

Political Knowledge in Coalition Democracies*

David Fortunato

UNIVERSITY OF CALIFORNIA, MERCED

Nick Lin

Randolph T. Stevenson

RICE UNIVERSITY

Abstract

Using original survey data from France, the Netherlands, and the United Kingdom, we investigate the extent and structure of core political knowledge in coalitional democracies - knowledge needed to participate effectively in politics. Our analysis yields several interesting insights into what voters actually know about politics. We find, first, that there is pronounced variation in the types of knowledge voters possess across contexts. We also find that at the aggregate level several core pieces of political knowledge, factors that voters are so ubiquitously assumed to possess that we rarely, if ever, measure them, can be strikingly low. We also find that, while there is substantial variation across contexts, the structure of political knowledge *within* is remarkably unidimensional, but that the content of many, if not most, political knowledge batteries is ill-suited to tap into this underlying structure of core knowledge.

*Paper prepared for the Annual Meeting of Midwest Political Science Association, April 2014, Chicago, IL.

Introduction

From the point of view of the average citizen, parliamentary democracies, especially those that form coalition cabinets, are complex. Policy outcomes in such systems are the result of bargaining among multiple parties, with different policy positions, which join together in various kinds of policy making coalitions. In elections, a prospective policy-oriented voter must think about not only which parties share his preferences, but whether they are likely to get into cabinet and what their influence will be over policy if they do. Likewise, the retrospective, performance voter must ask herself which parties in a coalition cabinet (which may hold a minority of seats in the legislature) are actually responsible for the record of performance she has observed. Even purely affective voters must at least recognize the various party options and their own affective orientations to them. Given such informational difficulties, political scientists have traditionally been skeptical of voters' abilities to manage the complexity of these systems and cast either prospective policy oriented votes (i.e., Downs 1957) or even retrospective performance-based votes (Powell & Whitten 1993), instead suggesting that many voters in these systems turn to affective voting or simply just abstain.

In contrast to this traditional view, we (and various colleagues) have argued in a number of papers that citizens in complex coalitional systems can utilize a small set of heuristics to cut through their complexity and operate effectively.¹ Specifically, we posit that in order to vote effectively in these systems, or even just to talk sensibly about politics to friends and family, citizens use heuristics that help them (1) predict which policy-making coalitions are more or less likely, (2) attribute policy-making responsibility either prospectively or retrospectively, given an electoral outcome and cabinet composition, and (3) understand the relative policy positions of parties and track these overtime.

Each of the heuristics posited in these papers (and many of the others suggested in the literature) has the same structure: a simple rule that citizens use (perhaps subconsciously) to map a specific (limited) set of political information into the more complex beliefs and expectations that one needs to operate in these systems. Thus, they only “work” when citizens actually have accurate information

¹“Operating effectively” can include the ability of casting strategic votes (which has been the focus of much of the relevant literature) but also includes simply being able to discuss politics sensibly.

about their informational inputs.² For example, Duch & Stevenson (Forthcoming) argue that citizens in coalitional systems attribute responsibility to coalition cabinets using two simple rules: “The prime minister’s party is responsible for policy outcomes” and “The largest party is responsible for policy outcomes.” They show in a variety of lab and survey experiments that voters apply these two rules so consistently that if the same party holds both the prime ministry and the most seats in the legislature, almost no punishment accrues to anyone else. Such rules are clearly simple and (by definition) ignore a lot of potentially relevant information (like the identity of junior cabinet partners, the sizes of other parties, which party is median, etc.) in favor of a much smaller (ostensibly useful) information set.

The main challenge in this emerging literature is identifying exactly which heuristics are used in which contexts. The general expectation is that a heuristic will be used when it is “ecologically rational” given the context – that is, it accurately tracks outcomes over the long-term so that voters who use it tend to be satisfied with their inferences and decisions *ex-post* (Gigerenzer 2010). A variety of analytical strategies, ranging from statistical analyses of survey data to randomized lab experiments, have been used both to identify which candidate heuristics are being used in a given context and to determine whether these heuristics correspond to the ones that the “ecological rationality” argument predict should be operative. Fortunato & Stevenson (2014), however, point out that another way to explore which heuristics might be useful in a given context is simply to look at the distribution of knowledge of their informational inputs. If a heuristic, by definition, applies a simple rule to some input information set (in order to produce a more complex cognition), then it can hardly be widely useful in a given context if most citizens lack that input information.

Indeed, we might go further to suggest that the full set of “input information” to the most relevant and useful heuristics in a given context constitutes a kind of “core” political knowledge – i.e., the set of knowledge one would need to produce (using various heuristics) all of the key cognitions that are generally thought necessary to operate effectively in a given political context (that is, to

²Of course, individuals can feed incorrect information into a heuristic and thus make incorrect inferences from it. Indeed, Fortunato and Stevenson have argued that these misapplications of a heuristics are useful to political scientists in identifying whether or not an individual is actually applying a particular heuristical rule. In the long-run, however, we would not expect heuristics that routinely led voters astray (i.e., when they consistently used incorrect informational inputs) to continue to be used. Thus, in this paper we assume citizens who use a given heuristic are likely to gather the necessary information to use it accurately.

be seen as an informed or sophisticated citizen). As suggested above the literature on coalitional democracies suggest that these key cognitions include the ability to talk about likely policy-making coalitions (again, either predictively or in attributing responsibility for previous joint policy outputs), to identify the relative ideological positions of parties (at least on aggregate dimensions like the left-right, but perhaps also party positions on specific policies) and how these positions are moving over time, and to identify one’s affective orientations toward the parties and possible coalitions of parties.³

In this paper, then, we used new survey data on citizens’ political knowledge in three countries (the UK, France, and the Netherlands) to map which kinds of “core political knowledge” they have and which they do not. Table 1 maps out the various proposals for heuristics, their rules, and informational inputs that have been proposed in the literature and that guided our identification of the dimensions of core political knowledge we examine in this paper (and which we attempt to measure in our surveys).

Table 1: Heuristics Proposed by the Literature

| Authors | Target Cognition | Informational Input | Hueristical Rule |
|--|--|--|---|
| Armstrong & Duch (2010) | Expectations about which governing coalitions will form | Cabinet participation history | “Parties that have governed together in the past will govern in the future” |
| Fortunato & Stevenson (2013) | Updating the Left-Right policy positions of parties | Cabinet membership | “Parties in cabinet have converged together ideologically” |
| Fortunato, Stevenson & Spiegelman (2011) | Expectations about which party will be PM | Which party has the most legislative seats | “Largest party will be PM” |
| Duch & Stevenson (Forthcoming) | Policy making responsibility attributions | Which party has the most legislative seats | “The largest party is responsible for which policies are chosen” |
| Duch & Stevenson (Forthcoming) | Policy making responsibility attributions | Which party is the PM | “PM’s party is responsible for which policies are chosen” ⁴ |
| Fortunato & Stevenson (2014) | Expectations about which parties will join the PM in cabinet | Ideological positions of a parties relative to the PM | “Parties closer to the PM are more likely to be in cabinet” |
| Lau & Redlawsk (2006) | Perceptions of policy positions of a given candidate | Which politicians have which party labels | “Politicians’ policy positions are those of their party” |
| Green, Palmquist & Schickler (2002) ⁵ | Affective orientation to a given politician or policy | Which politicians have which party labels and which parties advocate which positions | “I like politicians from party X, and policy Z they advocate.” |

³These affective orientations are probably not strictly “cognitions”, but for convenience we will use the label to include them here.

The target cognitions in Table 1 include all those that Downs (1957) and Powell & Whitten (1993) and others have identified as necessary to cast either prospective policy votes or retrospective performance votes, as well as affective orientations to parties and policies. Likewise, the hueristical rules that allow the average voters to produce these cognitions with only limited investment make it clear that the key informational inputs citizens need to be able to engage in satisfying political discourse are (1) a recognition of the parties operating in a given context, (2) the relative ideological positions of these parties, (3) the current roles of the parties in government, and (4) the sizes of the parties. It has also been proposed that citizens need to know the historical record of government roles Armstrong & Duch (2010), but this has been disputed by Fortunato & Stevenson (2014).

But, of course, the question remains: do citizens in parliamentary systems even possess the information necessary to use these heuristics? Do they have only part of this information set and/or is its content context specific? If so, perhaps (as Fortunato & Stevenson (2014) have suggested) this variation could help us determine where one or another of these heuristics might be operative? More generally, do different dimensions of core knowledge scale together – e.g., if you know cabinet roles you are likely to know sizes as well? And, finally, how does this “core knowledge” map to other kinds of political knowledge that has been more often emphasized in the literature (e.g., institutional knowledge, knowledge or current events, knowledge or political personalities, etc.)? We do not know the answers to these questions because the kinds of survey questions necessary to answer them have not often (and for some dimensions never) been asked.⁶

In the rest of this paper, then, we take up these questions. We begin by describing the surveys we conducted in France, the UK, and the Netherlands, and then examine the extent to which citizens in these countries possess each of the informational inputs listed in Table 1 (which we are calling the dimensions of core knowledge). Following this variable by variable exploration, we then ask whether

⁴Duch and Stevenson’s experimental results really only apply to a situation in which one decision-maker is the agenda setter (out of five). However, they interpret this result in the lab as applying to coalition cabinets mainly through the oversized influence of the PM as the cabinet agenda setter. As they point out, however, one can also interpret their results as indicating the cabinet as a whole is the agenda setter.

