Can (and do) Voters Use Simple Mental Models to Forecast Cabinet Composition? Experimental Evidence from Ten Parliamentary Systems

David Fortunato
ORCID: 0000-0003-2200-4080
Associate Professor
School of Global Policy and Strategy
University of California, San Diego
San Diego, CA, USA
dfortunato@ucsd.edu

Seonghui Lee
ORCID: 0000-0003-4562-2121
Associate Research Fellow
Institute of Political Science,
Academia Sinica
Taipei, Taiwan
seonghui.lee@gmail.com

Randolph T. Stevenson
ORCID: 0000-0003-2565-1355
Radoslav Tsanoff Professor of Public Affairs
Department of Political Science
Rice University
Houston, TX, USA
randystevenson@rice.edu

Abstract

Elections in parliamentary democracies rarely yield single-party majorities, requiring voters to form expectations about potential governing coalitions to cast outcome-oriented votes—a task often dismissed by scholars as too difficult. This paper proposes that voters overcome this challenge using an ecologically rational mental model that relies on simple cues derived from party sizes and ideologies. Using original survey experiments across ten countries and more than 20,000 subjects, we demonstrate that voters expect coalitions that include the largest party as prime minister, are of "winning" size, and are ideologically compatible. Our theoretical model deterministically predicts survey respondents' expectations in up to 41% of cases. Furthermore, the weights voters assign to these cues vary contextually, mirroring the actual empirical regularities of cabinet formation in their respective countries. Ultimately, we provide robust experimental evidence that many voters can, and do, generate sensible coalition expectations using only a few simple cues.

> *"...each voter's ballot does not support the policies of any one party. Instead it supports the whole coalition that party joins. Thus the meaning of a vote for any party depends upon what coalitions it is likely to enter..."*
>
> Downs (1957, 162-163)

Elections in parliamentary democracies rarely produce single-party majorities, creating a problem for voters seeking to influence government composition: casting "coalition-directed" ballots requires voters to have expectations over the likelihood of different potential cabinets forming. Without such expectations, citizens cannot use their vote to increase the chances of their preferred cabinets taking power or calculate the policy implications of their vote. However, inferring "coalition probabilities" appears intractible. Coalition formation is often complex and opaque, and the number of potential cabinets routinely runs into the thousands, making the mere enumeration of possibilities, let alone inferring their likelihoods, extraordinarily difficult.

For decades, this presumed difficulty led researchers to dismiss the possibility that voters could forecast coalition outcomes. Downs famously argued that, faced with doing so, "...each voter either abstains, votes after cutting off his deliberation at some unpredictable point, or decides it is easier just to vote for his favorite party" (1957, 163). Decades later, Powell (2000) clarified the troubling normative significance of this conclusion: if voters cannot infer which governments are likely to form, their ability to issue prospective mandates or coordinate on viable alternatives to underperforming incumbents is severely hindered, degrading democratic responsiveness and accountability.

We propose that voters can solve this problem using an abstract mental model of coalition formation. Building on interdisciplinary research on the ecological rationality of heuristics (e.g., Gigerenzer and Sturm 2012; Todd and Gigerenzer 2012), including recent applications in political science (e.g., Fortunato and Stevenson 2019, 2021; Fortunato et al. 2021), the model posits that voters infer coalition probabilities using observable characteristics ("cues") of potential cabinets that depend on only two features of party systems: relative party sizes and ideological positions. The model suggests voters aggregate these cues using three simple rules: 1) the largest party will provide the prime minister; 2) the coalition will be of "winning" size; and 3) to reach a winning size, the PM party will (if necessary) coalesce with "ideologically compatible" partners.

Given the ubiquity of these factors in real-world multiparty government formations and the widespread availability of information on party size and ideology, we expect this model to apply across contexts. However, ecologically rational models of heuristic inference depend on empirical regularities linking cues to the

inferential target (cabinet formation probabilities) *within a specific context*. Thus, we hypothesize that the weights voters assign to different cues will vary contextually in response to differing local empirical regularities associated with cabinet formation.

To test these arguments, we conducted three series of original survey experiments that present subjects with hypothetical parliamentary compositions (randomized party sizes and ideological positions) and elicit their expectations about which cabinet would form. The first was a two-trial experiment with five fictional parties embedded in a multi-purpose, nine-country survey. The second was a 12-15 trial version of this same experiment that used five to nine parties across three countries. The third replaces the fictional parties with real parties.[1] In total, we analyze over 80,000 coalition expectations from over 20,000 subjects in ten countries: Denmark, France, Germany, Hungary, Italy, the Netherlands, Poland, Spain, Sweden, and the United Kingdom.

