Gamson's Law and voters' perceptions of portfolio allocation

NICK C. N. LIN, ¹ RANDOLPH STEVENSON, ² MATHIAS WESSEL TROMBORG² & DAVID FORTUNATO³

¹University of Mannheim, Germany; ²Rice University, Houston, TX, USA; ³Texas A&M University, College Station, TX, USA

Abstract. The assignment of ministerial portfolios to parties is one of the most contested and consequential processes in coalition politics. Accordingly, a great deal of scholarship has investigated how many portfolios different parties obtain in coalition negotiations as well as which parties are assigned which portfolios. However, to our knowledge, no one has ever examined how voters perceive the outcomes of this process – perceptions which must be fundamental to any assessment of policy responsibility in systems with coalition government. This article uses original survey data from four Western European countries to examine voter perceptions of the distribution of cabinet portfolios across parties. In addition to describing the extent to which voters know this distribution, the article also examines whether their perceptions are consistent with a number of different heuristics that voters might use to infer characteristics of the cabinet portfolio distribution. The results suggest that many voters use party role and size heuristics to infer the number of portfolios allocated to different parties hold which portfolios, but also that a significant number of voters have direct knowledge (not inferred using heuristics) of which parties hold which parties hold which ministries.

Keywords: Gamson's Law; portfolio allocation; voters' perceptions

Introduction

Holding elected leaders accountable for policy outcomes requires relatively well-formed and accurate perceptions of the relative distribution of policy-making power across political parties. In parliamentary democracies, policy-making power is tied inextricably to a party's relative influence within the cabinet, which may itself be a function of both its overall share of ministerial portfolios and the specific government ministries it controls. Indeed, theories of performance voting in coalitional systems (i.e., multiparty parliamentary democracies where coalition cabinets are the norm) often assume that voters know the distribution of portfolios over parties and use this information to allocate policy-making responsibility among coalition parties (e.g., Duch & Stevenson 2008) – or suggest that voters can anticipate the distribution of portfolios for many different potential coalition governments that have not been realised yet, and factor this information into their vote choices (e.g., Duch et al. 2010; Meffert & Gschwend 2010).

To date, however, no study has investigated voters' perceptions of (or expectations about) the distribution of portfolios among parties in coalition governments. Do voters directly learn the true distribution of cabinet portfolios over parties from, for example, media reports? Do they instead use simple heuristics to form their perceptions and expectations in the absence of such information? If so, do these heuristics conform to the strong real-world empirical regularities that appear to govern real portfolio allocations (i.e., Gamson's Law)? Or, are voters' perceptions and expectations systematically flawed – and, if so, do these flaws reveal anything about how they may perceive the relative policy-making influence of parties more generally? Are their perceptions simply whimsical or random? And, is there evidence that individuals with different levels of political interest or sophistication form their perceptions differently?

In this article we present new data that allow us to answer these salient open questions. We rely on the heuristics literature and the well-known theoretical and empirical literature on cabinet formation to articulate a set of expectations about how voters might form their beliefs about the distribution of portfolios over parties. The coalition bargaining literature is useful because it tells us which variables determine the true distribution of portfolios in coalition governments; whereas, the heuristics literature provides us with a realistic model for how voters may use the enduring empirical regularities produced by the process of coalition formation to make educated guesses about portfolio distributions that may be quite costly to observe directly. More specifically, we examine the possibility that voters use various heuristic rules to infer the relative share of cabinet seats (quantitative portfolio allocation) and which specific ministries each party holds (qualitative portfolio allocation). With respect to portfolio shares, these include heuristics based on Gamson's Law (i.e., voters assume that parties' portfolio shares are proportional to the seat shares that they bring to the government), seats shares more generally, and the special role of the prime minister. With respect to the allocation of specific ministries to parties, the heuristics include a simple 'importance rule' in which voters assume the most important ministries are assigned to the prime minister's party, a 'topical affinity' rule in which voters assume that parties are assigned ministries consistent with the policy topics that are most central to their 'brands' (e.g., environmental parties are allocated environmental ministries) and a 'historical regularity rule' in which voters expect parties to receive ministries that they have held often in the past.

To evaluate these possibilities for the quantitative allocation of portfolios (i.e., each party's portfolio share), we employ original survey data from Denmark, Germany, Italy and the United Kingdom; while for the qualitative allocation of portfolios (i.e., which parties get which portfolios), we use original data from Denmark. The results show that there is a strong correlation between perceived party seat share and portfolio share as implied by Gamson's Law, but also that prime ministerial parties are attributed a disproportionate share of portfolios that is larger than what is suggested by Gamson's Law or what could realistically be achieved in cabinet bargaining. Likewise, we find that parties perceived to hold the prime ministry are typically attributed the portfolios that are traditionally considered most important, but also that perceptions about the qualitative portfolio distribution depend on the topical affinity of party brands and the histories of parties holding particular ministries. Finally, we find evidence that some voters also have direct knowledge, not inferred using heuristics, about the true quantitative and qualitative portfolio distributions.

Voters, heuristics and coalition politics

For decades, the prevailing wisdom among students of comparative political behaviour was that coalition formation and policy making were simply too complex for voters to understand. Consequently, canonical research on both prospective voting (Downs 1957) and retrospective voting (Powell & Whitten 1993) concluded that voters are unlikely to factor post-electoral bargaining into their evaluations of parties on election day. This pessimistic conclusion has since been challenged (e.g., Duch et al. 2010; Kedar 2005; Meffert & Gschwend 2010), and the discipline now generally accepts that voters possess at least an elementary understanding of post-electoral bargaining, which is evident in their voting behaviour. Because of this, political scientists are now no longer asking 'Do voters understand coalition politics?', but rather 'How do voters understand coalition politics?' (Duch et al. 2015; Fortunato & Stevenson 2013).

The answer that has been offered to this question is that voters can make some sense of complex political contexts through the use of simple heuristics. These heuristics are typically rooted in the enduring empirical regularities of the political system and synthesise a small set of low-cost informational inputs into more complex, relatively accurate cognitions. This generalised heuristics model has been used to study how voters update their perceptions of political party ideologies (Fortunato & Stevenson 2013) and how they allocate responsibility for collective decisions (Duch et al. 2015). Following this recent work, we explore whether citizens are similarly capable of using simple rules to form relatively accurate perceptions about the distribution of portfolios in their polity.

Before we move forward, it is important to note that our attention to the heuristics voters might use to make inferences about portfolio allocation does not rule out the possibility that many of them do not use such heuristics, either because they obtain information about the distribution of cabinet portfolios directly (e.g., from the media) or because they do not even attempt to obtain such information (i.e., to infer it from a heuristic). If the former is true, we should see that the true distribution of portfolios influences voter perceptions even after accounting for heuristic rules; if the latter is true we should simply see inaccurate and unstructured answers on our survey questions, including large numbers of 'don't know' responses.

In the rest of this article, we examine the nature of voter perceptions of parties' relative shares of cabinet portfolios and their perceptions of which parties receive which portfolios. In each case, we review the theoretical literature to identify several candidate heuristics voters might use and evaluate the evidence for these against the possibility that voters do not use such heuristics, either because they obtain knowledge of portfolios directly or not at all.

Perceptions of the portfolio parties control

The modern literature on heuristics has two distinct faces, both of which are relevant here. On one side, there is an older literature that emphasises the role that heuristics play in producing error, misperceptions, consistent recurring mistakes and 'irrationality' (e.g., Tversky & Kahneman 1974). On the other side is the more recent approach championed by Gigerenzer and the ABC Group at Max Planck Berlin, in which individuals rationally use heuristics as efficient and generally accurate ways to make sense of complex situations.¹ In the latter argument, the heuristics that are most often applied are those that are 'ecologically rational' in a given context – that is, they are cheap (requiring easily acquired information as inputs), simple (easily applied to a given situation) and accurate on average over time and

across populations (Fortunato & Stevenson forthcoming). Empirically, there is evidence that both kinds of heuristics exist. Tversky and Kahneman's famous experiments make it clear that there are biases in perceptions that are real and persistent. Likewise, the ever expanding empirical case for the ecological rationality of many heuristics has shown that heuristic use is often conditioned on exactly the kinds of contextual variables that Gigerenzer and the ABC Group suggest – that is, those that regulate the expense, simplicity and accuracy of a given heuristic. Our purpose here is not to settle, or even contribute to, the ongoing debate between the rationality and irrationality of heuristics in general, but simply to test a set of heuristics that individuals might plausibly use to make inferences about the distribution of cabinet portfolios over parties. Thus, we do not take an *a priori* theoretical position about whether or not we expect one or the other of these kinds of heuristics to be used – though we do clearly point out whether such use would likely lead to accurate or inaccurate inferences.