⁵See also Koch (2001) and Rahn (1993), for instance.

⁶Our suspicion is that previous survey researchers simply assumed that some questions would be answered almost universally correctly (who is the PM) and others almost universally incorrectly (what are the specific sizes of the parties) and so were not asked. As we show below, neither of these assumptions appears to be correct in all contexts.

knowledge of these different inputs scales together or form separate knowledge dimensions (e.g., perhaps there are significant numbers of individuals who know about cabinet roles and not about sizes or ideological positions, or perhaps knowledge of one of the inputs reliably predicts knowledge of the others). Finally, in the penultimate section, we ask whether individual variation in core knowledge predicts variation in other kinds of political knowledge. We conclude with some observations about what our analysis implies for how one might want to try to measure political knowledge more generally (and especially if one want to compare knowledge across countries) and what it says about our conception of what it means to be politically sophisticated in these systems.

The Extent and Structure of Core Knowledge Items

The Surveys

In order to investigate the extent and nature of core political knowledge across different political contexts, we conducted three internet surveys on representative samples of adults in France, the UK, and the Netherlands.⁷ The UK survey was conducted in May of 2012 (two years after the 2010 general election that resulted in a Conservative - Liberal Democrat coalition cabinet). The French survey was conducted in January of 2013, almost nine months after the election of the Socialist Francois Hollande to the French Presidency and the installation of a coalition cabinet between the Socialists (30 cabinet seats), the Greens (2 cabinet seats) and the Radical Party of the Left (the PRG, with 2 cabinets seats) under the Socialist Prime Minister Jean-Marc Ayrault. The Dutch survey was conducted two weeks before the general election of September 2012 (so the incumbent cabinet at the time of the survey was the VVD led coalition with the CDA – with outside support from the PVV).

In each survey, we included a variety of demographic questions as well as questions tapping respondents' recognition of parties, their knowledge of parties' roles in cabinet or opposition, their knowledge of parties' histories of cabinet participation, their knowledge of party sizes, and their knowledge of party ideology. In the next section, we explore each of these in turn and following that

⁷The details of the samples are available from the authors.

look at the structure of core knowledge as a whole.

Party Recognition

The most basic piece of input information that in one's need to apply all the various heuristics in Table 1 is simply recognizing that a party is active in a given context. If an individual does not recognize a party, then she will clearly have little knowledge of that party's ideological positions, role in cabinet, or size. So, do most citizens at least recognize the parties that are active in politics in their systems? We examined that this question by asking:

Several political parties are listed below. Please tell us the ones you recognize as being active in [Dutch, French, British] politics, regardless of whether you can recall anything else about the party.

We then listed the full names of all the parties' active in the respective contexts, as well (in the French and British cases) several "fake" parties, with plausible sounding names. We allowed one of two answers: "Yes, I recognize the party" or "No, I do not recognize the party". "Don't Know" responses or skipping the question were not allowed.

Table 2 provides the percentage of the respondents who said they recognized each party, for the real parties. The table tells an interesting story about the relevance of party labels in different contexts that it will be important to keep in mind as we move forward. Specifically, while recognition of the main parties in each country is almost universal in the UK and the Netherlands, recognition of party labels is quite low in France. Only half of our French respondents said they recognized the UMP (*Union pour un Mouvement Populaire*), the main political party on the right in France. The UMP was formed in 2002 as a merger of several center-right parties under President Jacques Chirac and so it is both a relatively young party and a conglomeration of different previous parties and tendencies on the right. Still, at the time of our survey, it had been organizationally unified and had run candidates in both presidential and legislative elections under the UMP label for over a decade. Indeed, even when we consider much older French parties with clearly established "brand names", we find that the French respondents recognize them at lower levels than similar parties in the other countries. For

example, almost 20% fewer respondents recognize the French Socialist Party than the British Labour party or the Dutch PvdA.

Table 2: Recognition Rate of Political Parties

| | | | | | | | | | | |
|-------------|------|--------|--------|------|-------|--------|------|-------|------|------|
| Netherlands | VVD | PvdA | PVV | CDA | SP | D66 | GL | CU | SGP | PvdD |
| | 91.3 | 96.7 | 91.1 | 87.9 | 92.3 | 90.5 | 92.6 | 89.6 | 72.6 | 84.7 |
| France | PS | UMP | Greens | PRG | NC | PR | FN | MoDem | | |
| | 82.7 | 48.6 | 63.8 | 70.6 | 55 | 40.4 | 82.4 | 75.9 | | |
| UK | Tory | Labour | LDP | SNP | Plaid | Greens | UKIP | BNP | | |
| | 97.5 | 98.1 | 97.1 | 89.4 | 69.1 | 93.2 | 82.4 | 93.7 | | |

Note: Parties are ranked by their real seat shares, from the largest to the smallest.

Dutch Parties: Volkspartij voor Vrijheid en Democratie(VVD), Partij van de Arbeid(PvdA), Partij voor de Vrijheid(PVV), Christen-Democratisch Appellocal(CDA), Socialistische Partij(SP), Democraten 66 (D66), GroenLinks(GL), ChristenUnie(CU), Staatkundig Gereformeerde Partij(SGP), Partij voor de Dieren (PvdD).

French Parties: Parti socialiste(PS), Union pour un mouvement(UMP), Les Verts(Greens), Parti radical de gauche(PRG), Nouveau centre(NC), Parti radical(PR), Front National(FN), Mouvement democrate(MoDem).

British Parties: Conservative Party(Tories), Labour Party, Liberal Democrat Party(LDP), Scottish National Party(SNP), Plaid Cymru(Plaid), Green Party, United Kingdom Independence Party(UKIP), British National Party(BNP).

A comparison of recognition rates for the Green parties in the three countries tells the same story. Thirty percent fewer respondents recognized the French green party in the 2012 than recognized the British green party (which elected only one representative in 2010) or the Dutch Greens.⁸


Of course, caution is called for in interpreting these results. These are (as far as we have been able to determine) the only surveys to ever ask this kind of question and so we will need to see the French result repeated in other surveys. And, of course, this result begs for more subtle questioning to try to uncover the sources of the differences apparent across countries in Table 2. That said, we think these results are intriguing in that they suggest a potential difference in the usefulness of simple party heuristics across countries.

In Figure 1 we look at these same recognition rates (y-axis) plotted against each party's history of service in the cabinet (the x-axis). In addition, the size of the plotted circles indicated the party's

⁸The French Green party that we asked about was Europe Ecologie which had been formed in 2010 from a collection of other green parties, including the most prominent, Les Verts (which, of course, means "green"). Mindful of the possibility that French voters might not recognize Europe Ecologie by its actual name and just think of them generically as the greens, we listed the option as "Europe Ecologie - Les Verts".

share of seats in the legislature at the time of the survey. Clearly, there are few systematic patterns other than the one noted above (i.e., all French parties have recognition rates less than almost all of those in the other two countries). While having spent more time in cabinet may result in marginally higher recognition in the UK, there is not even a small effect apparent in the Netherlands or France. Further, while many of the largest parties appear toward the top of the graph, so do many small parties: At least in the UK and the Netherlands, small size is no barrier to very high recognition rates.

Figure 1: Party Recognition



Finally, in addition to recognition of actual parties, we also asked our respondents whether they recognized several fictional parties.⁹ These are helpful in that they let us establish a baseline level of “guessing” – for example, if we included only active parties and respondents simply claimed they recognized all parties (when they did not), we could not separate this from respondents with real knowledge. Thus, one way to use our results on fictional parties, is to subtract the percentage who “recognized” fictional parties from those percentage reported above to establish a lower bound on recognition (i.e., if all those who recognized fictional parties also did not really recognize the real parties).

⁹These were included along with the non-fictional parties and the order of the lists was randomized.

In the UK, we included three fictional parties: The “United Left”, the “Party for Progress”, and “England Alone” and only 9.36%, 9.56%, and 7.07% of respondents, respectively, reported that they recognized these parties. These percentages are similar to the percentage of British respondents who did not recognize the main parties. In contrast, and again consistent with the emerging idea that there is something quite different about the distribution of knowledge of party labels in France than the other countries in our sample, the French respondents reported recognizing fictional parties at twice the rate: “The Struggle of Farmers” party (24.70%), the “Union of French Republican Citizens” (21.44%), the “Rally for French Independence” (17.24%), and the “Liberal Party” (10.62%).¹⁰

With these results, we can thus think of the “lower bound” for recognition (averaged across all parties) as about 75-80% in the UK (85-90% recognition minus 10% who “recognized” even fictional parties) and perhaps as low as (55-70% recognition minus 20%) 35% to 50% for the French.¹¹

Cabinet Roles

The next domain of core political knowledge we investigated concerned parties’ current roles in government. Do citizens know which party hold the presidency or prime ministry, which are in the cabinet, which in the opposition, and which are not even in parliament? To get at this, we asked respondents in the UK and the Netherlands the following question:

Please choose the option which best describes each party’s role in the most recent government — the government formed after the last election: 1) Prime Minister, 2) In the current cabinet but not the current Prime Minister, 3) In the opposition, 4) Has no seats.