These experiments build on previous empirical efforts to explore coalition expectations using observational data that argue that expectations are formed via three (non-exclusive) mechanisms: 1) adaptive expectations based on historical trends; 2) message-based expectations derived from elite signals or media; and 3) model-based expectations utilizing a mental model of the world to transform inputs into inferences (e.g., Duch and Stevenson 2008; Eberl and Plescia 2018; Bowler, McElroy and Müller 2022). These studies have established robust associations between coalition expectations and the size and ideological characteristics of those coalitions, media or elite signals about the likelihood of different potential cabinets, and (in one unpublished study) patterns of historical cabinet compositions.[2] Together, they reveal that the largest and most robust correlations are for size and ideological variables.

Recognizing the difficult inferential challenges that such observational studies face—e.g., is a potential coalition's size and ideology a cause or consequence of its history of formation? Are media/elite messages about a potential cabinet's chances driven by its size and ideology?—scholars in this area have also begun to use survey experiments to try to sort out the causal impact of one of these factors: media/elite messages (Falcó-Gimeno and Muñoz 2017; Bahnsen, Gschwend and Stoetzer 2020). The results are revealing, since they provide much less support for the causal impact of media/elite messages on coalition expectations than the correlational evidence would suggest.

In the appendix we provide a more thorough review of this empirical literature and argue that it motivates

[1] We discuss these experiments, which allow us to address external validity concerns raised by using fictional parties in the main experiments, in detail in Appendix 9 and summarize them in the section on external validity below.

[2] Armstrong and Duch (2010) is often cited in the literature on coalition directed voting in support of the idea that voters use history to generate coalition expectations. That paper, however, does not (and did not intend to) test that hypothesis—providing only a review of historical patterns of cabinet formation.

the need for a study, such as ours, that can provide a clean answer to the causal question: Can voters use cues about parties' sizes and ideologies to generate sensible coalition expectations in a setting in which we can definitively rule out the influence of historical patterns of cabinet formation and media/elite messages? And, if so, can we then generalize that finding to more realistic settings in which voters can substitute size and ideological cues for historical and mediated alternatives?

To preview, our results provide unambiguous evidence that many of the respondents in our experiments possess an abstract mental model of coalition formation that generates sensible expectations using only party size and ideology. Indeed, up to 41% of our subjects expect coalitions to form that are consistent with our proposed model. Further, we find these effects of size and ideology consistently across the all of our experiments and even when we replace fictional parties with real ones, providing evidence that the results of our purely hypothetical experiments can teach us about to real-world environments.

In the remainder of this manuscript, we briefly review why coalition expectations are central to the large literature on coalition-directed voting, describe our heuristic model of these expectations, derive several testable hypotheses, detail our research designs, analyze the results, and finally conclude with a discussion of our study's implications.

Coalition-directed voting and coalition expectations

Coalition-directed voting is the act of casting a ballot for an individual party with the intention of increasing the likelihood that one's preferred coalitions come to power.[3] To do this optimally, one must consider all the coalitions that each party could join, their associated formation probabilities, and how these probabilities are likely to change if one votes for each party (or not). For example, Duch and Stevenson (2008) formalize the influence of coalition expectations in economic voting, finding that if voters expect a party to be excluded from *all* viable cabinets, a coalition-directed vote for that party is "wasted" and thus uncast. Conversely, if a party is certain to enter government, maximizing coalitional impact requires voting for a different party. Like Duch and Stevenson (2008), every model of coalition-directed voting of which we are aware assumes that voters know or can infer coalition probabilities (e.g., Gschwend 2001, 2007; Kedar 2005; Duch, May and Armstrong 2010; Indridason 2011; Meffert et al. 2011).

These models represent a clear departure from the empirical literature's previous consensus that voters

[3]Many theories of coalition-directed voting are policy-motivated (e.g., Kedar 2005; Duch, May and Armstrong 2010), but not all (e.g., Gschwend 2001, 2007; Duch and Stevenson 2008). Because coalition expectations are required for both policy and non-policy oriented coalition-directed voting, our discussion will focus on the more general idea that outcome-oriented voting in multiparty democracies in general requires voters to form coalition expectations.

could not calculate or intuit the coalition expectations necessary for coalition-directed voting.[4] This change began in the early 2000's (e.g., Gschwend 2001, 2007; Kedar 2005), as some researchers developed a new approach to studying coalition-directed voting empirically: assume voters have well-formed coalition expectations, produce estimates of those expectations from a stipulated rationalist model, assign these estimated expectations to respondents in an election survey, predict their votes under the model, and examine if actual voting behavior is consistent with the model's predictions.