Party seat share

Gamson's (1961: 376) general observation that '[a]ny participant will expect others to demand from a coalition a share of the payoff proportional to the amount of resources which they contribute to a coalition', as applied to the distribution of cabinet portfolios in multiparty government in which parties tend to receive a share of cabinet portfolios proportionate to the share of legislative seats they bring to the cabinet, is one of the most reliable and well-known empirical regularities in comparative politics. Though scholars have sometimes questioned whether Gamson's Law is truly as robust as commonly thought (e.g., Bäck et al. 2009; Falcó-Gimeno & Indridason 2013; Warwick & Druckman 2001), the fact remains that the naïve correlation between the relative seat shares of cabinet partners and their portfolio shares is close to one across many democracies and many years (Bassi 2013; Browne & Frendreis 1980). This correlation suggests, according to Gamson, that parties ascribe to a 'norm of fairness', though subsequent work suggests that parties need not seek 'fairness' in order arrive at a proportional allocation of ministries (e.g., Bassi 2013).

This empirical regularity could drive a norm among voters to expect, or infer, that parties' portfolio shares will reflect their legislative seat shares. The informational inputs to this rule are certainly not trivial, though perhaps less costly to obtain than directly observing portfolio shares.² Further, there is evidence that in at least some contexts voters are reasonably informed about both the composition of the cabinet and the sizes of the parties (Fortunato et al. 2015), though this varies contextually (e.g., voters in Italy and the Netherlands know substantially less about party sizes than those in Germany and Denmark). Thus, the first hypothesis we will examine in our empirical work is:

Gamson's Law hypothesis: Voters' perceptions of a party's portfolio share will be proportional to the legislative seat share it contributes to a cabinet.

One caveat to this hypothesis, however, comes from a closer examination of the empirical work on Gamson's Law (e.g., Browne & Franklin 1973; Browne & Frendreis 1980; Schofield & Laver 1985). While this work has certainly found a strong and robust relationship between legislative and cabinet seat shares, it has also consistently found that these shares deviate systematically from proportionality. Specifically, this work finds that smaller parties get a small, but significant, portfolio bonus at the expense of larger parties. For example, in

one of the most comprehensive and recent tests, Warwick and Druckman (2006) find that a regression between legislative and cabinet seat shares produced a slope coefficient of only 0.793 and a y-intercept of 0.069, implying that parties that contribute less than 33 per cent of the cabinets seats get a cabinet seat bonus. This robust empirical finding leads us to also examine the possibility that voters' key on the actual empirical relationship between portfolios and seat shares rather than the theoretical one argued by Gamson. While few would argue that voters consciously understand Gamson's Law, modern students of heuristics argue that individuals subconsciously use such strong empirical relationships to shape relevant perceptions and inferences.³ Thus our second hypothesis is that:

Gamson's Law hypothesis with small party bonus: Voters' perceptions of a party's portfolio share will be less than proportional to the legislative seat share it contributes to a cabinet, with a bonus for small parties.

One strike against the idea that voters will use Gamson's Law as a heuristic, at least if one ascribes to Gigerenzer's (2010) view that only ecologically rational heuristics will be used, is that it may be fairly difficult for a typical voter to apply the rule, even if they have information about legislative seat shares. The problem is that the rule does not actually map legislative seat shares to portfolio shares, but rather standardised shares of legislative seats, where the standardisation is over the set of parties in the cabinet. Thus, for voters to use the rule, they must first translate the shares of legislative seats that are widely reported on election night to the shares of seats the parties contribute to the cabinet. For example, suppose the voter knows that the Social Democrats control 35 per cent of the seats in the legislature, the Christian Democrats 45 per cent and the Liberals 20 per cent, and that the Social Democrats and Liberals have entered cabinet. To apply the Gamson's Law heuristic, they need to calculate that the Social Democrats contribute 64 per cent of the legislative seats in the cabinet [35%/(20%+35%)] while the Liberals contribute 36 per cent [20%/(20%+35%)]. It may be implausible that voters would include such a calculation, even approximately, in a rule intended to facilitate quick and easy inferences about portfolio shares.

The role of the prime minister

Another set of plausible heuristics that individuals might use to infer the distribution of cabinet seats over parties ignores information about the relative sizes of parties, and the identity of cabinet partners, and instead uses only the identity of the prime ministerial party as a very inexpensive and widely available input. One such heuristic, which we call the 'prime ministerial bonus' rule, assigns a disproportionate share of portfolios to the prime ministerial party. Indeed, in its most extreme version, the rule simply assigns all portfolios to the prime minister, which is, by definition, incorrect for all coalition cabinets, but in some cases would not be wildly inaccurate.

There are two different strains of work that lead us to examine this candidate heuristic. The first case draws from (1) the voluminous empirical work on cabinet formation that shows that the *formateur* party – the party which leads, and controls the agenda of, the cabinet formation process – almost always provides the eventual prime minister; and (2) theoretical work that suggests this agenda control should allow such parties to shape the cabinet, at least to some degree, to their liking (e.g., Baron & Ferejohn 1989; Snyder et al. 2005). Thus, a voter

who is not aware of Gamson's Law, as surely many voters are not, may well assume prime ministers or their *formateur* precursors get their way in cabinet negotiations. Of course, this is a case in which voters who use such an assumption will clearly be wrong empirically – at least with respect to the number of cabinet seats. Thus, modern theories of heuristic use would discount the long-term accuracy of such a heuristic, and thus its frequency of use.

The second reason that voters might assume prime ministers control a disproportionate share of cabinet seats stems from a very different logic. This reason does not have to do with their agenda control during cabinet negotiations, but rather the privileged position that prime ministers hold in policy making after cabinets form. Specifically, Duch et al. (2015) have shown that when individuals attribute policy-making influence for collective decisions, they often attribute a disproportionate share of that influence to decision makers with agenda control. This experimental finding has also been supported in subsequent observational research (Tromborg et al. 2015). And, of course, these findings are consistent with a great deal of scholarship in comparative politics that finds prime ministerial parties do dominate the legislative agenda, even in contexts where their formal powers over the agenda are comparatively weak. Indeed, several empirical examinations of lawmaking in parliamentary democracies suggest that prime ministers nearly always get what they want (Cheibub et al. 2004; Saiegh 2009).

Given this, it is possible that voters simply assume, or apply a useful heuristic that suggests, that prime ministers are disproportionately responsible for policy outcomes and subsequently assume that a party with this kind of policy-making influence also must control a large number of portfolios. From the point of view of the modern theory of the ecological rationality of heuristics, a prime ministerial bonus rule might be used, despite its inaccuracy in the distribution of actual portfolios, because of its consistency with an arguably more important (to voters) construct: prime ministerial influence, not to mention its relatively low informational and cognitive costs.

Prime minister bonus hypothesis: Voters assign more portfolio shares to the prime ministerial party than to junior member parties in cabinet.

It is, however, also possible that voters use the same set of informational cues about party roles in the government to generate completely different predictions about the allocation of portfolios. Specifically, Indridason (2011) argues that if the prime ministerial party foresees the threat of government dissolution, it may be willing to offer its partner(s) more portfolios in order to keep the government alive. He further shows that there is a systematic formateur penalty in portfolio allocation (but see Ansolabehere et al. 2005; Warwick & Druckman 2006). The underlying logic here is also consistent with Browne and Franklin's (1973: 461–463) explanation of small party bias in portfolio payoffs. The dominant party, most likely that of the prime minister, may be motivated to surrender several cabinet posts to its coalition partner(s) in exchange for necessary support. This 'prime ministerial penalty heuristic' is likely to be used from the perspective of the modern theory of heuristics as it uses the same inexpensive, informational inputs as the prime minister bonus rule, is equally simple and is more accurate – at least for inferences about the number of portfolios, if not for inferences about influence – in the long term.

Prime ministerial penalty hypothesis: Voters assign less portfolio shares to the prime ministerial party than that to junior member parties in cabinet.

An important aspect of our empirical testing will be accounting for the presence of true information in our estimates of heuristic application. Some voters may have direct knowledge of the true quantitative distribution of portfolios across cabinet parties, and possibly the true qualitative distribution as well, and therefore do not require any heuristics to generate perceptions of the distribution. Such information would likely come directly from media reports of the outcome of coalition negotiations, such as 'The Social Democrats capture 8 of 15 cabinet seats ...', as well as subsequent reminders in the course of normal political reporting. We account for this possibility by simply including in our models the true value of the target of interest: either a party's true portfolio share or the identity of the party holding a particular portfolio. Further, by structuring the models in this way, we can evaluate which voters (e.g., more or less sophisticated) are more likely to have some true information on the distribution of portfolios and which are more likely to rely on heuristics.⁴

Data and variables

In order to examine our hypotheses empirically, we conducted four Internet surveys on representative samples of adults in Denmark, Germany, Italy and the United Kingdom. Of course, there is a clear tradeoff in evaluating these hypotheses observationally, rather than experimentally. On the one hand, observational data do not allow us to identify the application of a particular heuristic in a controlled setting (e.g., Lau & Redlawsk 2001). On the other hand, these observational data do allow us to simultaneously test for the presence of true information on the real distribution of portfolios while assessing the heuristic roots of voter perceptions (e.g., Ansolabehere & Jones 2010). The survey in Denmark was fielded in July 2014, which was three years after the 2011 general election, and six months after the *Socialistisk Folkeparti* withdrew from the coalition government, but remained an external legislative support party, keeping the government alive. The German survey was launched in August 2014, which was about nine months after the 2013 general election. The Italian survey was also conducted in August 2014, about 15 months after its last general election in February 2013. The survey in the United Kingdom was administered in May 2012, two years after the 2010 general election.⁵

The potential heuristics we discussed above together suggest two specific informational inputs that may shape voters' perceptions of parties' portfolio shares: status in government (i.e., prime ministerial party or junior partner) and legislative seat shares. Consequently, we designed specific questions in the surveys in order to capture respondents' perceptions of these party characteristics.