Given the structure of the question, it was possible for respondents to name more than one party (or no parties) as the PM and Table 3 lists the number of voters who listed different numbers of parties as PM. In both the Netherlands and the UK, about 75% of respondents named only one party as PM. Likewise, almost 20% of Dutch and 11% of the British respondents did not name a PM (many of

¹⁰Due to a misunderstanding about with the survey company implementing the Dutch survey, our fictional parties were (in an attempt to be helpful. . .) deleted from the survey and not asked.

¹¹This strikes us as very low and so we are somewhat skeptical – clearly more survey work will be needed before we would be willing to vigorously defend this conclusion.

these answered don't know to all or may of the parties in this question). Finally, the sizable number of respondents in the UK naming two prime ministers may be a symptom of British respondents' lack of experience with coalition cabinets before 2010 (so that to be a party "in government" or "in cabinet" was synonymous with being PM), so that some named both cabinet partners as the PM. A similar phenomenon did not happen in the Netherlands, though they too had a coalition cabinet (between the VVD and the CDA, with outside support from the PVV).¹²

Table 3: Knowledge on Number of Prime-Minister

| | % Named Each Number of PMs | |
|----|----------------------------|----------------|
| | Netherlands | United Kingdom |
| 0 | 17.48 | 10.65 |
| 1 | 75.97 | 74.33 |
| 2 | 3.28 | 13.73 |
| 3+ | 3.28 | 1.29 |

Turning now to the question of which parties citizens thought held the prime ministry (and the presidency in France), Table 4 gives the relevant results, sorted by how many parties were named as PM. We see that in all three countries, about 90% of the respondents who named only one PM, named the correct one (thus, about 70% of the respondents in the Netherlands and the UK named a single correct PM).

If we allow that some voters in the UK and the Netherlands may have named more than one PM and only require that one of the parties that was named was the actual PM, the rates of "correct" responses go up in these two countries to about 83% for the UK and 75% for the Dutch. Still, at this suggests that at least 20% simply could not name the party of the PM and closer to 30% were sufficiently hazy about it that they could not settle on a single party.

In France, the results suggest that knowledge of the party of the PM and President is considerably more widespread than in the UK and the Netherlands.¹³ Indeed, about 90% of French

¹²The French question was structured differently so that multiple parties could not be identified as the PM. Specifically, the French survey included one question asking the respondent to identify the party of the PM and another asking the party of the President.

¹³The number reported in the first part of Table 5 comes from a question asking respondents which party was the PM or President, respectively. Thus, there was no opportunity to name more than one party (and no Don't Know option).

Table 4: Knowledge on Prime-Minister and which Party holds it

| | Named 1 PM | |
|--------------------|-----------------|-------------|
| | Named Correctly | Named Wrong |
| Netherlands | 91.76 | 8.24 |
| UK | 94.24 | 5.76 |
| France (PM) | 89.10 | 10.90 |
| France (President) | 88.82 | 11.28 |
| | Named 2 PM | |
| | Named Correctly | Named Wrong |
| Netherlands | 75.76 | 24.24 |
| UK | 91.30 | 8.70 |
| | Named 3+ PM | |
| | Named Correctly | Named Wrong |
| Netherlands | 81.82 | 18.18 |
| UK | 83.33 | 16.67 |

Note: Question wording in France is different from the other two countries.

respondents knew the partisanship of the Prime Minister, and similarly for the President. This then is different from our earlier conclusion that the French were less able to recognize party labels and suggests that that result (or this one) may not be a symptom of a general (across knowledge domains) deficit in knowledge between the French and the other countries in our sample.

Turning now to knowledge about other cabinet roles, Table 5 summarizes the number of parties named as cabinet members (including the PM). The correct number for the UK was two (the Conservatives and the Liberal Democrats) and 70% of respondents chose two parties. For the Dutch case, the correct answer for was either two parties (the VVD and the PVDA) or three - if one counts the outside support of the PVV as being “in cabinet” (as many respondents may have done). But even if one accepts both answers as correct, only 57% of Dutch respondents identified the correct number of parties in the cabinet.¹⁴

¹⁴The French survey, unfortunately, did not ask respondents about cabinet roles other than the prime ministry and presidency. The cabinet did however contain two minor parties besides the socialists, each of which had two seats.

Table 5: Political Knowledge on the Number of Cabinet Parties

| | % Parties Named | |
|----|-----------------|----------------|
| | Netherlands | United Kingdom |
| 0 | 13.01 | 7.86 |
| 1 | 3.67 | 3.88 |
| 2 | 14.20 | 70.05 |
| 3 | 42.60 | 7.76 |
| 4+ | 26.51 | 10.45 |

The ambiguity of the Dutch case (in terms of which parties were actually incumbent complicates the presentation of the data on respondents identification of cabinet partners, so Table 6 simplifies this presentation by categorizing each respondents' answers as correct or not on three different questions (where we give results for three different ways to code correct responses for the Dutch case): Did the respondent name the correct PM? Did the respondent name the correct PM and the correct partners? And, did the respondent name the correct cabinet members (even if she switched which one was the PM)?¹⁵

Table 6: Political Knowledge on Party Roles, Summary

| Named Roles Correctly | Netherlands | | | United Kingdom |
|-----------------------------|--------------|---------------|------------------|----------------|
| | Government I | Government II | either GI or GII | |
| PM | 69.71 | – | – | 70.05 |
| Cabinet (w/ PM identified) | 9.04 | 30.98 | 40.02 | 60.90 |
| Cabinet (w/o PM identified) | 9.53 | 32.67 | 42.20 | 67.76 |

Note: Government I = VVD + CDA; Government II = VVD + CDA + PVV.

The results here reinforce the general impression from this section that the Dutch underperform the UK in knowledge of cabinet roles (as well as the French for knowledge of the PM). Indeed, even when we allow for the possibility of including the PVV as a cabinet or not, only about 40% of respondents identify the cabinet as either being led by the VVD with either the CDA alone or the CDA and the PVV as partners. Likewise, if we relax the requirement of naming the PM correctly, the percentage of correct responses still only reached 42% (the increase comes mostly from respondents

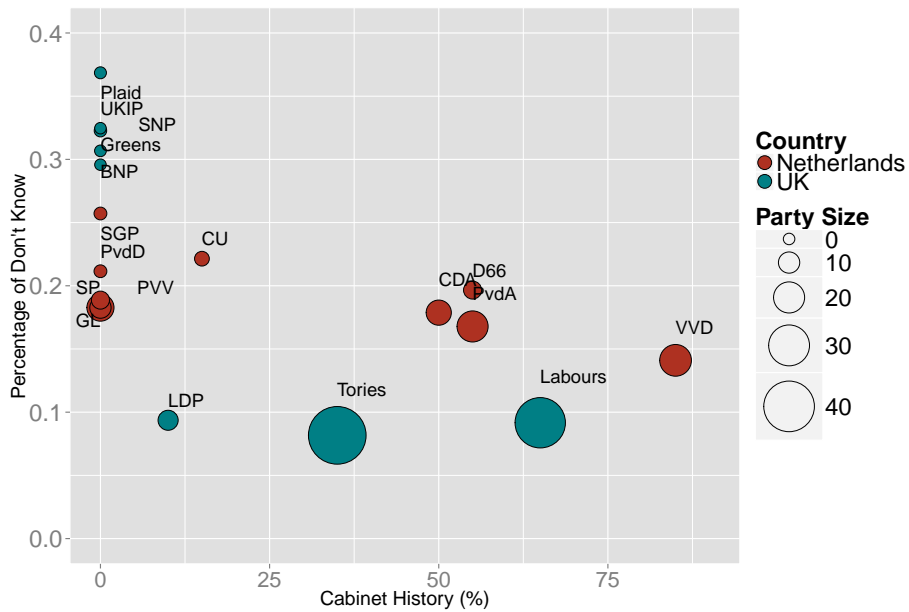
¹⁵In each case, we included only exactly correct answers – so the answer was wrong if the respondent names other parties in addition to the correct ones to a given role.

who think the CDA was the PM - as it had been PM most often since its formation in the 1970's). This contrasts to some degree with the UK, where about 61% of respondents named the correct PM and the cabinet. But, still if only 67% of British voters know who is in the cabinet (out of three major parties to choose from) then it suggests real limitations on the percentage of voters who have the basic information necessary to attribute policy making responsibility or to forecast the policies they should expect out of cabinet – even if they might otherwise use simple heuristics to accomplish these tasks.