For example, Duch, May and Armstrong (2010) assume that voters have a mental model of coalition formation mirroring political scientific statistical models of coalition formation. They then estimate coalition expectations for their survey respondents from the Martin and Stevenson (2001) empirical model of cabinet formation and use those estimates to calculate how a rational voter *should* have cast a coalition-directed vote given those coalition expectations and a Gamsonsian policy-making model. Influential work by Kedar (2005) and Duch and Stevenson (2008) take a similar approach which we refer to as 'Voter as Political Scientist" (VaPS) models.

This shift to assuming voters hold VaPS-level expectations, however, occurred without any evidence refuting the previous consensus. Instead, scholars relied on an "as-if" justification for the VaPS model, tolerating tenuous assumptions about expectations as long as they produced empirically verified downstream predictions about voting behavior (and they did).

Therefore, one goal of this manuscript is to evaluate whether voters can infer cabinet probabilities that conform to political scientists' rationalist assumptions (without relying on coalition histories or elite messages). If the answer "yes," theorists can confidently assume that voters do have the ability to infer coalition probabilities and build those assumptions into rational models of coalition-directed voting.[5] That said, exploring the usefulness of the VaPS model of coalition expectations is not the primary goal of this paper. Instead, in the next section, we advance our own model of coalition expectations, which is rooted in the interdisciplinary literature on ecologically rational heuristics, and is much simpler than the VaPS model (e.g., Gigerenzer and Sturm 2012; Fortunato and Stevenson 2019; Fortunato et al. 2021).

[4]Many theorists were undeterred by Downs' skepticism, however. See Austen-Smith and Banks (1988).

[5]Beyond that, if our results can be generalized from the finding that voters *can* form sensible coalition expectations using size and ideology, to the assertion that they actually do so "in the wild" (likely in combination with other factors like history and media messages) then our study also speaks directly to scholarship that seeks to understand the way voters use all the information available to them to form expectations.

A heuristic model of coalition expectations

Heuristics are effort-reducing strategies for inference based on limited information. They rely on cues (here, observable characteristics of potential cabinets), cue-weights, and aggregation rules. A heuristic model is "ecologically rational" if it meets four criteria: 1) it uses simple aggregation rules (e.g., linear additive or lexicographic "take-the-best" rules); 2) it discounts costly and complicated cues; 3) it avoids redundant or uninformative cues; and 4) it produces accurate inferences, at least relative to its cost and the cost of available alternative models.

We build a core model of coalition expectations based on accessible cues, that are strongly associated with government formation across contexts, and present in all government formation episodes. The theory of ecologically rational heuristics then produces the testable hypothesis that these cues impact inferences proportionally to their real-world empirical associations with cabinet formation outcomes, discounted by the cost of learning and/or calculating their values.[6]

While we do not comprehensively review the cabinet formation literature here—see Martin and Stevenson (2001, 2010), Müller and Strøm (2003), and De Marchi and Laver (2020, 2023)—we can quickly summarize the characteristics of potential cabinets that political scientists have found to be predictive of cabinet composition by relying on Kayser et al.'s (2023) recent comprehensive review and reanalysis of cabinet formation in 31 democracies.[7] These variables fall into six classes:

1. size characteristics (e.g., minority/minimal-winning/surplus majority or number of parties)
2. ideological compatibility (e.g., divisions within coalition or opposition)
3. bargaining power (e.g., plurality and median status of member parties)
4. electoral performance (e.g., changes in seat share)
5. coalition history (e.g., incumbency and record of co-governance)
6. elite signals (e.g., pre-electoral pacts)

[6] See the definition of "validity" in Gigerenzer (2019, 3559) or the derivation of the "ecological rationality hypothesis" in Fortunato and Stevenson (2019, 313-315).

[7] Kayser, Orlowski and Rehmert (2023) do not account for recent multidimensional agent-based models (De Marchi and Laver 2020, 2023). While these models improve our understanding of coalition formation and offer substantial promise for future revelation, it is presently unclear if they reveal new empirical associations between observable cues and coalition formation that voters may find useful to their expectation formation process. We discuss these issues in more detail in the appendix.

Categories 1-3—size, ideological compatibility, and inclusion of parties with strong bargaining power—are universal to coalition formation and consistently associated with cabinet formation (Martin and Stevenson 2001; Kayser, Orlowski and Rehmert 2023) and prime ministerial selection (Glasgow, Golder and Golder 2011). Specifically, the most predictive characteristics of cabinets are ideological compatibility, minimal-winning status, and inclusion of both the largest party (which enjoys formateur/informateur privileges) and the median party. For prime ministerial selection, plurality and median status are the only consistently predictive variables across contexts. Crucially, variables in categories 1-3 are directly observable or easily calculable from party size and position—information voters know well (e.g., Fortunato, Stevenson and Vonnahme 2016; Lee, Haime and Stevenson 2019, respectively). Our heuristic model is built upon these variables.