First, we asked respondents which parties they thought were in the cabinet. Specifically, 'Please choose the option which best describes each party's current role in the government.' The question was followed by a list of all the parties and allowed four possible responses: (1) Party of the current prime minister; (2) Party is in the current cabinet but is not the party of the prime minister; (3) Party is currently in the opposition; and (4) Party has no seats in the parliament, or lower house. Since we are concerned with the perceived allocation of ministerial portfolios and parties in opposition cannot, by definition, be allocated portfolios,

our analysis excludes all parties that the respondent identified as not in the cabinet. Thus, our unit of analysis within each country is a respondent-party, but where the set of parties for each respondent may differ depending on which parties they identified as being in the government. Next, we created a dummy variable that indicates whether each perceived cabinet party is also *perceived* to be that of the prime minister ('1' for parties that are perceived to be the party of the prime minister, '0' otherwise).

In order to measure our dependent variable – voters' perceptions of each party's share of cabinet seats – we asked survey respondents the following question: 'What percentage of the cabinet ministries do you think each of the following parties holds?' Respondents were allowed to give any answer between '0' and '100' for each party and there was no constraint that these shares sum to 100. Nevertheless, 58 per cent of respondents gave answers that did sum exactly to 100, while 70 per cent summed to within 20 per cent of 100.⁶ In order to facilitate comparisons of this variable to the shares of legislative seats that each perceived cabinet party contributed to the cabinet, which is necessary to test the Gamson's Law hypothesis, we examine shares in the raw form described above as well as in a rescaled form. In the later version, we divided the perceived share of portfolios of each party that the respondent perceived to be in cabinet by the sum of the shares over all the parties they perceived to be in cabinet.

We also asked respondents to report their perceptions of parties' legislative seat shares. Specifically, we asked: 'What percentage of seats in the [name of national parliament] do you think each of the following political parties currently holds?' Respondents were allowed to give any answer between '0' and '100' for each party and, again, we did not impose a constraint that these answers sum to 100. In this case, 39 per cent of respondents gave answers that did sum exactly to 100, while 67 per cent summed to within 20 per cent of 100. As with portfolio allocations, we also created a rescaled version of this variable in which we divided each perceived cabinet party's perceived legislative seat share by the sum of the shares the respondent allocated to all his or her perceived cabinet parties. This variable is then directly comparable to the rescaled portfolio shares described above and it is this comparison that is, in our view, the fairest test of the Gamson's Law hypothesis. With this rescaling, the Gamson's Law hypothesis is that the covariate parameter estimate will be 1, and the intercept 0, if we regress our dependent variable on perceived seat share in government.

The accuracy of perceptions of portfolio shares

We begin our empirical analysis with a description of the distribution of values on the dependent variable. This description speaks directly to the question of how accurately respondents perceive the distribution of cabinet portfolios across parties. As we mentioned above, theories of coalitional voting have often assumed that voters have perfect knowledge about this distribution (e.g., Duch et al. 2010). However, when examining the shares that voters actually report, we see that the story is mixed: there is both broad qualitative support for the assumption and reasons to be cautious.

To facilitate understanding of the data, we begin with a subset of the data that is the most straightforward to examine and then move on to more complex analyses. Figure 1 shows the distribution of perceived (rescaled) portfolio shares for those respondents who identified the



Figure 1. Respondents' perceptions of portfolio shares for perceived cabinet parties (respondents who correctly identified the cabinet).

Note: Solid lines indicate the mean of perceived portfolio shares and dashed lines represent the true shares.

cabinet composition correctly (and allocated portfolios only to those parties). Solid vertical lines indicate the average response and dashed lines show true shares. A wider variance in the distribution and a greater distance of the perceived average share from the true portfolio share indicate that voters are collectively less accurate about a party's portfolio share.⁷

We can immediately see from this graph that perceptions among the 72 per cent of respondents who only allocated cabinet seats to parties actually in the cabinet are quite accurate overall. In three of the cases (Germany, Italy and Denmark) the average perceived portfolio shares (again, for those who identified the cabinet parties correctly) deviate from the true shares by no more than 6 percentage points. That said, in each of the cases these respondents under-allocate portfolios to the prime ministerial party, and by definition over-allocate an equal amount to partners, compared to the true shares and, for Britain, this under-allocation is fairly large at 12 per cent. The graph also shows that while average perceptions are generally accurate for this sample, there is substantial inter-respondent variation.

Figure 2 is similar to Figure 1 but includes all the parties and all respondents, even those who allocated portfolio shares to parties that were not in the cabinet (this breaks the mirror



Figure 2. Respondents' perceptions of portfolio shares for all perceived cabinet parties. Note: Solid lines indicate the mean of perceived portfolio shares and dashed lines represent the true shares.

image in allocations between partners and prime ministers). Here again, we see broadly accurate perceptions, though there is clearly more variation. Further, the under-allocation of seats to prime ministers that we saw in Figure 1 is again apparent and increases substantially in the wider sample. For example, the difference between the true and average perceived share of the British Conservatives increases from 12 per cent in Figure 1 to 30 per cent in Figure 2.⁸ Indeed, using the full sample, only about 6 per cent of our respondents allocate more than 50 per cent of the cabinet portfolios to a single party, even when the leading party of the cabinet actually controlled a significantly larger share.

What accounts for perceptions of portfolio shares

While perceptions of portfolio shares are not wildly inaccurate, at least on average, it is clear that there is considerable variation across respondents and that perceptions do depart from the true values, even when examining mean perceived allocations. Furthermore, the extent of these deviations varies across parties. To examine the sources of this variation, and to test the hypotheses we described above, we estimate a number of statistical models.

Variables	Model 1 Denmark	Model 2 Germany	Model 3 Italy	Model 4 United Kingdom	Model 5 Pooled
Perceived seat share	0.83*** (0.01)	0.92*** (0.01)	0.89*** (0.01)	0.86*** (0.01)	0.88*** (0.01)
Constant	6.29 (0.51)	2.97 (0.44)	3.17 (0.40)	5.48 (0.52)	4.24 (0.23)
N of respondents	2,500	2,431	3,093	2,151	10,175
\mathbf{R}^2	0.67	0.79	0.72	0.75	0.74

Table 1. Gamson's Law and perceptions of portfolio shares

Note: ***p < 0.01; **p < 0.05; *p < 0.1.

These models use both country-specific and pooled data from our four country surveys. The data are structured so that each row is an individual respondent's response for one of his or her perceived cabinet parties.⁹ The dependent variable records a respondent's perception of each perceived cabinet party's portfolio share on a 0–100 scale. Likewise, our major independent variables capture our respondents' perceptions about which party or parties held the prime ministry, as well as their perceptions of each perceived governing party's relative legislative seat share, also measured on a 0–100 scale.

Do voters use Gamson's Law?

Gamson's Law holds that the percentage of portfolios allocated to a party should be proportional to the legislative seats that party contributed to the cabinet. In a bivariate regression model of legislative seats shares on portfolio shares, with both variables measured on the same scale, this is equivalent to the joint hypothesis of an intercept of 0 and a slope coefficient of 1 on legislative seat shares.

To test the relationship between perceived legislative seat shares and perceived shares of cabinet portfolios, we thus estimate simple bivariate regressions of the rescaled perception variables described above. This corresponds to the usual tests of Gamson's Law, except that we replace cabinet and legislative shares with the perceived shares. Thus, an estimated slope coefficient of 1 and an intercept of 0 would provide unequivocal support for the idea that voters use the theoretical Gamson's Law as a heuristic for inferring portfolio shares from seat shares. Likewise, an estimated slope coefficient that is less than 1, as well as an estimated positive y-intercept would lend support to the idea that voters base their inferences about portfolio shares, in some way, on the empirically accurate version of Gamson's Law, with its a small party bonus.

Table 1 provides the regression estimates for each of our countries and the pooled sample. The evidence clearly supports the idea that many voters have perceptions of cabinet seat shares that closely fit the predictions of Gamson's Law, falling somewhere between the pure theoretical version and the empirically verified version (i.e., with a small party bonus). Indeed, given the widespread skepticism of voters' ability to navigate the complexities of multiparty, coalitional systems, the simple, strong linear relationships (with R^2 values around 0.75) captured in these equations are somewhat remarkable. Figure 3 provides the corresponding scatter plots, with the dashed lines indicating perfect Gamsonian perceptions



Figure 3. Gamson's Law and perceptions of portfolio shares. Note: Solid lines indicate the regression lines and dashed lines represent the perfect Gamsonian perception.

and the solid lines being regression lines. Except for Denmark, they clearly suggest that perceived seat and portfolio shares fit the theoretical Gamsonian predictions better than the true allocation of portfolio shares (i.e., slopes around 0.9 versus 0.8 for actual seat and portfolio shares), while not quite matching theoretical predictions: in each case, the slope coefficients are statistically different from 1.