We conclude this section by examining the distribution of Don't Know responses to the cabinet role question for the UK and the Netherlands. Figure 2 plots these DK responses by cabinet history and party size. Several conclusions are immediately apparent. First, consistent with the conclusions above, the baseline level of DK responses for the Netherlands and the UK differ, with the Dutch choosing DK at rates of about 20% for even fairly large parties with significant histories of cabinet service. The corresponding rate in the UK is closer to 10%. That said, for small parties with no (or little) history of cabinet service, the relative frequency of DK responses between the Dutch and the British switch, with rates above 30% for the British respondents and about 20% for the Dutch. To put it another way, the percentage of respondents answering DK for the British is highly conditional on party size and/or history of cabinet service, while for the Dutch the rate of DKs is close to 20% no matter how big the party or how often it has served in cabinet.

When taken together with the data on party recognition that was reported above, the results for the UK are interesting: while almost 90% of Britons say they recognized the smaller parties (and we are inclined to believe this number is actually close to this – say between 80 and 90% - given that only 10% “recognized” the fictional parties), less than 70% were willing to guess whether these parties were in cabinet – suggesting that these respondents either did not have a sense that they parties were perennial opposition parties (none have ever been in cabinet) or were unwilling to use that sense to guess about these parties current status. One intriguing possibility is that we are seeing in these data exactly the kind of informational updating that happens when new events shake a party system. After all, the first coalition cabinet in the lifetimes of most Britons had formed in 2010, so that a previously perennial opposition party – the Lib Dems – was now in cabinet. So perhaps this caused

Figure 2: Party Recognition



our respondents to question what they thought they “knew” about other parties – perhaps they should be more uncertain whether these small and opposition parties were part of the cabinet – perhaps in this new world of coalitions that was now possible.

History of Cabinet Participation

Armstrong & Duch (2010) have shown that there is considerable regularity in the composition of cabinets in the Western parliamentary democracies and conjecture that these “patterns of coalition formation provide voters with information that facilitate the kinds of strategic calculations that are prominent in rational voting theories.” But, of course, such historical regularities can only be the basis for voters’ expectations of the future (i.e. a useful heuristic for anticipating which cabinets will form and perhaps facilitating strategic votes) if voters actually know these patterns. To examine this question we asked our respondents in the following question (in all three countries):

What percentage of time over the last 20 years would you say each of the following parties have been in the Cabinet (either alone or in coalition with other parties)?

Respondents were asked to provide a number between 0 and 100; given an example; encouraged to guess if they were unsure; but told to leave the question blank if they would rather not answer.

Figure 3 provides the results for the UK, the Netherlands, and France, respectively. The graphs show the distribution of responses from 0 to 100 (for those that answered¹⁶) for each party and the parties are ordered on the graph from those with the least history of cabinet participation to those with the most. Note that over 75% of the UK respondents answered 0 for all the 5 small parties (i.e., SNP, Plaid, Greens, UKIP, and BNP), we decide not to present these distributions for graphing convenience.

While there is a lot of information in these graphs, there are two prominent conclusions that seem to apply across countries. First, voters do seem to be able to categorize parties into two sets of parties: those that often get into cabinets (in cabinet greater than about 30% of the last 20 years), and those that never get into cabinets (in cabinet 0% of the last 20 years, and perhaps including those less than 3% of this time in France).¹⁷ The median response in each country for perennial opposition parties is always essentially zero with a tight distribution around that number. Likewise, parties that often get into cabinet (above 30% of the time over the last 20 years for all parties) have median responses that are much larger than the perennial opposition parties. In this sense then Armstrong & Duch (2010) are right: the history of cabinet participation does seem to make its way (at least in the broad categorization) into the minds of citizens.

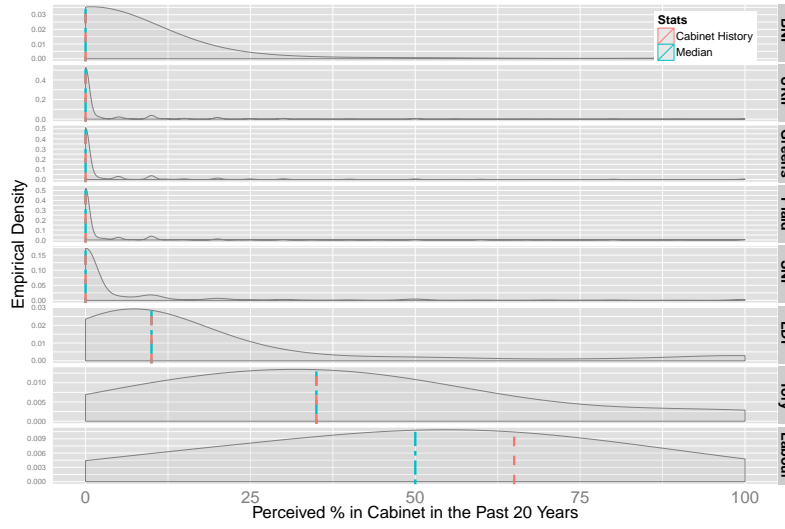
That said, however, if we go beyond this dichotomous classification to examine the full distribution of responses for those parties that often get into cabinet, it's clear that voters know little more than this broad dichotomous classification of parties into perennial parties of government or opposition. Examine, for example, the distribution of responses for the Labour and Tory parties, who served in cabinet in the UK for 65% and 35% of the last 20 years, respectively, is almost uniform over the whole range of possible responses. The same holds for the VVD, CDA, and PvdA, in the Netherlands and the Socialists in France. Likewise, and to a lesser extent for the Dutch D66 and French UMP.

¹⁶Given the encouragement to guess, DK responses for these questions were quite low, so we ignore them here.

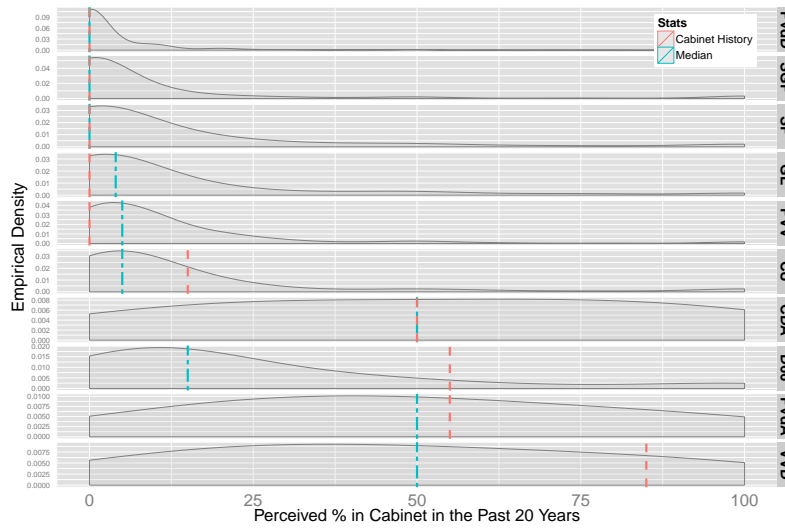
¹⁷There are only two parties in our sample that have served in cabinet more than 3% and less than 30% of the time in the last 20 years - the LDP in the UK (10%) and the CU in the Netherlands (15%). Thus, while it may be that citizens can distinguish a third category - those parties that are "sometimes" in cabinet, we really don't have enough evidence to say.

Figure 3: Knowledge on History of Cabinet Participation


(a) United Kingdom



(b) Netherlands

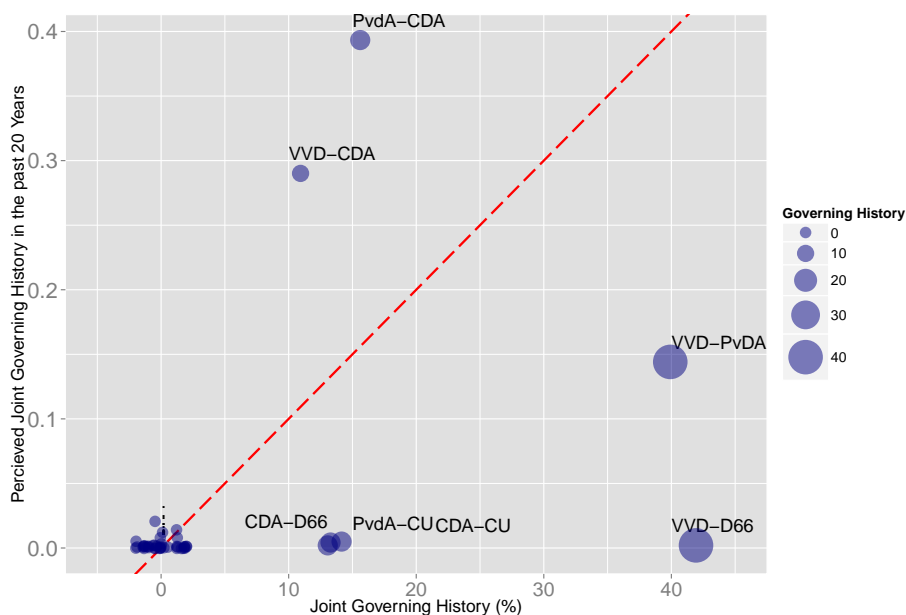


(c) France



Finally, given the emphasis on history as a guide to predicting the composition of cabinets, we were curious whether our respondents knew the history of joint cabinet service between different parties. So do citizens know which parties have been in cabinet together? To examine this, we ask respondents in the Netherlands an additional question similar to the one above, but directed it towards particular pairs of parties, which we selected from all the possible pairs for their relevance and to get some variation on length of real joint service. The results are graphed against the real history of joint cabinet participation in Figure 4 below.