We exclude categories 4-6 from our heuristic model on theoretical grounds. Category 4 (electoral performance) is uncorrelated with coalition formation (Kayser, Orlowski and Rehmert 2023). Categories 5 and 6 (coalition history and elite signals) represent the prevailing non-rational explanations of expectation formation (adaptive or message-based) and are intentionally held constant in our research design.[8]

Core model and expectations

Our heuristic model generates expectations regarding which cabinet will form and which party will provide the prime minister (PM). We begin with PM selection because most real-world governments are led by a formateur party that secures the premiership.

Research on PM selection distills context-free predictors to two reflections of bargaining power: size and ideological centrality. The empirical literature unanimously identifies plurality status as the strongest predictor of PM selection (Bäck and Dumont 2008; Glasgow, Golder and Golder 2011). In the ParlGov database of 1,031 governments formed in 36 post-war parliamentary democracies (Döring and Manow 2021), the largest party secures the premiership in 78% of governments, making plurality status a highly efficient predictor given that the PM is chosen from an average of 7 parties.[9]

Controlling the median legislator is also predictive, but less so, and we expect voters to rely less on this cue for three reasons. First, plurality strongly outperforms median status in predicting real-world outcomes

[8] We include the negative association between pariah parties and coalition formation in these excluded categories, as this association derives from history and concurrent anti-pacts. Related factors like incumbency or pre-electoral pacts may supplement mental models "in the wild" (see Eberl and Plescia 2018; Bowler, McElroy and Müller 2022).

[9] Restricting to more complex party systems, say, those of 6+ parties where the plurality party is not a majority (which corresponds to 562 of the 1,031 governments in the ParlGov archive) the coalition that ultimately forms contains the largest party as PM 69% of the time.

and frequently discriminates to a single alternative, making consideration of median superfluous. Second, median status requires aggregating size and position information, increasing its relative cognitive cost. Third, the median voter theorem is an academic concept rarely discussed in popular media. As such, it is unclear that voter understand the concept. Confirming this, Fortunato et al. (2021) demonstrate that voters do not use median status to heuristically infer policy influence in multiparty governance. In contrast, plurality status is intuitive, widely discussed, and easily observable (see, e.g., Duch, Przepiorka and Stevenson 2015). While our design allows us to estimate the weight respondents give median status, our central prediction is that plurality will be the primary cue driving expectations about which party provide the PM:

H_1: voters will expect the largest party to provide the PM

Empirical models of government formation typically identify the constellation of coalescing parties without distinguishing cabinet roles, or identify the set of partners joining a specific PM (Martin and Stevenson 2001). Regardless of approach, consensus exists on the most important cues: ideologically cohesive potential cabinets are much more likely to form than divisive ones,[10] and majority coalitions, specifically minimum-wining majorities, are more likely to form than minority coalitions (though this varies in interesting ways, see Strøm 1990). In addition to being the most predictive, these cues are also cheap, intuitive, and likely to generate "good enough" expectations.

H_2: voters will expect coalition parties to be ideologically compatible

H_3: voters will expect the coalition to be of a winning size

We have not yet rigidly defined "compatible" and "winning." This is intentional as we lack guidance from existing research on how voters understand these concepts. Voters might conceptualize ideological (in)compatibility as total variance in party positions, the distance between the most dissimilar parties, average distance from the PM, etc. Similarly, voters might define "winning" as majority control, minimal-winning, or minimum-safe.[11] Voters may even use more exotic definitions, perhaps intuiting that a near-majority is "close enough." Rather than imposing ex ante definitions, we evaluate our model against various plausible conceptualizations of compatibility and winning. To be clear, each of these definitions of ideological

[10] While most models measure compatibility unidimensionally, recent work takes a multidimensional approach De Marchi and Laver (2020, 2023) and Martin and Vanberg (2025). We discuss multidimensionality in the appendix, but for experimental tractability, we do not derive multidimensional hypotheses.

[11] A "minimum-safe majority" is a majority that is large enough to tolerate expected ad hoc defection (Provins, Monroe and Fortunato 2022, 1701).

compatibility and winning are consistent with our hypotheses. Further, we expect heterogeneity in the use of these different concepts across voters. Consequently, in testing our hypotheses and evaluating the predictive accuracy of the model, we will consider all of them separately and in concert.[12]

There are two other covariates from the categories of size, compatibility, and bargaining power that have been consistently associated with the chances a potential coalition will form: whether it includes the largest party or the median party. However, their impacts depends on whether the focal party is PM or a junior partner. Including the largest party as PM positively impacts formation probabilities, but including it as a junior partner has a negative effect.[13] Median status has a generally positive (if modest) impact for both roles. Nevertheless, we doubt voters use median status as a cue for junior partners for the same reasons discussed regarding PM selection: it is less predictive, more costly, and less intuitive than other cues based on size and ideology (see Fortunato et al. 2021). That said, we will test this intuition by including median status in many of the empirical models presented below.