Accounting for the prime minister

The second set of heuristics that voters might use to infer the portfolio shares of cabinet parties relies on the identity of the prime ministerial party. To test these hypotheses, we simply add a dummy variable for the perceived prime ministerial party to the regressions of legislative and cabinet seat shares reported above. This addition to the model produces two changes shown in Table 2. First, there is a large and statistically significant effect of being the perceived prime ministerial party dummy variable results in a small decline in the proportionality of legislative and cabinet seats, which translates to a larger small-party bonus that is particularly manifest when comparing across junior partners.¹⁰

Variables	Model 1 Denmark	Model 2 Germany	Model 3 Italy	Model 4 United Kingdom	Model 5 Pooled
Perceived seat share	0.74*** (0.02)	0.85*** (0.01)	0.71*** (0.01)	0.77*** (0.01)	0.77*** (0.01)
Perceived prime ministerial status	6.12**** (0.80)	5.49*** (0.55)	16.24*** (0.78)	7.90**** (0.71)	8.83*** (0.36)
Constant	7.06*** (0.52)	3.63*** (0.43)	3.41*** (0.38)	5.56*** (0.51)	4.73*** (0.23)
N of respondents	2,500	2,431	3,093	2,151	10,175
\mathbf{R}^2	0.68	0.80	0.76	0.76	0.75

Table 2. Gamson's Law, prime ministerial status and perceptions of portfolio shares

Note: ***p < 0.01; **p < 0.05; *p < 0.1.



Figure 4. Portfolio shares of perceived prime ministers and cabinet partners. Note: Solid lines indicate the regression lines while dashed lines represent the perfect Gamsonian perception.

Clearly, the small-party bonus, or large-party penalty and the prime ministerial party bonus can either be reinforcing for perceived small prime ministerial party's or perceived large partners, or countervailing for perceived large prime ministerial party's and small partners. As Figure 4 shows, however, both effects are empirically discernible in our data

Variables	Model 1 Denmark	Model 2 Germany	Model 3 Italy	Model 4 United Kingdom	Model 5 Pooled
Perceived seat share	0.59*** (0.02)	0.73*** (0.02)	0.60*** (0.01)	0.68*** (0.02)	0.77*** (0.01)
Perceived prime ministerial status	3.33*** (1.08)	3.12**** (0.92)	5.41*** (1.01)	4.61**** (0.85)	4.47*** (0.48)
True portfolio share (A	Ref. = 0)				
SD (63)	17.78*** (1.29)				13.07*** (0.69)
RV (37)	15.37*** (0.65)				13.44*** (0.05)
CDU/CSU (62.5)		11.42*** (1.11)			14.40**** (0.67)
SPD (37.5)		8.55*** (0.65)			11.16*** (0.52)
PD (66.66)			19.62*** (1.17)		19.14*** (0.70)
NCD (20)			6.02^{***} (0.77)		6.36*** (0.63)
SC + UdC (6.67 each)			-1.71** (1.17)		-1.18*** (0.56)
Conservative (83)				12.34*** (1.16)	15.57*** (0.70)
LDP (17)				6.09^{***} (0.74)	7.98*** (0.50)
Constant	2.57*** (0.54)	2.35*** (0.44)	3.89*** (0.48)	3.84*** (0.62)	3.08*** (0.26)
N of respondents	2,500	2,431	3,093	2,151	10,175
\mathbb{R}^2	0.74	0.82	0.76	0.76	0.75

Table 3. Heuristic use, and direct knowledge of portfolio shares

Note: ***p<0.01; **p<0.05; *p<0.1.

since, despite the strong preponderance of large prime ministerial parties and small partners, there are cases in which respondents perceived small prime ministerial parties and relatively large partners.

Direct knowledge of portfolio shares

Finally, it could certainly be that rather than inferring cabinet portfolio shares from party sizes and party roles, some individuals who pay a lot of attention to political news simply have direct knowledge of these shares. We can examine this possibility by including the true portfolio shares in the model along with the variables capturing the perceptual inputs to our possible heuristics. If the estimated impact of these perceived inputs remains after controlling for true portfolio shares, it suggests that that individuals are using the corresponding heuristic rules. However, if controlling for true values of the heuristic inputs eliminates these effects, then that suggests that the previously reported positive relationships between perceived input variables (i.e., perceived legislative seat shares and perceived prime ministerial status) and perceived portfolio share is due to the simple fact that the true values of these variables are similarly related and many individuals have direct knowledge of the values of all these variables. Estimates for models that add measures of true portfolio shares are provided in Table 3. Notice that true portfolio shares, given their limited variation across the small number of cabinet parties in our sample, enter the model non-parametrically as 'fixed effects' for each party (the actual seat shares are in parenthesis next to the party name).

This is preferable to, for example, adding a linear term for true seat size, which would impose a strong parametric structure on the relationship between true portfolio size and perceived size. In addition, because in each case one, and only one, of these listed parties was also the true prime ministerial party (the first listed for each country), this set of fixed effects simultaneously controls for true prime ministerial party status.

Revisiting our explanations above, if we estimated this model for a set of respondents relying entirely on Gamson's Law, we would estimate '1' for the effect coefficient on our measure of perceived seat share and '0' on all others. In contrast, if our respondents used only a prime ministerial dominance heuristic, they would assign all portfolios to the perceived prime ministerial party and none to others, yielding a coefficient of '1' to '100' on the prime ministerial covariate and a '0' on all others. Finally, if all our respondents had accurate knowledge of the true portfolio shares of parties and simply reported this knowledge, then the coefficient on our measure of true portfolio shares would be equivalent to the true shares (in parentheses) and all other coefficients '0'.¹¹ Thus, the results clearly suggest that the perceived legislative seat shares and perceived prime ministerial party status continue to have large, statistically significant, effects on perceived portfolio shares even when controlling for real cabinet shares. This result is consistent with the heuristic hypothesis that many voters are indeed relying on these perceptions to form their perceptions of cabinet shares, even when the perceptions underlying the inputs are wrong. That said, the coefficients on our measures of the true values of these inputs are also large and statistically significant, suggesting that these values condition perceptions of portfolio shares in some way, even when controlling for perceived legislative seat shares.

We can examine the sources of these effects further by assuming that individuals who are more interested in, and pay more attention to, politics are also more likely to have direct knowledge of portfolio shares. In this case, we should expect the impact of true portfolio shares to be greater, and the impact of the two heuristic inputs – perceived legislative seat shares and the perceived identity of the prime ministerial party – to be smaller for individuals with greater interest in, and attention to, politics. Consequently, we asked all of our respondents about how interested they are in politics and so can divide our sample into those more or less interested in politics.¹² The results are given in Table 4 for each individual survey as well as the pooled data.

Looking first at the pooled models, the results are consistent with the idea that some individuals have direct knowledge of portfolio shares. Specifically, among the more informed there is: (1) a much weaker relationship between perceived seat share and perceived prime ministerial party status, on the one hand, and perceived portfolio allocation, on the other; and (2) a stronger association between true portfolio shares and portfolio perceptions. These are exactly the results we would expect if our model is really capturing the difference between using heuristics and direct observation.

Further, these relationships hold across the individual countries: the differences between high and low interest cases are in the right direction and statistically significant in almost all cases.¹³ Thus, in all four countries, these results support the idea that some voters are using heuristics rather than directly (and correctly) perceiving cabinet portfolios. We know this because the estimated effects of the heuristic inputs (perceived seats and prime ministerial party) do not become insignificant or substantively meaningless when true portfolio shares are included in the model. In addition, we find strong support across all four countries for