Figure 4: Knowledge on History of Joint Governing in the Netherlands



Correct responses would be on the 45-degree line of this figure. Clearly, respondents do not have a good sense of which parties have served together. Several pairs with significant histories of joint cabinet service (e.g., VVD-D66 at over 40% and PvdA-CU, CDA-CU, and CDA-D66, all at around 15%) are perceived as never having served together, while the PvdA-CDA pair, which served in cabinet together over 40% of the time period, was perceived as only having served on average 15%. Thus, our conclusion from this, as well as from the dichotomous nature of knowledge about individual party histories of service in the cabinet, is that citizens do not appear to know the history of cabinet service well enough for this to be helpful in heuristics aimed at producing anything but the

broadest expectations about the probabilities with which different governing cabinets form.

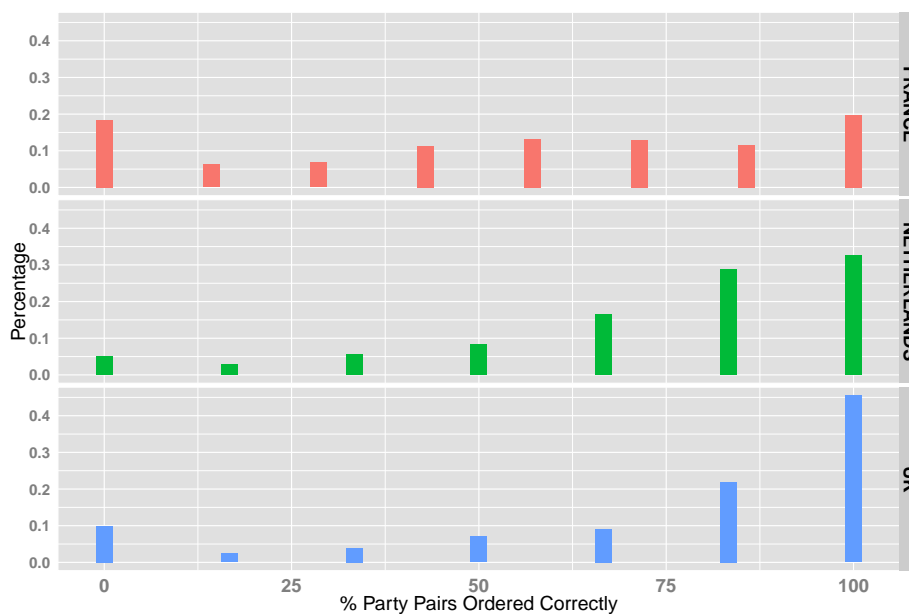
Party Sizes

To examine voters' knowledge on the sizes of political parties, we ask respondents the two questions. In the first question we provide 6 (French voters get 7) pairs of parties and ask respondents which one is bigger. This question is designed to address voters' knowledge on party sizes in a relative sense and included pairs of parties that covered the range of size differences. The question wording is as follows:

There are six pairs of parties listed below. For each pair of parties please click the button of the party in the pair that has more seats in the [lower house of the legislature].

Don't know responses were allowed. As shown in Figure 5, Dutch and UK voters seem to know the relative sizes of their parties well. Over 70% of the Dutch respondents and the UK respondents are able to get 75% of the pairs correct. In contrast, less than 30% of the French respondents can put three quarters of the pairs of parties in the correct size order.

Figure 5: Knowledge on Size via Pair Comparison



Further, we look a bit deeper into voters' knowledge on the sizes of political parties by asking a more specific question:

Now we would like to get more specific about the sizes of the parties. What percentage of the seats in the House of Representatives do you think each of the following parties holds? If you think a party has no seats in the House of Representatives, please indicate that by typing in a zero. Your answer for each party must be a number between 0 and 100. If you are not sure, please give us your best guess; but, if you really do not want to answer for a party, just leave the corresponding box blank.

Figures 6 give the answers to these questions. Like the graphs in the last section on party history, these provide the full distribution of responses for each party, the real seat shares, and the median response. Clearly, at least for the UK and the Netherlands, these distributions contrast greatly with those for cabinet histories. Especially in the Netherlands, respondents have a keen grasp of the exact seat shares of the parties, with quite accurate median responses and tight, normal distributions around these. In the UK, the distributions for larger parties are somewhat more spread out, but nothing like what we saw for cabinet histories.

This contrasts with the evidence for France. Consistent with the questions about the relative sizes of parties, the French do not do well at naming the specific seat shares of the parties. This is true across all the parties, but is especially true for the larger ones, for which we see distributions that cover the range of possible responses and are to some degree even bimodal (especially for the UMP).

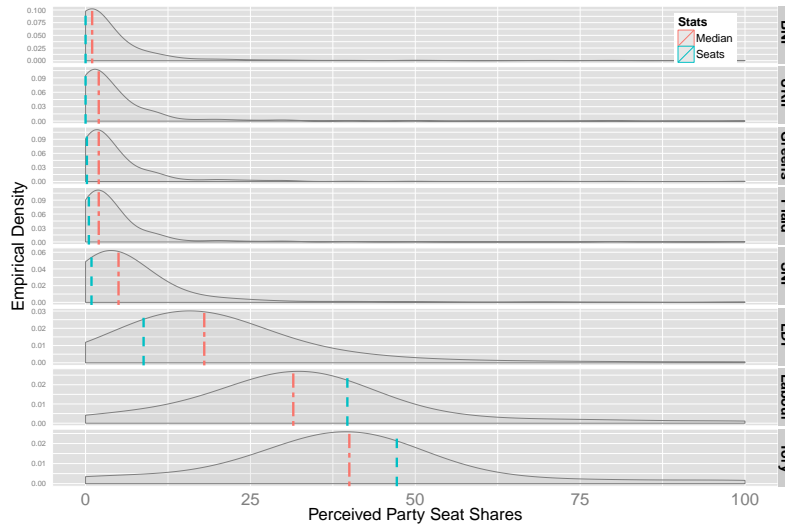
In sum, our investigation of respondents' knowledge of party sizes reveals two perhaps surprising conclusions: First, many individuals in the UK and the Netherlands seem to know the seat shares of parties, not just in a relative sense but the precise numbers. Second, there is a real contrast between the UK and Dutch respondent on the one hand, and the French on the others, with the latter voters unable to pin down either the exact seat shares of parties, or their relative sizes.

Left-Right Ideology


Besides the above factual knowledge, we also explore whether respondents are capable of using an abstract left-right ideological spectrum to comprehend the political system. For this purpose, we ask

Figure 6: Knowledge on Party Sizes

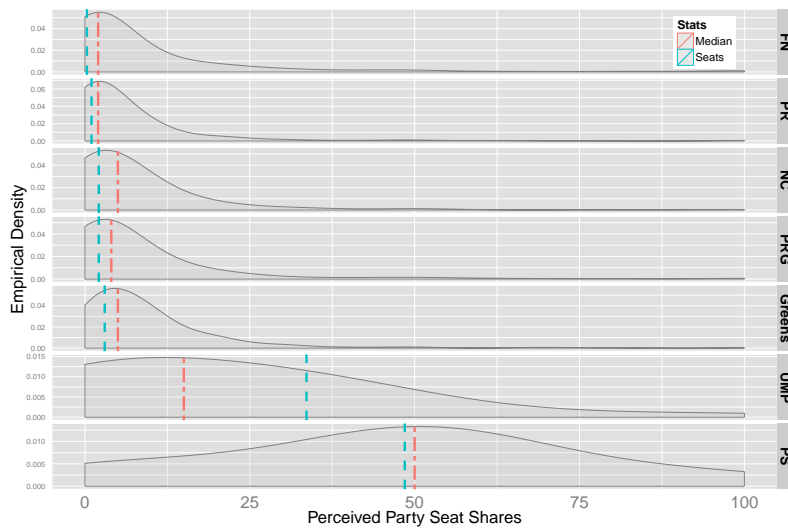
(a) United Kingdom



(b) Netherlands




(c) France



the respondents the traditional ideological placement question that requires them to place political parties on a 10-point scale, in which 1 indicates the most left position and 10 as the most right position. Like with party sizes, we divide our analysis into two parts. The first, asks if respondents know the relative positions of the parties and the second whether they know absolute positions.

To examine knowledge of relative positions, we translated responses about individual parties to responses about party dyads. And, we then coded a pair of answers as correct if the voters placed the two parties in the correct left-right order, taking the mean response of all respondents' placements of a party as our measure of the correct location.¹⁸ The results are presented in Figure 7.