Finally, other bargaining power indicators, like minimum integer weights, influence cabinet formation and the distribution of ministerial portfolios (Laver, De Marchi and Mutlu 2011; Cutler et al. 2016). However, given the complexity of calculating these weights, we are highly skeptical that voters use them, a skepticism supported by Fortunato et al.s' (2021) finding that voters do not utilize such complex cues to infer policy influence.

In sum, our heuristic model comprises three cues: Which party is largest? Which parties are ideologically compatible? And, which combinations of parties are of winning size? Ecologically rational voters then aggregate these using a linear combination (with weights proportional to real-world associations, discounted by availability) or a "take-the-best" rule, sequentially evaluating cues based on predictive power until a single alternative emerges (e.g., Gigerenzer and Goldstein 1996, 1999).[14] In our between-subjects experiment—which elicits a single response for the most likely cabinet—both aggregation rules should produce similar empirical patterns.

[12]For example, while we will be exploring the predictive accuracy of our model when we measure ideological compatibility and winning using each combination of definitions, the best test of our model's performance asks how well it predicts expectations if voters form expectations using *any* of them.

[13]As no study models formation and PM selection simultaneously, these effects are derived by comparing models of coalition formation with models of partner selection to a stipulated PM (e.g., Martin and Stevenson 2001). We confirm these associations below.

[14]A take-the-best process could involve: 1) identifying the plurality party as PM; 2) if the PM alone is not winning, adding the nearest ideological neighbor; 3) if still not winning, adding the next nearest, and so on.

Contextual correspondence

While our core heuristic model should be widely applicable, relative cue-weights should vary contextually based on their empirical associations with the composition of cabinets in the real-world. Consider party system: $\{A_1^{36}, B_3^{12}, C_7^8, D_8^{12}, E_{10}^{32}\}$, where letters identify parties, subscripts denote left-right position, and superscripts denote seat share. If a respondent expects party A to be PM, they might then add only B (falling just shy of a majority but remaining ideologically compact) or instead add B and C (securing a majority but tripling ideological spread). The specific coalition they expect will depend, in an ecologically rational heuristic theory, on the context in which they form their expectations—and most importantly how the empirical predictors of real-world formations vary across those contexts.

For example, a German respondent, accustomed exclusively to majority cabinets, may prioritize majority status. A Danish or Swedish respondent, accustomed to minority cabinets, may prioritize ideological compactness. Generally, cue-weights in voter expectations should positively correlate with the real-world predictive power of those cues within their country, since an ecologically rational model is one in which individuals are "selecting building blocks that are adapted to [a] particular environmental structure" (Todd and Gigerenzer 2007, 169). Thus, we predict that, within country, the relative weights placed on cabinet characteristics in voters' coalition expectations should positively correlate to the relative weights placed on cabinet characteristics in real-world government formation.

> H_4: cue-weights in respondents' coalition expectations will be positively correlated to cue predictiveness in real coalition formation in their country

Research design

As discussed above, previous empirical studies of cabinet expectations primarily use observational approaches, eliciting expectations in pre-electoral surveys (e.g., Eberl and Plescia 2018; Bowler, McElroy and Müller 2022). This provides both high external validity and allows simple accuracy assessments by comparing expectations to actual outcomes. Likewise, experimental studies that randomize hypothetical treatments about real-world parties similarly maintain external validity while allowing clean estimation of a causal effect (e.g., Falcó-Gimeno and Muñoz 2017; Bahnsen, Gschwend and Stoetzer 2020).

However, both observational designs and experimental designs that use real parties have significant drawbacks. First, tractability constraints may require the researcher to limit the number or type of coalitions included in the study to a very small (presumably non-representative) number. Second, subjects may have varying levels of background knowledge (e.g., of coalition histories or pre-electoral pacts) that they bring

with them to any observational design, as well as to experimental designs that use real parties, making it difficult to isolate the effects of different potential drivers.

While this is a mundane observation about observational studies, it is important to make clear that it also applies to experiments using real-world parties. When treatment conditions conflict with a subject's own information about the likelihood of potential coalitions, they may well act on their private information instead of the treatment. Coupled with the need in such experiments to constrain treatments to be plausible in context (i.e., treated elite/media signals, seats, or ideological distributions that are implausible may be rejected), it can be difficult to know if an estimated treatment effect is due to the treatment or its coincidence with other factors (that the respondent brought with them to the experiment) that lead to the same expectations.[15]