)	•						
		,					Un	ited		
	Den	ımark	Gem	nany	It:	aly	King	dom	Poc	bled
Variables	Low interest	High interest	Low interest	High interest	Low interest	High interest	Low interest	High interest	Low interest	High interest
Perceived seat share	0.69^{***} (0.03)	0.51*** (0.02)	0.73*** (0.03)	0.72*** (0.02)	0.64^{***} (0.03)	0.59^{***} (0.02)	0.71*** (0.02)	0.65*** (0.02)	0.68^{***} (0.01)	0.61^{***} (0.01)
Perceived prime	3.01 (2.00)	2.81^{**} (1.25)	2.61^{*} (1.52)	3.76^{***} (1.17)	7.32*** (2.35)	5.19^{***} (1.10)	5.68*** (1.22)	3.53^{***} (1.19)	4.90^{***} (0.84)	4.28^{***} (0.59)
ministerial status	v = 0									
The purifound state (We	(n - i)									
SD (63)	13.34*** (2.18)	22.94^{***} (0.76)							11.41^{***} (1.24)	14.64^{***} (0.84)
RV (37)	9.14^{***} (1.24)	18.91^{***} (0.76)							8.66^{***} (1.02)	15.52^{***} (0.57)
CDU/CSU (62.5)			11.07^{***} (1.90)	11.64^{***} (1.38)					11.48^{***} (1.20)	16.28^{***} (0.81)
SPD (37.5)			5.89^{***} (1.16)	9.99^{***} (0.79)					7.35*** (0.98)	13.20^{***} (0.61)
PD (66.66)					16.38^{***} (2.58)	20.79^{***} (1.31)			16.43^{***} (1.27)	21.04^{***} (0.84)
NCD (20)					4.25** (1.72)	6.68^{***} (0.85)			4.43^{***} (1.34)	7.27*** (0.70)
SC + UdC					-3.50^{**} (1.66)	-1.14 (0.78)			-2.97** (1.27)	-0.44 (0.62)
(6.67 each)										
Conservative (83)							8.96^{***} (1.70)	15.16^{***} (1.58)	11.18^{***} (1.17)	18.46^{***} (0.89)
LDP (17)							6.14^{***} (1.10)	5.89^{***} (1.00)	6.88^{***} (0.81)	8.44^{***} (0.64)
Constant	2.84^{***} (0.94)	2.34*** (0.66)	3.53*** (0.77)	1.65^{***} (0.53)	4.35^{***} (0.98)	3.75*** (0.55)	3.20*** (0.92)	4.39^{***} (0.84)	3.48^{***} (0.45)	2.87^{***} (0.31)
N of respondents	780	1,720	760	1,671	810	2,283	952	1,199	3,302	6,873
\mathbb{R}^2	0.73	0.76	0.81	0.82	0.77	0.79	0.77	0.76	0.77	0.79
Note: $**p < 0.01$; $*p < 0$	0.05; *p < 0.1									

Table 4. Political interest, heuristic use and direct knowledge of portfolio shares

NICK C. N. LIN ET AL.

the prediction that interested and informed individuals will rely less on heuristics and more on direct knowledge of portfolio shares. This result is not only interesting in and of itself, but, because it is consistent with previous work about who should use heuristics, it further increases our confidence that the basic heuristic explanation is sound for at least some voters.

Differences across contexts

While our four-country study was clearly not designed to explore variations across electoral or institutional contexts, some of the differences apparent in the various results above are suggestive and deserve further exploration. Most importantly, we saw that perceptions of portfolio shares were less accurate in the United Kingdom than in the other countries. The results in Table 4 are also marginally weaker for Germany, with differences in the politically interested and less interested somewhat less robust than elsewhere. However, if we rank-order the countries according to the impact of true shares on perceptions, we recover the same rank-ordering reported in Fortunato et al.'s (2016) cross-national analysis of partisan left-right knowledge: Danes seem to be best informed, then Italians and Germans, with British voters showing the lowest levels of true information in regards to both portfolio allocation and party positions. Of course, at the time of our survey the British respondents were represented by a coalition government for the first time (at least in the lifetime of the overwhelming majority of our respondents) and, as such, the lower accuracy of their perceptions is to be expected.

The extent to which heuristic use varies across contexts, however, varies much less and in a way that does not necessarily comport with expectations we may have from previous research. We believe that this may be a function of context, given that we have only one sample from each country. For example, the small differences between high and low interest respondents in Germany, as well as the reversed direction of difference in the relationship for perceived prime ministerial party status between high and low interest voters, suggest that there may be something about the context of the 'Grand Coalition' that heightens the informational challenges voters face. Perhaps without large size asymmetries in cabinet partners, the informational demands of distinguishing between partners are simply too high even for the most interested voters to achieve.¹⁴ Thus, because our sample is limited to a single time period, any differences we observe across countries may a be function of contextual factors at the time of our surveys, which makes us reluctant to draw strong inferences regarding these differences, save the exception of the lower accuracy in the case of the United Kingdom.

Perceptions of which portfolios parties control

We now turn to the second set of questions about portfolio allocation that we raised in the introduction: Which parties do voters think control which ministries and how do they form these perceptions? Again, our empirical investigation will allow for the possibility that voters obtain such information directly (most likely from the media), that they do not know it at all or that they infer it from more readily available information using heuristic rules. We begin, then, with a brief discussion of several such potential heuristics.

Potential cues and rules

Given the enormity of the research literature on government formation, there is surprisingly little research that directly examines the question of which parties obtain which portfolio and what drives this distribution.¹⁵ Nevertheless, the work that does exist (Browne & Franklin 1973; Budge & Keman 1990; Bäck et al. 2011; Laver & Schofield 1990; Warwick & Druckman 2006) highlights two important variables relevant to the allocation of portfolios that voters could plausibly use to inform their inferences about which parties hold which portfolios: the general importance of the ministry in question, and its topical salience to different parties.

First, a frequent refrain in discussions about which parties obtain which ministries, in both academic and editorial accounts, is that, *ceteris paribus*, parties prefer to control more 'important' ministries. Further, there appears to be a general consensus that this ministerial 'pecking order' is widely understood, with ministries such as foreign affairs and finance ranking above ministries such as culture and sport. Browne and Franklin (1973), for example, argue that some ministerial portfolios are more important than others, and they speculate that this matters when it comes to coalition bargaining. Specifically, they speculate that parties that are overpaid in terms of the number of portfolios are correspondingly underpaid in terms of importance of the portfolios that they get. Given this, one might expect parties with greater bargaining power in coalition negations, such as prime ministerial or larger parties, to obtain the more important ministries more often. And, indeed, one does find that prime ministerial parties, for example, do actually obtain important ministries more often than their partners. Given this empirical regularity, a plausible and reasonably accurate heuristic voters might use to inform their perceptions of which parties hold which ministries is just that 'the prime ministerial party is more likely to control ministries that are more important'.

Ministerial importance hypothesis: Voters tend to assign ministries that are more important to the prime ministerial party rather than to junior members.

The second variable scholars often discuss when accessing why different parties obtain different ministries is the different salience of the topical policy area to each party. Simply put, parties for which a particular policy area is very salient may have good reason to try to control those ministries. For example, Laver and Schofield (1990: 183) argue that 'each party has a particular set of policy concerns, seeking control over a specific portfolio as an instrumental means of advancing these'. Likewise, Budge and Keman (1990) rank ministries according to party policy objectives to indicate qualitative preferences of parties for individual ministries. Bäck et al. (2011) argue that government parties are more likely to get portfolios that are salient to them than portfolios that are not. Finally, Warwick and Druckman (2006) show that bargaining strength (i.e., seat contribution) influences the probability of getting a salient portfolio.

Given this, we suggest that voters may use the 'topical affinity' between a portfolio and a party to gauge how likely it is that the party controls the portfolio. The plausibility of this rule, of course, depends on voters having ready access to its informational inputs, which includes some knowledge of parties' topical affinities or policy 'brands' – that is, the issues or policy areas that are particularly salient to each party. Further, there is likely to be variation across parties in how clearly voters can identify clear topical affinities. For example, 'catch-all' parties are likely harder to link to particular policy areas than single issue parties (e.g., environmental parties or anti-immigration parties). Usefully, survey work asking voters to place parties relative to each other on various issue scales suggests that voters often do make distinctions between the policy positions, if not the priorities, of parties. Likewise, the voluminous literature on the relative emphasis that parties themselves place on different issues (e.g., the many analyses of the comparative manifestos project) makes it clear that parties often emphasise such differences in their communications. Indeed, the whole 'salience' theory of elections (e.g., Budge & Farlie 1983; Petrocik 1996) takes such differentiation as foundational.

Thus, if information about party brands is pervasive and cheap to obtain, then this information is potentially useful for voters who do not know directly which parties control which ministries, but want to make educated guesses using a simple heuristic based on topical affinity. For instance, voters using such a heuristic might guess that a labour party holds the ministry of labour or employment, or the greens hold the ministry of environment. Of course, for some ministries there may be no obvious topical affinities to the set of cabinet parties. For example, in our Danish case, there is no party that Danes particularly associate with education policy and so topical affinity is unlikely to be a useful guide to voters deciding which party might hold this ministry.

Topical affinity hypothesis: Voters are more likely to assign a ministry to a party when the topic of the ministry is more salient to the party.

In addition to the two heuristics mentioned above, voters may employ historical patterns as a guide to draw their inferences. For instance, it is possible that some parties tend to regularly capture some ministries rather than others across different cabinets formed over time and voters gradually come to learn these patterns. Armstrong and Duch (2010) make a similar argument when they suggest that even when party systems are large and complex, the number of parties that actually join cabinets or compete for the premiership is small and, because of this, voters will learn to disregard ahistorical potential cabinets (those composed of parties that do not regularly serve in government) when forming their coalition expectations. Of course, it is likely that if such historical regularities do exist in qualitative portfolio allocation, they are likely to be driven by the topical affinities of parties. Nonetheless, we have data on which portfolios each of the parties in our sample held in the past and so we can examine the impact of historical patterns in the qualitative distribution of portfolios on voter perceptions as a separate hypothesis.¹⁶

Historical regularity hypothesis: Voters are more likely to assign a ministry to a party that has held that ministry often in the past.