Figure 7: Knowledge on Left-Right Ideology via Pair Comparison



First, the large spikes at zero in these graphs are usually due to respondents who answered don't know for most parties (or sometimes gave the same score to all parties, since we counted equal positions as incorrect orderings in our coding). Second, unlike the cross-country patterns seen so far, ideological knowledge seems to be less widespread among respondent in the UK and more widespread in France and the Netherlands. While 20% of Britons could place no pairs of parties in

¹⁸Fortunato, Stevenson & Vonnamme (2013) show in a large number of election surveys that orderings of parties base on the median survey response is almost always exactly the ordering one gets from other sources like expert placements and placements based on party manifestos.

the correct order and the median respondent only got slightly more than half the pairs correct, in the Netherlands, these numbers are about 10% and 85% respectively. And, in France the ability to place parties in the correct left right order is only slightly less widespread than in the Netherlands.

In addition to these relative comparisons, we also show the distributions of voters' placements of each party in Figure 8. The important thing in these graphs is the extent of dispersion in opinion about the left right positions of each party and in general we see a lot of disagreement among respondents about where to place parties on a left right scale. Against, this background, however, we do see that there is relatively more dispersion of opinion about the placement of British parties than Dutch or French parties.

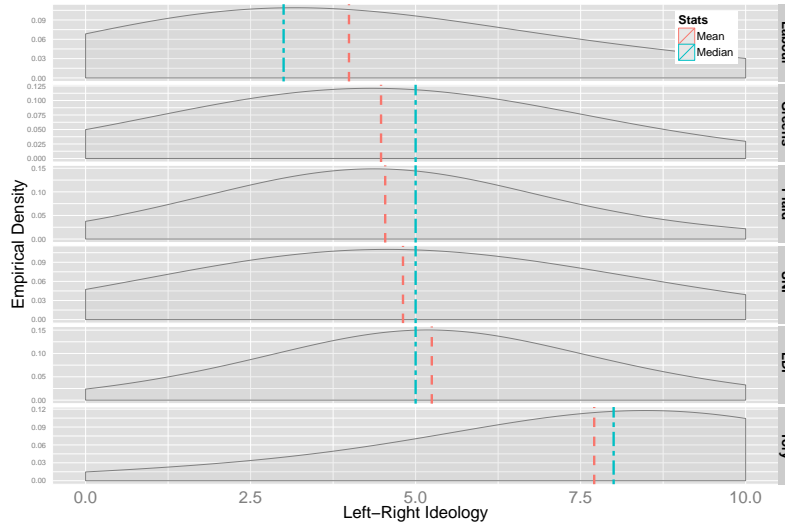
To make this plain, we further calculate the perceptual agreement score (PA) developed by Cees van der Eijk (2001) to capture the degree of agreement in voters' perceptions of parties' policy positions. The results are presented in Table 7. PA is designed for ordered rating scales such as the left-right scale, and it builds on the distribution of respondents' placement of parties on the scale. The PA coefficient ranges from -1 to 1, where -1 indicates a perfect bimodality, meaning an extreme disagreement among respondents, and 1 represents perfect unimodality, indicating a full agreement among respondents. A value of 0 captures a uniform distribution, which shows no agreement at all. As shown in Table 7, these scores are significantly smaller for the British Parties than the French and Dutch, with one major exception: the French UMP – the position of which is completely unclear to French voters.

Table 7: Perceptual Agreement of Party Ideology


| | | | | | | | | | | |
|-------------|--------|--------|-------|------|-----|------|-----|-----|-----|-----|
| Netherlands | SP | GL | PvdA | PvdD | D66 | CU | CDA | SGP | PVV | VVD |
| | .6 | .54 | .48 | .51 | .57 | .54 | .55 | .35 | .48 | .55 |
| France | Greens | PS | PRG | PR | NC | UMP | FN | | | |
| | .73 | .56 | .43 | .08 | .59 | .14 | .66 | | | |
| UK | Labour | Greens | Plaid | SNP | LDP | Tory | | | | |
| | .31 | .4 | .44 | .3 | .4 | .47 | | | | |

Figure 8: Knowledge on Party Left-Right Ideology

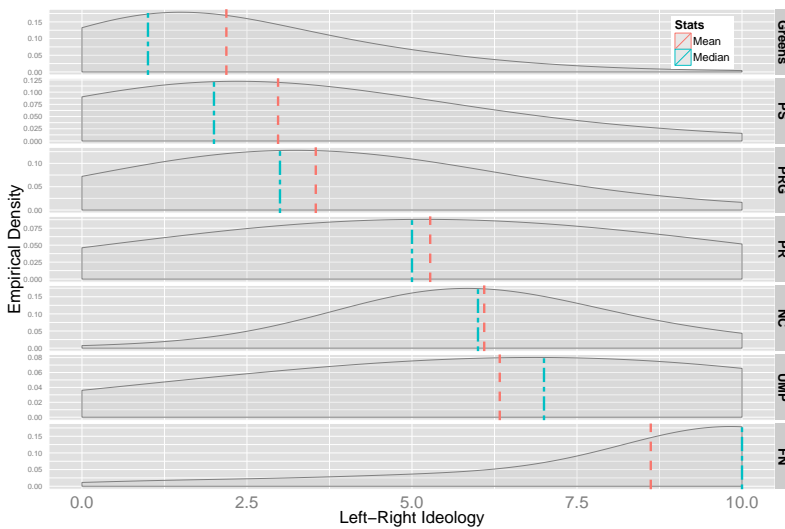
(a) United Kingdom



(b) Netherlands



(c) France



Section Summary

In this section, we have presented a lot of data from three original surveys designed to get at the extent and structure of core political knowledge among citizens in the UK, France, and the Netherlands. To try to highlight our broad conclusions that apply across countries as well as to highlight the differences in core knowledge between the countries, Table 8 summarized our main conclusions from this section.

Table 8: Summary of Political Knowledge in the UK, the Netherlands, and France

| Knowledge Dimension | UK | Netherlands | France |
|-------------------------|----------------------------------|-------------|--------|
| Party Recognition | High | High | Low |
| Party of the PM | Medium | Medium | High |
| Cabinet Composition | Medium | Low | – |
| Governing History | Only Gov./Opp. party distinction | | |
| Joint Governing History | – | Very Low | – |
| Party Size (relative) | High | High | Low |
| Party Size (exact) | High | High | Low |
| L-R Ideology (relative) | Low | High | High |
| PA on Party Position | Low | Medium | Medium |

As the table suggests, we can make one broad conclusion that seems to apply to all three countries: that history can provide a rough guide to the parties that are or are not potentially cabinet members, but does not provide much beyond that about likely cabinet compositions. All our other conclusions, however, vary by context. In the UK, our data suggest that the typical citizen can recognize the parties, and has a reasonable grasp of which parties are in the cabinet. She is also quite clear on party sizes but does not have a clear understanding of the relative positions of the parties. If that is so, then the literature on the use of heuristics in coalitional systems would suggest that British voters would have no trouble assigning responsibility for policy outcomes either retrospectively or prospectively (since they have information about cabinet roles and party sizes), but might have trouble using simple heuristics to predict which cabinet would form and to thus cast strategic votes aimed at shaping cabinet composition. This seems in line with the fact that the British voter has limited experience with the politics of cabinet formation and little reason previous to 2010 to even think about strategic voting over coalition cabinets.

We can contrast this with the Dutch voter who is in a good position to cast strategic votes over coalitions as well as to forecast the policy outputs of cabinets. Dutch voters have the best knowledge among our cases of the ideological positions of parties and also know their sizes and roles in cabinet. Finally, the French do not seem to know be able to recognize their parties as well as voters in the other two countries and they do not have a good idea of the sizes of the parties. That said they all know which party the President and PM come from (which happened to be the Socialists, the best recognized party as well) and understand the relative ideologies of all parties except the UMP. An obvious conclusion is that this pattern of knowledge is compatible with (and perhaps stems from) a heavy emphasis on presidential politics in France and the importance of the runoff system in elections. Such an emphasis enhances the importance of individuals (the Presidential contenders) over parties, encourages a seemingly never ending set of electoral coalitions that may obscure the distinctiveness of party labels, and may encourage elections towards being referenda on the performance of the presidential party.

The Dimensionality of Political Knowledge within Countries

Now that we have explored the extent and structure of political knowledge across our five different categories, we now turn our attention to the combined structure of these different knowledge types. That is, we want untangle the distribution of knowledge across individuals and categories. This investigation will allow us to determine whether certain types of knowledge are more or less likely to covary with other types.

There are several reasons why this investigation is important, but we will briefly discuss just three, omitting the obvious, normative reasons. First, if certain types of types of knowledge are more or less likely to covary, then we may conclude that there are multiple dimensions to political knowledge. The discovery of this type of dimensionality in political knowledge would imply that different people invest (or otherwise come to possess) different types of political information. This, in turn, would mean that different types of people are better or worse for different types of

political tasks. For instance, a person with high levels of ideological knowledge, but very low levels of historical knowledge, would be better suited to cast a coalition-directed prospective vote by generating coalition expectations according to the Fortunato & Stevenson (2014) framework, which requires ideological knowledge, than the Armstrong & Duch (2010) framework, which requires historical knowledge of co-governance. Likewise, a person that has little historical knowledge and little ideological knowledge, but perfect knowledge of the current roles of parties, would be well-suited to cast a performance-oriented retrospective vote, but ill-suited to cast a prospective vote. This means, of course, that our theoretical models of voting are not, and cannot be considered, “one-size-fits-all” explanations.