We designed our experiment to overcome these issues. First, we presented diverse party systems, sweeping through multiple iterations of voting weights (Laver and Benoit 2015) and ideological arrangements, varying the position of plurality parties and pivotal members, etc. Second, we controlled the internal information environment by directly providing information about party characteristics, preventing contamination from subjects' differing levels of political information. Third, we controlled the external information environment by using fictional parties (in our main experiments), eliminating possible (unmeasured) external signals like elite messages and co-governance histories.[16]

Our instrument presented subjects with a hypothetical legislative composition (party positions and seat shares) and asked them to form a government by assigning each party to a role: prime minister, cabinet partner, or opposition.[17] The only constraints were identifying exactly one PM and assigning every party a single role. We do not elicit coalition *probabilities* directly; rather, subjects identify the single most likely cabinet. We test our hypotheses by both estimating the correlates of these expectations with statistical models and evaluating the predictive accuracy of the heuristic model's deterministic theoretical predictions

[15]That coincidence is made much more likely when one is constrained to using plausible size and ideological treatments because these treatments, by definition, must overlap with the real-world seat and ideological distributions that (in part) produced the pre-treatment beliefs about coalition likelihoods that the subject brings to these kinds of experiments.

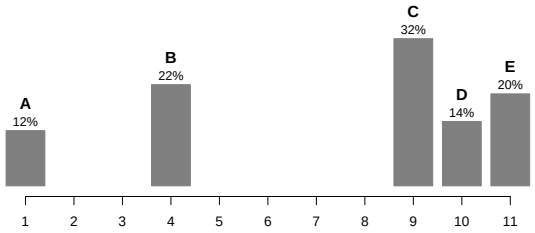
[16]In an unpublished study that is superficially similar to ours, Duch and Tyran (2012) present subjects with eight coalition vignettes, asking them to identify the most likely coalition. This study does not tend to these points as it does not allow unconstrained choice over parties and roles, does not present a diverse array of party systems, and specifically instructs respondents to choose a *majority* coalition.

[17]We collapsed a fourth category, "outside supporter," into "opposition" because extant research lacks clarity on how voters perceive external support parties (Tromborg, Stevenson and Fortunato 2019; Fortunato et al. 2021; Hjermitslev 2024). In the appendix we demonstrate that re-designating support parties as cabinet partners does not affect our central results, aside from redistributing minority cabinet probability to surplus majority and degrading model fit.

for each respondent (in each choice situation). Importantly, this experimental format asks the respondent to imagine the election result as given, so we do not account for the possibility that voters somehow account for noise in pre-electoral forecasts about post-election party sizes, though we note that evidence suggests that voters are able to accurately infer post-electoral party sizes from pre-election polling (Lee, Haime and Stevenson 2019).

There are two main benefits to this response structure. First, is allows respondents the freedom to identify any cabinet configuration with any party as PM. Second, by not providing the size and ideological characteristics of potential coalitions directly, but requiring respondents to work out the relevant relationships between parties for themselves, the experiment probes the precise reason many scholars doubt voters can from rational(ist) expectations about coalition formation—the cognitive complexity of translating a party system into comparisons between cabinet alternatives. The prompt (with a sample composition) reads:[18]

"Imagine that a legislative election produced a legislature composed of the following five parties."



"What role do you think each party above is most likely to play in the post-election government?"

We label parties with single letters that do not correspond to party abbreviations used in the respondents' countries. The graphic displays seat shares and ideological positions on a "left-right" x-axis, that is otherwise unlabeled, allowing respondents to interpret the dimension freely. While this constrains ideological dissimilarity to a single dimension, we recognize voters' political understanding may be multidimensional and so defend this design choice, in light of recent multidimensional models of cabinet formation by De Marchi and Laver (2020, 2023) and Martin and Vanberg (2025) in Appendix 3.

We constructed party sizes and positions by defining all possible unordered seat divisions for an n-party system with the following constraints:[19] seats must sum to 100%, allocations occur in 2% intervals, and no party has less than 2% or more than 52% of seats. We then defined all unordered ideological distributions,

[18] To ensure comprehension, we required subjects to correctly identify a party description (e.g., "party A is the leftmost party and has 12% of the seats") from four alternatives.

[19] By "unordered" we mean that sets $\{24, 24, 24, 26, 2\}$ and $\{26, 2, 24, 24, 24\}$ constitute a single distribution.

allocating positions between 0 and 10 in whole numbers with only one party per position. Because even a five-party system yields over 100 million unique configurations of these variables, we randomly sampled 1,000 seat distributions and ideological arrays, randomly matching them to create the electoral results presented to subjects.