Data and variables

To explore voters' perceptions of qualitative portfolio allocation, we rely on an additional Internet survey that we conducted in Denmark in July 2015. In this survey, and unlike the others, we asked respondents the following question:

As you may know, the current government divides functional responsibility over broad areas of policy and administration among various government departments, each headed by a cabinet minister from one of the government parties. In the last question, you indicated that the current governing cabinet includes [the list of parties the respondent indicated were in the cabinet]. Please tell us which of these parties heads each of the government departments listed below.

Following this question, the respondent was shown a random selection of five of the 19 Danish cabinet ministries and asked to choose from among the parties the respondent has previously said were in the cabinet which party held the ministry.¹⁷ To understand the analysis below, it is important that the structure of these data is clear. Specifically, since we only allow respondents to choose a party that they had identified in the previous question as being in the cabinet, it is possible that a respondent had to choose from among a set of parties that did not even include the true cabinet partners. While this may seem odd on the face of it, it is in fact very useful in exploring possible heuristic sources of respondents' perceptions of ministerial control since, for these cases, it is *impossible* that such perceptions could have been shaped by direct knowledge of ministerial control that is likely derived from media messages. That said, 46 per cent of our Danish respondents identified the partisan composition of the cabinet correctly and an additional 19 per cent included all the cabinet parties with the addition of one or more others. Thus, for 65 per cent of the sample, the parties that truly controlled each ministry were in the respondents' choice set. Finally, an additional 19 per cent of respondents correctly identified at least one of the two government parties as being in the cabinet.

Figure 5 illustrates this clearly. It shows the percentage of respondents who indicated that each listed party held the indicated ministry. The true government parties were the Social Democrats (SD, the prime ministerial party) and the Radical Liberals (RV); and even in the worst case (the climate, energy and building ministry) about 63 per cent of respondents chose one of these two parties. For all other ministries the number is far higher. In addition, when one of these two government parties is not chosen, it is much more likely to be because respondents chose the 'don't know' category than because they chose an incorrect party. This figure, which to our knowledge is the first of its kind, is remarkable given the pervasive scepticism among political scientists about the public's usual levels of political knowledge. The party that controlled each ministry is indicated by double asterisks, and we can see that in every case the plurality choice was the correct party. Given the high rate of accuracy in identifying the cabinet partners in the sample, most of our respondents had only the two government parties to choose from in answering this question. Thus, a reasonable approximation to what we might see if respondents were guessing randomly between these two would be that the bars for SD and RV would be the same height. Therefore, the difference in these heights is also a good measure of the extent to which Danish respondents correctly perceived which party held the ministry. In general, this difference is large. For nine portfolios it is larger than 60 per cent, in four it is larger than 40 per cent, in another three it is larger than 20 per cent and in only three (i.e., the higher education and science, the environment and the interior ministries) it is less than 20 per cent.

Further, it is hard to miss the patterns in Figure 5 that hint at an explanation of how voters form these perceptions. Three of the traditionally most important ministries (finance,



Figure 5. Accuracy of respondents' perceptions of which parties control which portfolios. Note: The asterisks indicate the true ministerial parties.

defence and justice) held by the prime ministerial party are all perceived very accurately. Further, the one very important ministry (foreign affairs) that is perceived somewhat less accurately is not held by the prime ministerial party. Likewise, it is interesting that a substantial percentage of respondents perceive that RV - a party that clearly markets itself as the green alternative among the mainstream parties in the party system – held the ministry of environment despite the fact that the environment minister was from SD. This percentage is substantially higher than any other ministry held by SD.¹⁸

In the next section we examine these impressions more formally in a controlled statistical model, but even uncontrolled examination of the data clearly reveals patterns consistent with the idea that voters think prime ministerial parties control the most important ministries and that topical affinity plays an important role.

Estimation strategy and constructing variables

In order to sort out the influence of the heuristics we posited above, as well as to allow for the possibility that voters have direct knowledge of qualitative portfolio allocation, we estimate a series of controlled statistical models. The dependent variable in these models is the respondent's choice of one party as the perceived ministerial party from a set of parties, where that set varies from respondent to respondent, depending on which parties each thought were in the cabinet. For each respondent, we have five such responses corresponding

Variables	Model 1	Model 2	Model 3
Perceived prime minister Prime minister*Portfolio importance Topical affinity Historical experience True ministry True ministry*Importance	-0.178 (0.117) 0.456*** (0.104) 0.798*** (0.071) 2.162*** (0.096)	-0.062 (0.126) 0.229** (0.113) 0.237*** (0.081) 0.804*** (0.119) 1.027*** (0.047)	-0.062 (0.126) 0.229** (0.113) 0.216** (0.086) 0.811*** (0.120) 1.115*** (0.132) -0.082 (0.116)
Observations Log likelihood	11,006 -3206.345	11,006 -2956.137	11,006 -2955.884

Table 5. Conditional logit models of qualitative portfolio allocation

Note: ***p < 0.01; **p < 0.05; *p < 0.1.

to the five ministries randomly assigned to the respondent so that the unit of analysis is a 'ministry choice opportunity' for a given respondent. Given this data structure, we use a conditional logit model (McFadden 1973), which allows us to estimate the impact of various characteristics of parties and ministries (e.g. their topical affinity and their importance) on the probability a respondent chooses a given party and is appropriate for shifting choice sets.

To examine the ministerial importance hypothesis, we first generate a dummy variable that indicates whether the respondent identified each party as the prime ministerial party and then interact this with a measure of the importance of each ministry. Since our survey did not ask respondents about their perceptions of the importance of each ministry, we rely on the portfolio salience index provided by Druckman and Warwick (2005) to capture the importance of ministries in Denmark.¹⁹ To investigate our topical affinity hypothesis, we create a measure for each ministry-party combination indicating whether that ministry dealt with topics that are particularly salient to the party's brand.²⁰

To assess the historical regularity hypothesis, we need a variable that captures the past experience of a party holding a particular ministry. For this purpose, we construct a simple measure that calculates the percentage of the total number of years that a party has held a ministry over the past 20 years.²¹ We obtain the information on portfolio allocation in Denmark from Seki and Williams (2014) and focus on the period between 1995 and 2015.

Finally, as we did in the last section, we can also estimate the extent to which individuals may have direct knowledge of cabinet composition, rather than relying on the heuristics we identified, by including a control variable in the model for the true party holding each ministry. We also include an interaction term between ministerial importance and the true controlling party in order to account for the possibility that voters are more likely to have direct knowledge about which party controls more important ministries that garner greater media coverage than less important ones.²²

Empirical results

Our statistical results are presented in Table 5 and each hypothesis has some support. First, the interaction between portfolio importance and perceived prime ministerial party is large



Figure 6. Difference in predicted probability of ministerial assignment. Note: Grey area indicates the 95 per cent confidence interval.

and significant across all three models, suggesting that respondents may be relying on the simple portfolio importance heuristic in which they assign the more important ministries to the prime ministerial party.

Figure 6 shows the substantive size of this effect (with 95 per cent confidence bands).²³ For less important ministries there is no difference between the predicted probabilities of a respondent assigning a ministry to the perceived prime ministerial party versus a perceived cabinet partner. However, as the importance of the ministry in question rises, this difference becomes larger (close to 20 per cent). Clearly, the result suggests that voters are substantially more likely to believe that the prime ministerial party controls more important departments than its partner does.

The estimates also support our topical affinity hypothesis. We find a positive and significant coefficient estimate on our measure of the topical affinity between a ministry and policy areas central to a party's brand. More precisely, when a voter perceives that a particular cabinet post deals with issues central to a party's brand, he or she is more likely to think that the party holds the ministry. The data similarly support the idea that history can provide guidance to voters in shaping their perceptions of qualitative portfolio allocation. One should be cautions, however, at taking this result at face value. An obvious interpretation is that voters learn the qualitative distribution of ministerial portfolios and

Variables	Model 1 Low interest	Model 2 High interest
Perceived prime minister	-0.094 (0.155)	0.027 (0.218)
Prime minister*Portfolio importance	0.306** (0.139)	0.099 (0.196)
Topical affinity	0.272**** (0.101)	0.172 (0.136)
Historical experience	0.670*** (0.149)	1.051*** (0.200)
True ministry	0.896^{***} (0.060)	1.233**** (0.078)
Observations	6,949	4,057
Log likelihood	-1964.932	-977.223

Table 6. Political interest, heuristics and direct knowledge of qualitative portfolio allocation

Note: ***p < 0.01; **p < 0.05; *p < 0.1.

retain this knowledge over time, using it to form adaptive expectations of the present allocation of ministerial portfolios. This, however, flies in the face of other work that shows voters' explicit knowledge of historical features of cabinets (e.g., their historical memberships) is abysmal (Fortunato et al. 2015). Another possibility, then, is that in the same way that an estimate of the impact of a lagged dependent variable includes the impact of unmeasured variables that have persistent effects over time, this historical variable could be picking up the impact of variables not included in our model, or measured only partially in our model, that are nevertheless predictive of real outcomes.²⁴ Whatever the particular mix between topical affinity and history as drivers of voters' perceptions of which cabinet parties hold which ministries, however, the overall message of these results is clear: voters' perceptions of which parties hold which ministries are both eminently sensible and plausibly due to simple heuristics that require relatively little current information as inputs.