On a second, related note, understanding the dimensionality of political knowledge can provide insight into the process political socialization or political attentiveness. If we discover a high dimensional structure in all contexts, then we may conclude that people come to possess different types of political knowledge in different ways. If we discover a low dimensional structure in all contexts, then we may conclude that people simply come to possess political knowledge generally, and it is the depth, rather than the type, of knowledge, where the meaningful distinctions lie. Finally, if we discover that dimensionality varies by context, then we would be alerted to the risks of designing cross-national surveys with uniform political knowledge batteries.

Third, the dimensionality of political knowledge can also serve as a more pragmatic guide for the types of data we wish to collect moving forward. It is perfectly clear that any data collected for the purpose of evaluating retrospective voting, must include questions regarding the current roles of parties (and likely size as well), just as it is equally clear that any data collected to study prospective voting must include questions regarding the ideological positions of parties. However, if we merely wish to measure political knowledge, discovering the dimensionality of political knowledge can be our guide to determining which type of knowledge questions to include (i.e., recognition, role, historical role, ideology, and/or size).

Given the importance of understanding the dimensionality of political knowledge, how, then, should we measure it? One option is to simply compare how the various categories of political knowl-

edge correlate to one another. If the correlations are positive, statistically robust, and of comparable magnitude, then we may conclude that the different categories are all tapping the same underlying dimension. This is the approach we take in the tables below.

Below, we simply calculate each respondent's unweighted mean level of political knowledge in each category. The process is fairly simple, but deserves explanation nonetheless. First, following from the discovery above that knowledge more categorical than cardinal, we reorganize all cardinal knowledge questions into series of categorical knowledge responses. For example, for the history question, we establish a true rank-ordering of the history of cabinet participation for all parties in each country. We then use the respondents' estimates of each individual party's history of cabinet participation to derive a series of pairwise comparisons. For example, say there are three parties (A, B, and C) where the true rank-order of cabinet participation is (A>B>C), that is, A has been in government more than B, and B has been in government longer than C. This number of parties yields 3 unordered pairs of correct rank-ordering (AB, AC, BC), that may be ranked either correctly or incorrectly by each respondent. A respondent that ranks two pairs correctly would be given a score of 0.667, for example. This process is repeated for size and ideological knowledge. The recognition and role questions, already categorical, need no such transformation, merely classification as correct or incorrect (for the purposes of this section, we do not distinguish between non-response and incorrect, as both indicate a lack of knowledge). These un-weighted means are then correlated to each other in each sample. The results are presented in Table 9 below.

First, we note that each correlation in each matrix is positive and each correlation is differentiable from 0 at $10e^{-15}$. Second, we note that, given the different number of responses in each category, across all contexts, a correlation of 1, or approaching 1, is virtually impossible. That said, the correlations and their robustness are still quite strong and seem to imply that political knowledge is fairly unidimensional.

But we can do better than correlation matrices. There exists a large psychological and political science and literature on data dimensionality and stimuli response and we would be remiss to fail to take advantage of those advances here. Thus, we borrow from the literature on the analysis of roll call

Table 9: Correlation between Political Knowledge Variables

| | | Recognition | Role | History | Ideology |
|-------------|----------|-------------|------|---------|----------|
| UK | Role | .521 | – | – | – |
| | History | .466 | .596 | – | – |
| | Ideology | .339 | .523 | .395 | – |
| | Size | .482 | .615 | .658 | .456 |
| Netherlands | Role | .344 | – | – | – |
| | History | .311 | .587 | – | – |
| | Ideology | .378 | .624 | .515 | – |
| | Size | .380 | .555 | .603 | .526 |
| France | Role | .275 | – | – | – |
| | History | .291 | .299 | – | – |
| | Ideology | .373 | .384 | .493 | – |
| | Size | .360 | .358 | .692 | .599 |

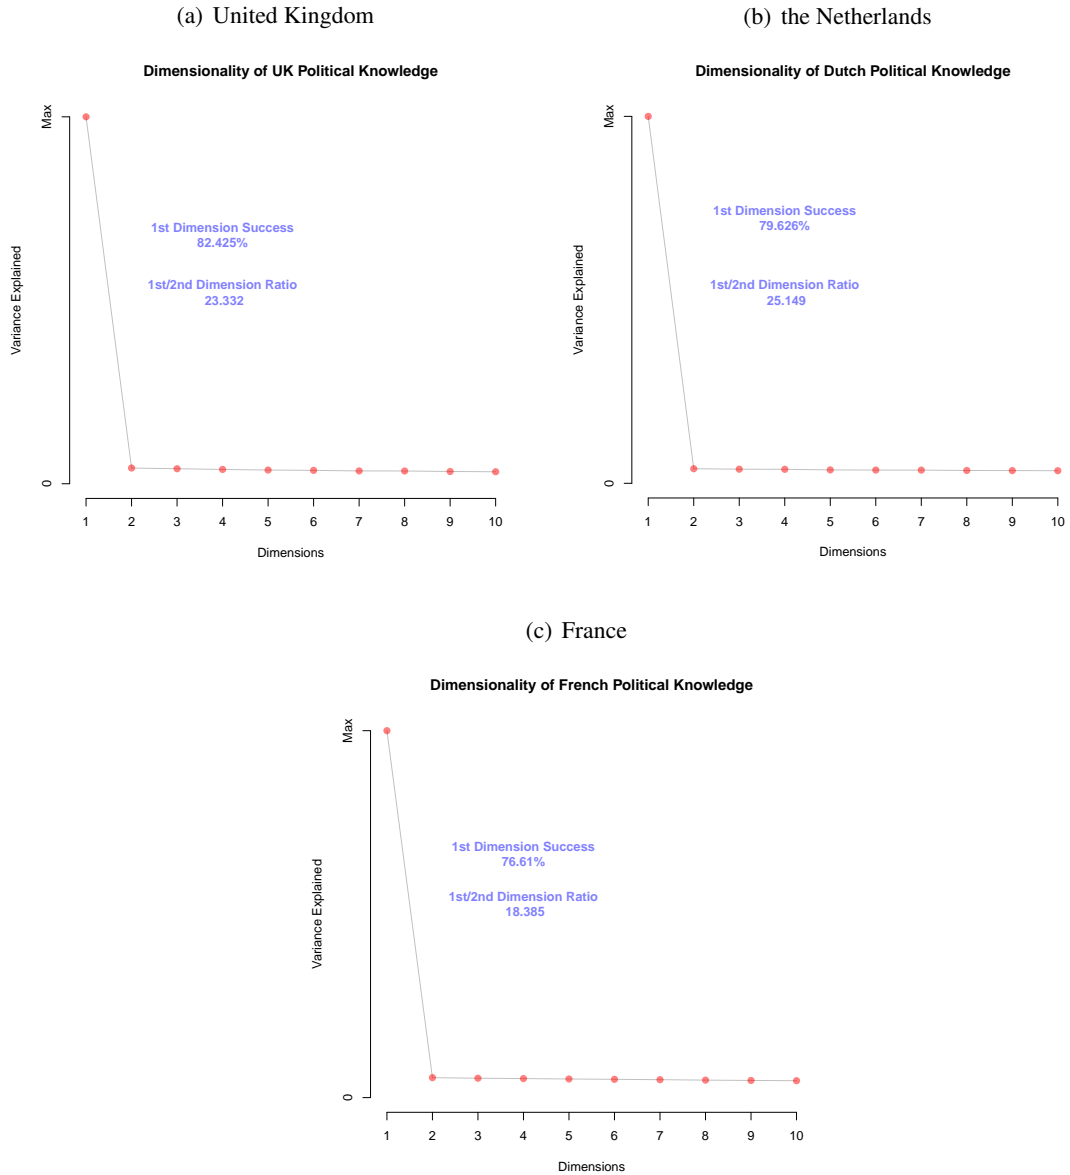
voting in the US Congress to assess the dimensionality of political knowledge across our categories.

As before, we rearranged all questions into a vector of ordinal responses where the respondent can give either a correct or incorrect response. This results in a matrix of 0,1 responses of the dimensions $n \times m$, where n is the number of survey respondents and m is the number of unique stimuli. This $n \times m$ matrix is then translated into an agreement score matrix (a simple estimate of how similar each respondent is in their responses) and we perform eigenvalue decomposition on the matrix to assess the level of dimensionality in each set of responses. The results of these decompositions are demonstrated in Figure 9. Each panel shows the ability to classify every response in a single dimension as well the ratio of first to second dimension variance explained. Given the relatively small number of stimuli, both of these figures are very high. Indeed, his analysis suggests that the structure of political knowledge, at least in terms of the categories discussed here, is overwhelmingly unidimensional *withincountries*. To be clear, each of these analyses are within country only, and, though, the structure of political knowledge is unidimensional within country, this does obscure the very strong cross-national differences in knowledge structures we discussed above.