It is critical to clarify what this experiment can and cannot demonstrate. Because subjects receive the seat shares and ideological positions of all parties, we can detect the application of cues derived from size and ideology (e.g., plurality or median status, ideological compatibility, minimal-winning status). However, we intentionally hold constant factors derived from history and elite messaging (e.g., incumbency, pre-electoral pacts, media forecasts). Consequently, we cannot estimate the weight voters assign to such information "in the wild," nor can we conclude voters rely *exclusively* on our model rather than alternatives. We can only evaluate the degree to which (or the share of respondents whose) expectations are 1) consistent with the predictions of our model and 2) covary across contexts as expected given the differing empirical regularities associated with cabinet formation in each country.[20]

We gathered data in two ways. First, we embedded the experiment in a large institutional survey of $\sim 2,000$ respondents each in France, Germany, Hungary, Italy, the Netherlands, Poland, Spain, Sweden, and the United Kingdom, where each respondent evaluated two episodes for a five-party system.[21] Second, we administered the experiment to $\sim 1,000$ subjects each in Denmark, Germany, and the Netherlands, where each subject evaluated 12-15 cabinet formation episodes featuring different legislative compositions of five, six, seven, eight, or nine parties.[22] These data allow us to compare cue usage as party system complexity increases and to analyze cue weights across a diverse set of countries.

Our analysis proceeds in three steps. First, we estimate conditional logit models of our subjects' expectations (identifying coalescing parties and the PM) and compare those estimates to a parallel model of real-world cabinet formation using identical covariates. This allows us to describe the correlates of subjects' expectations, evaluate the VaPS model by comparing the drivers of coalition expectations to the drivers of real-world formations, and to test H1-3. Second, we assess the core model's overall fit by calculating the proportion of respondents' expectations that are consistent with our model under varying definitions of compatibility and winning status. Importantly, these tests are not tests of statistical model fit—they are

[20] In a later section discussing external validity, we investigate whether size and ideology predict coalition expectations in experiments with real-world parties for which respondents may know coalition histories and may have received media and elite messages.

[21] This survey was administered by SFB 884 Political Economy of Reforms at the University of Mannheim.

[22] Respondents were recruited from Qualtrics Panels. Details of the survey administration and human subjects compliance are in the Appendix.

tests of how well our data fit the *deterministic* predictions of (various versions of) our theoretical model. Finally, we test H4 by correlating respondent cue weights with the real-world empirical associations between those cues and cabinet formation outcomes in each context. These analyses demonstrate that respondents' expectations comport well with our model (and political science models of actual cabinet formation), that use of our heuristic model is likely widespread, and that we can reject the null hypotheses for H1, H2, and H4 in every country, and H3 in all but one.

Analysis

Assessing a VaPS model of expectation formation

Political scientists studying coalition-directed voting often assume voters generate expectations using a mental model incorporating the same predictive factors used in empirical models of real government formation (e.g., Duch and Stevenson 2008; Duch, May and Armstrong 2010; Indridason 2011). Above, we called this the Voter as Political Scientist (VaPS) model. However, because the full VaPS model is cognitively demanding, scholars studying how voters actually generate coalition expectations have payed surprisingly little attention to exploring it empirically, instead assuming voters rely on elite or media signals (Bargsted and Kedar 2009; Irwin and Van Holsteyn 2012; Eberl and Plescia 2018; Bowler, McElroy and Müller 2022) or co-governance histories (Gschwend 2001, 2007; Armstrong and Duch 2010). Even Duch, May and Armstrong (2010), who in their empirical models actually impute coalition probabilities to voters using a VaPS model, argue against the idea that voters are actually using that model. Instead, they suggest that voters absorb this information from "the stability in coalition formations, public opinion polling... and party efforts that signal which coalitions are more or less likely to occur" (706).

We evaluate the VaPS model directly in four ways. First, following Martin and Stevenson (2001), we estimate conditional logit models on both our experimental expectations data and real-world cabinet formation data from the same 10 countries. This allows us to assess whether or not the predictors of our subjects' experimental coalition expectations are broadly similar to the real-world predictors of cabinet formation in those countries. Second, we compare the fully specified VaPS model to a limited specification that includes just the three covariates required for our heuristic model, demonstrating nearly identical predictive accuracy. Third, we estimate separate models using experimental data from hypothetical legislatures of five, six, seven, eight, and nine parties, finding essentially no change in expectation formation as bargaining complexity increases. Fourth, we estimate separate models for each country to assess contextual differences.

Our VaPS specification is based on Martin and Stevenson (2001) Model 7, the same specfication used by

Duch, May and Armstrong (2010). We adapt it in three ways. First, our experimental data only provides party sizes and positions, so we limit covariates to those derived from these features: inclusion of the largest party, inclusion of the median party, minority/minimal-winning status (relative to surplus majority), ideological divisions in the cabinet, ideological divisions in the opposition (interacted with minority status), and logged number of parties.

Second, we simultaneously model cabinet identification and PM selection. This requires restructuring the data so each potential PM in each potential cabinet has its own row—e.g., cabinet AB has two rows: 1) A as PM and 2) B as PM. It also requires distinguishing whether individual characteristics (plurality and median status) apply to the PM or a partner party.