Direct knowledge

The coefficient on the true identities of the parties controlling each ministry is large and significant in Table 5, which is consistent with the hypothesis that some of our respondents have direct knowledge about the true partisan control of the ministries. As in the last section, we can examine the logic of this interpretation further by testing the intuitive hypothesis that respondents that are more interested in politics will rely more on such direct knowledge, and less on heuristics, than those with less interest in politics.

Once again, the data bear this out. In Table 6, we provide separate analyses for the top and bottom half of the distribution of political interest.²⁵ Consistent with the hypothesis, our variable for true ministry holders is much larger among the politically interested and informed than those less so (odds ratio of 3.43 versus 2.45). Further, the results suggest that more interested and informed voters are also more likely to use the history as a guide, perhaps because they are more likely to have that information on hand or because they are using a larger set of accurate – in the sense of correctly predicting previous ministerial allocations – though unmeasured variables than those less interested and informed. Likewise, the coefficients on our heuristic indicators (portfolio importance

and topical affinity) are smaller for the more informed and larger for the less informed, indicating the latter group is relying more on these heuristics and less on direct knowledge. This, of course, is exactly the pattern we found when examining perceived portfolio shares. Since there is nothing about the measurement or analyses of these two different questions that should make them co-dependent, this consistency gives us additional confidence in both results.

Conclusion and discussion

Voter knowledge of the distribution of ministerial portfolios across cabinet parties is integral to several modern explanations of voter behaviour in coalitional systems, including those that depend on voters retrospectively assessing responsibility for policy outcomes and those in which they prospectively forecast the likely policy outcomes of potential cabinets. The informational demands that such models impose on voters, however, strike many critics as daunting. In this article, we provide the first systematic empirical investigation relevant to these assumptions and their critiques. Specifically, we have investigated voters' perceptions of parties' relative shares of ministerial portfolios, as well as of which parties control particular ministries.

Our approach allows for three possibilities: (1) that voters generate their perceptions of portfolio allocation by using one or more heuristics; (2) that voters possess information on the true portfolio allocation that is likely derived from media consumption; and (3) that voters' perceptions of portfolio allocation are whimsical, unsystematic or random. Our analyses reveal strong evidence that many voters have relatively accurate perceptions of the distribution of ministerial portfolios over parties as well as which parties hold which ministries. Further, these perceptions are consistent with the patterns we would expect if voters formed them by applying several plausible heuristics. Voters perceive parties to have a share of portfolios roughly proportionate to their share of seats in cabinet, though they give bonuses to small parties, which is consistent with the version of Gamson's Law that has been empirically verified repeatedly in the literature on cabinet formation. In addition, we find evidence that voters perceive prime ministerial parties to have more portfolios than their size (via Gamson's Law) would indicate - a result consistent with work that finds a prime ministerial heuristic operative in other kinds of voter perceptions (e.g. allocation of responsibility to parties). When it comes to the qualitative allocation of portfolios, voters are more likely to believe that the prime ministerial party controls portfolios that are important across parties, and to believe that parties are able to secure portfolios that are more salient to their brand (topical affinity). Likewise, the historical record of ministerial service is consistently related to perceptions, though it is unclear whether this suggests a direct role for history or the omission of other current, but previously predictive, variables in our perceptions model.

Importantly, despite clear consistency with a heuristic model of perception formation, we *cannot* reject the hypothesis that some voters possess true, direct information about the allocation of ministerial portfolios and so do not need to use heuristics to infer it. Indeed, we find fairly strong support that some voters possess such information, especially if they are interested in politics. In sum, the evidence examined paints a picture of coalitional voters that are largely well-informed about both the quantitative and qualitative aspects

of cabinets. While some of these voters likely obtain this information directly by consuming news reports and retaining that information, most voters likely infer it by applying a set of simple heuristics.

More generally, our findings contribute to the growing literature on how voters in coalition government contexts are able to cope with the complexities of their political environment and which tools they use to do so. At the same time, however, our findings expose just how much we have to learn about the coalition voter. There are essential questions regarding what voters do, and do not, know about coalition governance that have simply never been asked, including, until now, whether or not voters have reasonably accurate perceptions of the distribution of ministerial portfolios.

Supporting Information

Additional Supporting Information may be found in the online version of this article at the publisher's web-site:

Table A1: List of Parties included in the Four-Country SurveyTable A2: List of Ministries included in the Danish SurveyTable A3: Robustness Check of Table 5 with Seat Share PredictionsTable A4: Robustness Check of Table 6 with Seat Share Predictions

Notes

- 1. The Adaptive Behaviour and Cognition Group, behind such work as Gigerenzer (2010) and Goldstein and Gigerenzer (2009).
- 2. Though, to our knowledge, there is no direct evidence about, for example, the relative frequency of media mentions of cabinet sizes versus legislative sizes.
- 3. An anonymous reviewer notes that small parties may appear to receive bonuses because their seat share may be too small to warrant even one portfolio based on proportionality, even though all cabinet parties must receive at least one portfolio. While previous work on the portfolio allocation (e.g., Warwick & Druckman 2006) has addressed this possibility and shown that the small party bonus in allocations persists after accounting for it, we note that our survey questions ask voters directly about portfolio and seat percentages rather than converting from answers about numbers of seats, and so avoids an automatic effect of rounding to one for very small parties. That said, if voters perceive a small party bonus *and* that bonus is partially caused by rounding, then that effect should be reflected directly in the reported perceptions of seat and portfolio shares.
- 4. Unfortunately, our data do not allow us to further parse heuristic application according to, say, differences in respondents' personalities, as a reviewer helpfully suggested. However, past research (e.g., Fortunato & Stevenson 2013) has shown that, in parliamentary countries such as these, political interest is very salient in predicting heuristic application.
- 5. The parties included in these surveys are listed in the Online Appendix.
- 6. Respondents were not given an explicit 'don't know' option. Rather, they were explicitly encouraged to guess if they did not know for sure.
- 7. While calculating portfolio shares in Italy, we exclude those cabinet posts held by independents.
- 8. About 90 per cent of the British respondents thought correctly that the Tories were in the cabinet. Yet these respondents tended to under-estimate the true portfolio share of the party. Specifically, the average perceived share for respondents who thought the Tories were in the cabinet is 47 per cent (not rescaled) or 53 per cent (rescaled), while the average share for the respondents who did not include the party in the cabinet is 25 per cent (not rescaled) or 20 per cent (rescaled).

- 9. We dropped respondents who thought that there was only a single party in government since these logically must have all the cabinet seats.
- 10. While the perceived prime ministerial party variable clearly has a robust positive impact on a party's perceived share of cabinet seats, the addition of this variable only marginally improves the overall fit of the model (about a one-point increase in \mathbb{R}^2 in each country, except for Italy, where the increase is four points). This, along with the decrease in the size coefficient when this variable is included, results from the fact that most perceived prime ministerial parties are also perceived to be large. As explained in the text, however, there is enough independent variation in size and prime ministerial party status to estimate the separate effects of the two variables.
- 11. This conclusion assumes there would be random variation in perceptions of legislative seat shares and prime ministerial party status that is independent of perceived portfolio shares, as there indeed appears to be in our data. If all respondents who directly, and accurately, perceived the true portfolio shares and reported that number also accurately perceived legislative seat shares, then the coefficients on other variables would not be zero but inestimable separately from those on the true shares.
- 12. This is the standard question that asks how interested the respondent is in politics. The respondents answered on a six-point scale in Denmark, Germany and Italy and on a ten-point scale in the United Kingdom. For all countries we consider and code individuals in the top half of the distribution as 'highly informed'.
- 13. The only wronged signed results are for perceived prime ministerial party status in Germany and true portfolio shares for the LDP in the United Kingdom. Statistically *insignificant* differences in coefficients for individuals with high and low levels of political interest are confined to the perceived seat shares variables for Germany (p = 0.33).
- 14. Of course, it could also be that instead of high-interest voters looking like low-interest ones because accurate information is hard to obtain, low-interest voters look like high-interest ones because, during periods of grand coalition, it is easier to correctly guess true portfolio shares without relying on cues like legislative seat sizes and prime ministerial party status (say close to 50 per cent each).
- 15. This may stem from the tendency of theoretical models of coalition formation to treat portfolios as undifferentiated spoils and so focus only on the number of portfolios parties are allocated instead of qualitative allocation.
- 16. We are grateful to an anonymous reviewer for suggesting this.
- 17. The list of these ministries is provided in the Online Appendix.
- 18. This characterisation of the RV, based on our including our Danish expert's reading of recent Danish politics, is consistent with Benoit and Laver's (2006) expert salience coding of parties on the 'environment' policy dimension. To be sure, *Enhedslisten* and *Socialistisk Folkeparti* have a higher expert salience score on the environmental dimension than *Radikale Venstre*, but those two parties are usually not considered mainstream (e.g., Adams et al. 2006).
- 19. Note that because the portfolio importance variable does not vary over alternatives (parties) within observations, this variable cannot be estimated as a 'constituent' term to the interaction.
- 20. The way we constructed this variable is presented in the Online Appendix.
- 21. We also generate an alternative measure that captures the percentage of the number of times that a party has held a ministry in this 20-year period. This is calculated as a ratio of the number of times a party has held a given portfolio to the total number of times that the portfolio has been included in cabinets over the past 20 years. This alternative measure yields the same results.
- 22. Note that in this 2015 Danish survey we did not ask respondents about their perception of party seat shares and therefore we do not control for this in the analysis. That said, we asked respondents about their prediction of party seat shares in the upcoming election and including this prediction variable does not change the substantive results reported here. We include these models in the Online Appendix for interested readers.
- 23. We use model 2 and calculate the confidence interval for the first difference in the predicted probability of ministerial assignment via a parametric bootstrap (King et al. 2000).
- 24. See, e.g., Martin and Stevenson's (2010) discussion of why the incumbency status variables (i.e., 'historical' variables) usually included in models of cabinet formation are inflated due to exactly this kind of effect.