Given that the structure of political knowledge is unidimensional across our categories within countries, a fair question to ask is, how well do our measures, which were specifically selected to

measure the types of knowledge needed to vote or perform other behavioral tasks associated with democratic participation, correlate to the types of material that we typically see on political knowledge batteries?

Figure 9: Dimensionality of Political Knowledge in the UK, the Netherlands, and France



To address this question, we assess three survey modules from the most recent CSES study: the 2010 Netherlands module and the 2005 and 2009 German modules. We selected these modules because a) we want to compare apples with apples between our three surveys and these three, thus

the Dutch case is ideal and the German case is quite similar and b) they exemplify what we think of both very useful and practically useless knowledge modules. More specifically, the 2009 German knowledge battery asks question about the institutional structures of German politics – pieces of information that are both enduring and pragmatic for voters to possess. The three questions ask citizens the minimum electoral threshold parties must surpass to gain seats in the lower house (5%), which of their two votes determines the overall seat distribution in Germany’s mixed member electoral system (their PR, or “second,” vote), and in which German elections are non-German European citizens allowed to vote in (municipal, or “communal,” elections).

In contrast, the Dutch battery utilized photographic knowledge questions, asking respondents to identify photographs or MPs or party leaders — knowledge that is relatively useless to understanding political processes or outcomes. The battery for Germany 2005 is perhaps even worse. It asks voters to attribute three campaign proposals (or promises) to party groups, only one of which came from a member of cabinet. This type of knowledge is fleeting, perhaps even arbitrary or idiosyncratic to the audience (i.e., PDS supporters should score higher as two of the proposals came from the PDS, and, indeed, there is a robust empirical relationship between PDS support and correct answers on these two questions). Thus, our expected rank-ordering for how well our left-right knowledge scales predict performance on the batteries would be: Germany 2009 > Netherlands 2010 \geq Germany 2005.

To assess the correlation between our measures and these types of political knowledge, we simply assess how well our scaled measure of partisan ideological knowledge predicts these knowledge questions. Thus, for each survey, we rearrange the left-right party placements into a vector of pairwise comparisons where the respondents can either rank-order the party pair correctly or incorrectly. This matrix of 0,1 data is then decomposed and scaled parametrically using the WNOMINATE (Poole 2005) estimation algorithm. This scaling procedure produces a set of “ideal points,” or, a set of [-1, 1] party-knowledge values, where greater numbers indicate higher levels of knowledge.

These party knowledge estimates are then used to predict the voters’ responses on each survey’s knowledge battery. Each battery had three items and the only independent variables used were the party knowledge estimate and a binary indicating whether or not the respondent declined to place all

parties - this was included to capture a propensity for non-response, as there are several “don’t know” responses in each of the knowledge batteries as well as the party placements. After each model was estimated, a summary of its fit was recorded. This summary is simply the proportion of probability correctly allocated across response categories for each observation in each model. You can think of this as the likelihood, the complement to absolute residual, etc. The measure is, precisely, $1 - |y - \hat{y}|$, thus, 1 indicates a perfect prediction and 0 indicates the worst possible fit. Table 10 shows the results for each survey, and Table 11 summarizes the difference in fit between each pair.

Table 10: Model Fit Summaries

| Survey | Mean | Median | SD |
|-------------|------|--------|------|
| Netherlands | .626 | .627 | .113 |
| Germany '05 | .608 | .600 | .112 |
| Germany '09 | .653 | .640 | .145 |

Table 11: Comparing the Predictive Ability

| Comparison | Difference | <i>t</i> Statistic |
|---------------------------|------------|--------------------|
| Germany '09 - Germany '05 | .046 | 19.504 |
| Germany '09 - Netherlands | .028 | 12.012 |
| Netherlands - Germany '05 | .018 | 8.813 |

Given the questions asked in each module, the data suggest that measuring left-right party knowledge, as we have here, correlates much more closely to structural, institutional knowledge than to temporary knowledge of policy promises or knowledge regarding individual political personalities.

Concluding Remarks

There are several important takeaways from our exploration of political knowledge across France, the Netherlands, and UK. The first is that there is tremendous variation in the types of knowledge voters possess across contexts. This variation, in turn, implies that the set of heuristics voters use to understand political processes or make political decisions should also vary dramatically from context to context. Compare the Dutch, who have very high levels of party recognition, ideology, and size

knowledge, but very little knowledge of party role, to Britons, who have high levels of party recognition, size, and role knowledge, but very low levels of party ideology knowledge. Clearly, the Dutch are much better suited to the task of prospective voting, while the Britons are much better suited to vote retrospectively. Or consider the case of the French, who have very high levels of partisan ideological knowledge when they recognize parties, but on average, recognize far fewer parties (proportionately) than the British or the Dutch. This is precisely what we would expect given France's unique electoral institutions that favor pre-electoral coalitions and lead to several parties abandoning several districts in legislative elections, and run-off presidential elections (which were previously non-concurrent to parliamentary elections), as Cox (1997) points out, discourage national linkage. Indeed, the structural variations in political knowledge reinforce our previous findings that demonstrate the institutional roots of partisan left-right knowledge (Fortunato & Stevenson 2014).

Another salient takeaway from this exploration is that the data suggest the most basic informational assumptions in our voting models are very unlikely to hold for large swaths of voters. The data suggest that many voters are unlikely to know the partisan composition of their cabinet and that almost all voters are unlikely to know the relative histories of cabinet participation or co-governance. Thus, many voters are ill-equipped to cast outcome oriented retrospective voters or formulate perceptions of the likelihood of various cabinet formations given the histories of co-governance of the relevant parties. Further, our analysis shows that the informational commonality assumptions of many strategic voting models (i.e., Cox 1997) are, at least in reference to left-right positioning, unlikely to hold in many contexts. This does not mean that our individual-level theoretical models of voter behavior are flawed, rather, that the way we derive aggregate expectations from individual predictions must respect these variations in individual information sets.

Finally, we hope that these explorations can serve as a guide to future electoral survey construction. These findings, as well as our previous work imply that many empirical models of retrospective voting have underestimated the retrospective voting mechanism by making informational assumptions that did hold (the cabinet composition assumption; Fortunato & Stevenson 2013). We believe that the data we present here will encourage to survey curators to consider the types of behavioral

models their survey data will test and craft political knowledge batteries that accurately assess the informational prerequisites of those the models. More specifically, we hope that each survey will continue to ask respondents to place their parties on the left-right ideological spectrum, but also add questions asking respondents to identify party roles.

References

- Armstrong, David A. & Raymond M. Duch. 2010. "Why can Voters Anticipate Post-election Coalition Formation Likelihoods?" *Electoral Studies* 29(3):308 – 315.
- Cox, Gary W. 1997. *Making Votes Count: Strategic Coordination in the World's Electoral Systems*. Cambridge University Press.
- Downs, Anthony. 1957. *An Economic Theory of Democracy*. Harper and Row, New York.
- Duch, Ray & Randolph T. Stevenson. Forthcoming. "Responsibility Attribution for Collective Decision Makers." *American Journal of Political Science* .
- Fortunato, David & Randolph T. Stevenson. 2013. "Perceptions of Partisan Ideologies: The Effect of Coalition Participation." *American Journal of Political Science* 57:459–477.
- Fortunato, David & Randolph T. Stevenson. 2014. "Heuristics in Context." *Manuscript, Rice University* .
- Fortunato, David, Randolph T. Stevenson & Andrew Spiegelman. 2011. What do Voters Know about Cabinet Formation. In *Annual Meeting of American Political Science Association, Seattle Washington*.
- Fortunato, David, Randolph T. Stevenson & Greg Vonnahme. 2013. Context, Heuristics, and Political Knowledge: Explaining Cross-National Variation in Citizen's Left-Right Knowledge. In *Annual Meeting of American Political Science Association, Chicago Illinois*.
- Gigerenzer, Gerd. 2010. "Moral Satisficing: Rethinking Moral Behavior as Bounded Rationality." *Topics in Cognitive Science* 2(3):528–554.
- Green, Donald P., Bradley Palmquist & Eric Schickler. 2002. *Partisan Hearts and Minds: Political Parties and the Social Identities of Voters*. Yale University Press.
- Koch, Jeffrey W. 2001. "When Parties and Candidates Collide: Citizen Perception of House Candidates' Positions on Abortion." *Public Opinion Quarterly* 65(1):1–21.
- Lau, Richard R. & David P. Redlawsk. 2006. *How Voters Decide: Information Processing in Election Campaigns*. Cambridge University Press.
- Poole, Keith T. 2005. *Spatial Models of Parliamentary Voting*. Cambridge University Press.

- Powell, G. Bingham, Jr. & Guy D. Whitten. 1993. "A Cross-National Analysis of Economic Voting: Taking Account of the Political Context." *American Journal of Political Science* 37(2):pp. 391–414.
- Rahn, Wendy M. 1993. "The Role of Partisan Stereotypes in Information Processing about Political Candidates." *American Journal of Political Science* 37(2):pp. 472–496.
- Van der Eijk, Cees. 2001. "Measuring Agreement in Ordered Rating Scales." *Quality and Quantity* 35(3):325–341.