 6,000 evaluated spending profiles. The order of spending domains displayed on the screen was randomised across respondents yet not within.

To estimate the effects we used the *cjoint* package in R (see Hainmueller et al. 2014) with clustered standard errors by respondent. Our main quantity of interest, the average marginal component effect, or AMCE, is obtained by regressing the dependent variable on dummy variables for each spending level, except for one per domain to facilitate comparisons with the baseline.

Results

To start we investigate to what extent there is a structure to citizens' preference for spending profiles – whether citizens hold preferences of specific levels of spending. For this we estimate the AMCE for each spending level across domains. The lowest spending level was set as the baseline in each domain. The upper panel in Figure 1 shows the results alongside the 95 per cent confidence intervals for each spending level when respondents were forced to make a decision between the two spending alternatives, while the lower panel reports the results when respondents rated the different profiles. We rescaled the ratings to range from 0 (“should definitely not be adopted”) to 1 (“should definitely be adopted”).

Figure 1. Effects of spending levels on approval of spending profiles.



Note: Spending levels printed in bold represent the current level of government spending in each domain.

The depicted AMCEs in the upper panel show the expected change in the probability that a profile is chosen when a given spending level is included by comparison with the domain baseline. For example, the results show in the upper panel that spending profiles with larger spending on healthcare almost linearly increase the likelihood of choosing a spending profile. Specifically, a spending profile with the highest spending level on healthcare of 80 billion SEK instead of the baseline amount of 60 billion SEK increases the probability that the spending profile is chosen by 0.37 (SE = 0.02), or 37 per cent. This is the strongest effect of all four domains. The results also show that profiles containing higher levels of spending on defence and education are associated with higher probabilities of support. However, in both cases the AMCEs for the two most expensive spending levels are almost the same, which indicates that respondents do not prefer maximal spending in those two domains, i.e., preferences are not monotonic. Finally, the findings also reveal that there is no initial structure to respondents' preferences on labour market spending. All effects associated with the four proposed spending levels are indistinguishable from the baseline level of 70 billion SEK. It suggests that respondents are indifferent to levels of labour market spending.


It is also important to note that current levels of government spending do not show the highest level of support in any of the spending domains. Instead, the results suggest that citizen support would be highest for a budget with 55 or 60 billion SEK for defence, 75 or 80 billion SEK for education, and 80 billion SEK for healthcare. Such a budget would be at least 20 billion SEK more expensive than the current one. This general preference for more expensive budgets is also clear from the fact that all effects, except those for labour market spending, are positive and thus indicate that any level of more spending compared to the lowest level per domain is preferable. It also highlights the importance of incorporating broader trade-offs between spending, taxes and deficits themselves, which we consider in the conclusions.

Overall, these findings illustrate citizens' preferences for governmental spending profiles follow a structure that is dominated by a lot higher spending on healthcare, some more spending on defence and education, while labour market spending preferences do not follow a specific structure. It means that citizens not only have relative spending preferences but also have preferences for different levels of spending in many domains, in our analysis all but

labour market spending. Notice also that the results do not change meaningfully when using ratings as the dependent variable, shown in the lower panel of Figure 2. That said, it is important to note that, despite almost identical numbers of cases, the results here are consistently less reliable and this is true for all of the analyses that follow, evidently the result of not forcing respondents to choose.

In a second step we investigate respondents’ party preferences as a potential covariate for spending profile approvals. It is one of the central findings of previous research that ideology or party preferences drive spending preferences. In this analysis we are interested to see whether this result holds when respondents are confronted with the multidimensionality of spending decisions and exact amounts. We used respondents’ party identification and recoded it into a dummy variable representing the two major political blocs in Swedish politics. Party identifications for the Left Party, the Social Democrats and the Green Party (plus the Feminist Initiative) were coded as “left-leaning” (= 0) and party identifications for the Moderates, the Liberals, the Centre Party, the Christian Democrats (plus the Sweden Democrats) were coded as “right-leaning.” We used this new variable and interacted it with the spending dimensions to obtain conditional AMCEs for those two subgroups of respondents.

Figure 2. Effects of spending profile levels on preferences for budgets by party identification (left panel = left-leaning respondents; right panel = right-leaning respondents).