25. For this analysis, we made an additive interest scale from the usual four-category interest in politics question and a four-category self-assessment of how informed the respondent is about politics.

References

- Adams, J. et al. (2006). Are niche parties fundamentally different from mainstream parties? The causes and the electoral consequences of Western European parties' policy shifts, 1976–1998. *American Journal of Political Science* 50(3): 513–529.
- Armstrong, D. & Duch, R. (2010). Why can voters anticipate post-election coalition formation likelihoods? *Electoral Studies* 29(3): 308–315.
- Ansolabehere, S. & Jones, P.E. (2010). Constituents' responses to congressional roll call voting. American Journal of Political Science 54(3): 583–597.
- Ansolabehere, S. et al. (2005). Voting weights and formateur advantages in the formation of coalition governments. *American Journal of Political Science* 49(3): 550–563.
- Bäck, H., Meier, H.E. & Persson, T. (2009). Party size and portfolio payoffs: The proportional allocation of ministerial posts in coalition governments. *Journal of Legislative Studies* 15: 10–34.
- Bäck, H., Debus, M. & Dumont, P. (2011). Who gets what in coalition governments? Predictors of portfolio allocation in parliamentary democracies. *European Journal of Political Research* 50: 441–478
- Baron, D.P. & Ferejohn, J.A. (1989). Bargaining in legislatures. American Political Science Review 83(4): 1181–1206.
- Bassi, A. (2013). A model of endogenous government formation. *American Journal of Political Science* 57(4): 777–793.
- Benoit, K. & Laver, M. (2006). Party policy in modern democracies. London: Routledge.
- Browne, E.C. & Franklin, M. (1973). Aspects of coalition payoffs in European parliamentary democracies. *American Journal of Political Science* 67(2): 453–469.
- Browne, E.C. & Frendreis, J.P. (1980). Allocating coalition payoffs by conventional norms: Assessment of the evidence for cabinet coalition situations. *American Journal of Political Science* 24: 753–768.
- Budge, I. & Farlie, D.J. (1983). *Explaining and predicting elections: Issue effects and party strategies in twenty-three democracies*. London: Allen & Unwin.
- Budge, I. & Keman, H. (1990). Parties and democracy: Coalition formation and government functioning in twenty states. Oxford: Oxford University Press.
- Cheibub, J.A., Przeworski, A. & Saiegh, S.M. (2004). Government coalitions and legislative success under presidentialism and parliamentarism. *British Journal of Political Science* 34(4): 565–587.
- Downs, A. (1957). An economic theory of democracy. New York: Harper & Row.
- Druckman, J.N. & Warwick, P.V. (2005). The missing piece: Measuring portfolio salience in Western European parliamentary democracies. *European Journal of Political Research* 44(1): 17–42.
- Duch, R.M. & Stevenson, R.T. (2008). The economic vote. Cambridge: Cambridge University Press.
- Duch, R.M., May, J. & Armstrong, D.A. (2010). Coalition-directed voting in multiparty democracies. *American Political Science Review* 104(4): 698–719.
- Duch, R.M., Przepiorka, W. & Randolph, T.S. (2015). Responsibility attribution for collective decision makers. *American Journal of Political Science* 59(2): 372–389.
- Falcó-Gimeno, A. & Indridason, I.H. (2013). Uncertainty, complexity and Gamson's law: Comparing coalition formation in Western Europe. West European Politics 36(1): 221–247.
- Fortunato, D. & Stevenson, R.T. (2013). Perceptions of partian ideologies: The effect of coalition participation. American Journal of Political Science 57: 459–477.
- Fortunato, D. & Stevenson, R.T. (forthcoming). Heuristics in context. Political Science Research and Methods.
- Fortunato, D., Lin, N. & Stevenson, R.T. (2015). What Do Voters Know in Coalitional Democracies? Manuscript, Rice University.
- Fortunato, D., Stevenson, R.T. & Vonnahme, G. (2016). Context, heuristics and political knowledge: Explaining cross-national variation in citizens' left-right knowledge. *Journal of Politics* 78(4): 1211–1228. Gamson, W.A. (1961). A theory of coalition formation. *American Sociological Review* 26(3): 373–382.

- Gigerenzer, G. (2010). Moral satisficing: Rethinking moral behavior as bounded rationality. *Topics in Cognitive Science* 2(3): 528–554.
- Goldstein, D.G. & Gigerenzer, G. (2009). Fast and frugal forecasting. *International Journal of Forecasting* (25): 760–772.
- Indridason, I.H. (2011). Live for Today, Hope for Tomorrow? Rethinking Gamson's Law. Manuscript, University of California, Riverside.
- Kedar, O. (2005). When moderate voters prefer extreme parties: Policy balancing in parliamentary elections. *American Political Science Review* 99(2): 185–199.
- King, G., Tomz, M. & Wittenberg, J. (2000). Making the most of statistical analysis: Improving interpretation and presentation. *American Journal of Political Science* 44(2): 347–361.
- Lau, R.R. & Redlawsk, D.P. (2001). Advantages and disadvantages of cognitive heuristics in political decision making. *American Journal of Political Science* 45(4): 951–971.
- Laver, M. & Schofield, N. (1990). *Multiparty government: The politics of coalition in Europe*. Oxford: Oxford University Press.
- Martin, L. & Stevenson, R. (2010). The conditional impact of incumbency on government formation. *American Political Science Review* 104(3): 503–518.
- McFadden, D. (1973). Conditional logit analysis of qualitative choice behavior. In P. Zarembka (ed.), *Frontiers in econometrics*. New York: Academic Press.
- Meffert, M.F. & Gschwend, T. (2010). Strategic coalition voting: Evidence from Austria. *Electoral Studies* 29(3): 339–349.
- Petrocik, J.R. (1996). Issue ownership in presidential elections, with a 1980 case study. *American Journal of Political Science* 40: 825–850
- Powell, G.B. Jr. & Whitten, G.D. (1993). A cross-national analysis of economic voting: Taking account of the political context. *American Journal of Political Science* 37(2): 391–414.
- Saiegh, S.M. (2009). Political provess or 'lady luck'? Evaluating chief executives' legislative success rates. *Journal of Politics* 71(4): 1342–1356.
- Schofield, N. & Laver, M. (1985). Bargaining theory and portfolio payoffs in European coalition governments, 1945–83. British Journal of Political Science 15(2): 143–164.
- Seki, K. & Williams, L. (2014). Updating the party government dataset. *Electoral Studies* 34(2): 270–279.
- Snyder, J.M. Jr., Ting, M.M. & Ansolabehere, S. (2005). Legislative bargaining under weighted voting. *American Economic Review* 95(4): 981–1004.
- Tromborg, M.W. et al. (2015). The Mechanism Linking Context to Economic Voting: How Voters Attribute Responsibility. Manuscript, Rice University.
- Tversky, A. & Kahneman, D. (1974). Judement under uncertainty: Heuristics and biases. *Science* 185: 1124–1130.
- Warwick, P.V. & Druckman, J.N. (2001). Portfolio salience and the proportionality of payoffs in coalition governments. *British Journal of Political Science* 31(4): 627–649.
- Warwick, P.V. & Druckman, J.N. (2006). The portfolio allocation paradox: An investigation into the nature of a very strong but puzzling relationship. *European Journal of Political Research* 45(4): 635–665.

Address for correspondence: Nick C.N. Lin, University of Mannheim, L13, 15-17 – Room 421, D-68131 Mannheim, Germany. E-mail: nicklin@uni-mannheim.de