The results in Figure 2 show the expected differences. In particular, spending profiles with higher spending on defence are favoured by right-leaning respondents, whereas spending profiles with higher spending on labour market are preferred by left-leaning respondents. Respondents belonging to the two groups represent quite clearly the often-observed differences in preferences between “guns” and “butter.” That said, profiles with higher spending on education and especially healthcare receive support from both groups, albeit much stronger from left-leaning respondents. For example, a spending profile that includes the highest possible spending level on healthcare as opposed to the lowest possible amount increases the probability of being chosen by left-leaning respondents by 0.54 (SE = 0.03). The corresponding effect for the group of right-leaning respondents is 0.23 (SE = 0.03), which is exceeded only by the effects of the two highest spending levels on defence. These results underscore the earlier finding that higher healthcare spending is seemingly universal importance for Swedish respondents’ choice of spending profiles. Note also that although preferences for labour market spending did not follow any particular structure in the unconditional model, the conditional results do indeed show a preference structure for more labour market spending amongst left-leaning respondents.

Table 2. Size of chosen budget by party bloc preference (percentages and absolute numbers).

		Party bloc preference		
		<i>left-leaning</i>	<i>right-leaning</i>	<i>Total</i>
Size of chosen budget	<i>smaller than or equal to alternative</i>	37.8 % (398)	43.8 % (414)	40.6 % (812)
	<i>larger than alternative</i>	62.2 % (656)	56.2 % (532)	59.4 % (1188)
	<i>Total</i>	100 % (1054)	100 % (946)	100 % (2000)

In addition, a common finding in comparative survey research is that left-leaning citizens prefer overall more governmental spending compared to right-leaning citizens. Our data also allow us to investigate the empirical basis for this in a multidimensional context that is closely mimicking the budget decision. We aggregated the spending levels of each profile that a respondent evaluated and created a variable that indicates whether the chosen profile was also

the more expensive profile. The cheapest budget a respondent could see amounted to 230 billion SEK whereas the most expensive budget cost 310 billion SEK (mean = 269.8; s.d. = 14.13). In 9.8 per cent of the cases the two budget alternatives were equally expensive. The univariate distribution shows that around 59 per cent of respondents chose the more expensive budget alternative over the less or equally expensive one. In combination with the above information on respondents' party bloc preference, Table 2 illustrates that left-leaning citizens were indeed more likely to prefer an expensive budget than right-leaning citizens. Around 62 per cent of left-leaning citizens opted for the more expensive budget while only 56 per cent of right-leaning citizens favoured the expensive profile. This is in line with the expectation that left-leaning citizens prefer higher governmental spending compared to right-leaning citizens.

One of the major advantages of conjoint experiments is the possibility to study how different attributes of the experiment interact. To what extent does the effect of one characteristic on the profile's support depend on the value of another? In the case of spending profiles it means that the interactions of different spending domains and levels allow us to see the extent to which respondents make trade-offs in their decision-making for a particular spending profile. Do they prefer more spending in one domain as long as it comes with less/more spending in another? We estimated the effects in four models, where each spending domain was interacted with each of the other three. It means that we estimated 48 interactions for each of the domains.

Table 3. Significant interactions of spending domain levels and their main effects.

	Attribute	Level	Estimate	Std. Err	z value	Pr(> z)	
<i>main effects</i>							
base = 40							
	defence	45	0.061	0.019	3.1763	0.001	**
	defence	50	0.082	0.019	4.242	0.000	***
	defence	55	0.122	0.019	6.280	0.000	***
	defence	60	0.121	0.019	6.295	0.000	***
base = 60							
	health	65	0.091	0.019	4.884	0.000	***
	health	70	0.180	0.019	9.472	0.000	***
	health	75	0.285	0.019	14.827	0.000	***

	health	80	0.373	0.019	20.290	0.000	***
<i>interactions</i>							
base defence:labour = 40:70; defence:health = 40:60							
defence (range: 40-60)	defence:labour	45:85	-0.143	0.061	-2.349	0.019	*
	defence:health	50:70	0.133	0.060	2.235	0.025	*
	defence:health	55:75	0.158	0.061	2.592	0.010	**
	defence:health	55:80	0.126	0.058	2.163	0.031	*
	defence:health	60:80	0.141	0.057	2.461	0.014	*
health (range: 60-80)	health:defence	70:50	0.140	0.060	2.343	0.019	*
	health:defence	70:60	0.118	0.059	2.005	0.045	*
	health:defence	75:55	0.162	0.061	2.661	0.008	**
	health:defence	80:55	0.124	0.058	2.144	0.032	*
	health:defence	80:60	0.143	0.057	2.497	0.013	*
Number of Obs. = 6168							
Number of Respondents = 3084							
Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05							

Table 3 only shows the results of the very few statistically significant interactions. As can be seen, it is predominantly the dimensions of healthcare and defence that interact, meaning that respondents' preferences for defence spending are dependent on the level of healthcare spending, and vice versa. It also means that to the extent respondents' have preferences for spending on education and labour market they are largely independent of other spending dimensions (except for the above reported interaction between defence and labour market on a single spending level).

Substantively, the coefficients on the interactions between defence and healthcare all indicate an increase in the average probability that a profile is chosen with those spending levels as opposed to a profile with the lowest possible spending on defence and healthcare (i.e., 40 and 60 billion SEK). The sign and size of the coefficients suggests that total spending is an important factor: support for combinations of defence and healthcare spending increase as the total spending exceeds current spending levels (120 billion SEK). For instance, a spending profile that puts 80 billion SEK on healthcare and 60 billion SEK on defence is associated with an increased probability of 0.14 (SE = 0.06) of being chosen as opposed to a profile that includes the two lowest categories on both domains. The effects are replicated with other

combinations, which increase support by 13-16 per cent. It means that respondents made implicitly or explicitly a positive connection between healthcare and defence spending as total spending increases.

Finally, we investigate to what extent respondents' preferences for spending levels and profiles matches their stated relative preference for higher spending in a domain. This not only serves as a validity test for the relative spending questions but also as a second test of possible trade-offs between spending domains. Before entering the experiment respondents of the survey were also asked about their relative spending preferences in the same domains. These questions closely resemble established survey items in the field on relative spending preferences. Respondents were asked in a grid "What is your opinion on the following proposals?" with the proposals of "to increase labour market spending," "to increase healthcare spending," "to decrease defence spending" and "to decrease educational spending" and the following response options: "very good idea," "good idea," "neither nor," "bad idea," "very bad idea." For the analysis we reversed the coding for the first two proposals, yet not the latter two, so answers to all questions measure the level of support for *increasing* spending in that particular domain.

Table 4 reports the correlations between relative spending preferences. Here we can see that all coefficients except for the correlation between education and defence spending are statistically significant at the .05 per cent level. The results show moderate to strong associations between welfare state domains, such as health and education spending ($r = 0.24$) as well as healthcare and labour market spending ($r = 0.45$). This matches previous findings in the field. Similarly, our data reproduce a negative relationship of medium strength between defence and labour market spending ($r = -0.27$) and a weaker relationship in the same direction between defence and health care spending ($r = -0.15$). Relative preferences for labour market spending and for education spending share a weaker positive relationship ($r = 0.16$).

Table 4. Correlations between relative spending preferences.

	more labour market	more healthcare	more defence	more education
more labour market	1			
more healthcare	0.45	1		
more defence	-0.27	-0.15	1	
more education	0.16	0.24	0.01 (n.s.)	1

Similar to the earlier analysis on respondent’s party identification, we interact each of the relative spending preferences with the spending domains to see the extent to which relative revealed preferences match absolute preferences in a multidimensional choice. The results are presented in Figure 4A-D.

Consider first only the interactions in matching spending domains across the four panels, which are always presented in the upper part of each panel. According to panel A, as respondents’ expressed relative preference for more defence spending increases across the five frames, their latent preferences for more defence spending also increases. For example, respondents who think it is a “very bad” idea to increase spending on defence show an increased probability of 0.25 (SE = 0.05) of rejecting a spending profile that includes the highest level of defence spending. Conversely, respondents who think it is “very good” to increase defence spending show an increased probably of 0.47 (SE = 0.03) to accept a spending profile that includes the highest level of defence spending. Similarly, close associations between expressed relative preference and latent absolute preferences can be observed for healthcare spending in panel B, for labour market spending in panel C and for education spending in panel D. The pattern of results suggests that latent preferences for absolute spending levels in a multidimensional spending are largely consistent with expressed relative preferences in isolated survey questions.

The results also inform us about the connections between spending domains citizens make when expressing preferences for a spending profile. Panel A, for example, shows a negative association between relative preferences for defence spending and latent preferences for healthcare spending, as indicated by decreasing effect sizes and standard errors from the left

most frame to the right most frame. It means that people who think that increasing defence spending is a good idea also express a latent preference for less spending on healthcare. This result is different from what we found above when interaction the absolute spending domains with each other. While the results above suggested a positive connection between healthcare and defence spending, the results here indicate a negative relationship and indeed a trade-off between relative preferences for defence spending and absolute preferences for healthcare spending. Beyond that, there does not seem to be any further interactions between relative spending preferences on defence and other latent preferences.


When moving to panel B and the interaction of relative spending preferences for healthcare with latent spending preferences for other spending areas, the results indicate if anything an association with preferences for education spending, yet not for defence spending. It suggests that the relationship between defence and healthcare spending preferences is unidirectional, where relative preferences for defence might be driving latent healthcare preferences but not vice versa. This is somewhat different from the results above where we found a positive relationship in both directions. It suggests that latent preferences for precise levels of spending might be picking up a different aspect of citizens' spending preferences than relative spending questions.

The results of panel C for labour market spending preferences suggest a weak negative association with latent preferences for defence spending as well as a positive relationship again with healthcare preferences. Finally, relative preferences for education spending do not indicate any significant interactions with the other latent spending preferences, as reported in panel D.


Figure 4. Effects of spending domain levels on approval of spending profile by relative spending preferences.

Note: Conditions 1 = “very bad proposal,” 2 = “bad proposal,” 3 = “neither nor,” 4 = “good proposal,” 5 = “very good proposal”.


A)




B)



C)



D)



Conclusion

This paper investigated for the first time the extent to which citizens hold preferences for absolute spending levels. The existing literature on spending preferences assumes, if only implicitly, that citizens could only possibly hold relative spending preferences since absolute levels of spending in the millions or billions are too abstract. We employed conjoint experiments in a Swedish setting to unveil citizens' latent preferences for absolute spending levels in four spending domains -- labour market, healthcare, defence, and education.

The findings show that citizens do indeed hold preferences for absolute spending levels when faced with a multidimensional decision to choose one spending profile over another. Interestingly, respondents do not show the highest level of support for the current budget of the Swedish government but would instead prefer a budget that is more expensive and puts a lot more emphasis on healthcare spending and some more on defence and education. In addition, we found that party identification is predictably associated with the size and the composition of the spending profile and that the connections that citizens make between spending domains are important but complex. The results suggest a positive connection between absolute spending preferences where support for combinations of defence and healthcare spending increase as the total spending exceeds current spending levels (120 billion SEK). However, the relationship between relative defence spending and absolute healthcare spending is not only negative but also unidirectional. At the same time, relative preferences for spending in specific domains correlate highly with the corresponding measures of absolute preferences, which suggests considerable validity of the measures but also that absolute spending questions are picking up something else than relative spending preferences.

With these findings we provide for the first time insights into the extent and patterns of citizens' preferences for absolute governmental spending. They provide new evidence that citizens do hold predictable preferences on absolute spending levels when faced with a multidimensional decision. That said, our analysis still focuses on what are largely unconstrained preferences, and in future research we will address the consequences of imposing real budget constraints and allowing for higher/lower taxes.

Bibliography

Busemeyer, M. R., Garritzmann, J. L. and Neimanns, E., 2016. Public Opinion towards Policy-Tradeoffs: Investigating Attitudes on Social Investment and Compensatory Welfare Policies with a Comparative Survey. *Working Paper*.

Eismeier, T.J., 1982. Public preferences about government spending: Partisan, social, and attitudinal sources of policy differences. *Political Behavior*, 4(2), pp.133–145.

Franchino, F. & Zucchini, F., 2015. Voting in a Multi-dimensional Space: A Conjoint Analysis Employing Valence and Ideology Attributes of Candidates. *Political Science Research and Methods*, 3(02), pp.221–241.

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

Hainmueller, J. & Hopkins, D.J., 2015. The Hidden American Immigration Consensus: A Conjoint Analysis of Attitudes toward Immigrants. *American Journal of Political Science*, 59(3), pp.529–548.

Hainmueller, J., Hopkins, D.J. & Yamamoto, T., 2014. Causal inference in conjoint analysis: Understanding multidimensional choices via stated preference experiments. *Political Analysis*, 22(1), pp.1–30.

Hansen, J.M., 1998. Individuals, Institutions, and Public Preferences over Public Finance. *American Political Science Review*, 92(3), pp.513–531.

Jacoby, W., 1994. Public Attitudes toward Government Spending. *American Journal of Political Science*, 38(2), pp.336–361.

Jacoby, W., 2000. Issue framing and public opinion on government spending. *American Journal of Political Science*, 44(4), pp.750–767.

Rehm, P., 2011. Social Policy by Popular Demand. *World Politics*, 63(2), pp.271–299.

Rudolph, T.J. & Evans, J., 2005. Political Trust, Ideology, and Public Support for Governmental Spending. *American Journal of Political Science*, 49(3), pp.660–671.

Soroka, S. & Wlezien, C., 2010. *Degrees of Democracy. Politics, Public Opinion, and Policy*. Cambridge: Cambridge University Press.

Wlezien, C., 1995. The Public as Thermostat: Dynamics of Preferences for Spending. *American Journal of Political Science*, 39(4), pp.981–1000.