Third, we employ a different measure of "ideological compatibility" than is often used in cabinet formation models. Lacking consensus in the literature on how voters conceptualize multiparty ideological compatibility, we use the measure providing the best empirical fit: the sum of all partners' distances to the PM—e.g., if the PM coalesces with one party that is 1 unit away and one party that is 3 units away, the measure takes on a value of 4. Our substantive conclusions are robust to alternative measures.

We estimate this specification with conditional logit regression on both the real-world cabinet formation data and our experimental expectations data, and include a limited specification containing only the core covariates used in our heuristic model, which we call "VaPS-lite": largest party as PM, total distance from PM, and minimal-winning size. Continuous covariates are standardized (mean=0; SD=1). Table 1 records the experimental and real-world results.[23]

The real-world and experimental models show important similarities. Immediately clear is that the *signs* on all coefficients are identical across the models. It is also clear that the inclusion of the largest party as PM is as critical to coalition expectations as it is to real-world cabinet formation. 80% of our respondent's expected coalitions include the largest party as PM, mirroring the rate of 81% in our sample of real-world cabinets. Ideological compactness is also highly salient in both models. Our subjects have evidently intuited the salience of two of the most robust predictors of cabinet formation.

However, notable differences exist. Median party status as PM is highly predictive of real cabinet compositions (third most salient covariate) but less relevant in respondent expectations (ninth most salient), consistent with our discussion of the complexity of the median cue. Additionally, while the negative estimate on minority status is large and significant in both models, the positive estimate on minimal-winning status

[23]We use ParlGov data (Döring and Manow 2021) from 1970-2020 to model real-world cabinet formation. The real-world model yields more correct predictions, driven in part by the much higher frequency of single-party majorities in the real-world data compared to the experimental sample.

Table 1: Conditional logit regressions of cabinet formation and prime ministerial selection using real-world and experimental data. Experimental data only include 5-party systems.

	Real-world cabinet formation	Subjects' coalition expectations Full VaPS	VaPS-lite
Largest party is PM	2.792	2.358	2.795
	(0.237)	(0.012)	(0.009)
Largest party is a junior partner	−0.372	−0.524	
	(0.392)	(0.018)	
Median party is PM	1.503	0.393	
	(0.273)	(0.014)	
Median party is a junior partner	1.194	0.568	
	(0.238)	(0.012)	
Minimal-winning coalition	0.214	0.660	1.048
	(0.233)	(0.010)	(0.007)
Minority	−1.389	−0.730	
	(0.303)	(0.015)	
Total distance from PM	−5.647	−1.629	−1.706
	(0.503)	(0.011)	(0.007)
Opposition span	−0.021	−0.264	
	(0.120)	(0.006)	
Opposition spread × minority coalition	−0.161	−0.427	
	(0.161)	(0.010)	
Number of parties in coalition (logged)	−0.470	−0.438	
	(0.084)	(0.006)	
Episodes	209	80,873	80,873
N	6,617,804	35,918,432	35,918,432
Log Likelihood	−763	−296,326	−302,188
$E[\hat{y}\|y]$	0.354	0.205	0.211

The VaPS-lite estimates produce odds ratios of 16.36, 2.85, and 0.18, for the three variables in that model (in the listed order). This corresponds to a change in the odds of formation of 1,536% , 185%, and -85%, respectively, for a one unit change in each variable (recall that the continuous variables are standardized). Model construction based on Martin and Stevenson (2001), Duch, May and Armstrong (2010), and Kayser, Orlowski and Rehmert (2023).

is large and significant for our subjects, but smaller and statistically insignificant in the real-world data.

The third column provides results for the VaPs-lite model containing only the three focal heuristic cues. Overall fit as estimated by the log likelihood is *virtually identical* to the full VaPS model. Further, the predictions from the VaPS-lite model of expectations are significantly more accurate than the full VaPS model of expectations.[24]: $E[\hat{y}|y]_{\text{VaPS-lite}} = 0.211 > E[\hat{y}|y]_{\text{VaPS}} = 0.205; t = 3.242$. This is consistent with an ecologically rational model, built upon just a few cues. The other factors that we exclude from the limited model *are* predictive, but the data suggest that they are primarily influencing expectations at the margins,

[24]Here we predict the expected alternative to be the one with the highest probability. Other metrics produce the same conclusion).

perhaps when the core cues fail to discriminate to a single alternative. In sum, Table 1 suggests subjects possess a mental model that is broadly consistent with political scientists' empirical models of real-world cabinet formation. However, the these results also suggest that voters' mental models are likely simpler, relying heavily on a small set of accessible, highly predictive cues.

Before disaggregating the experimental data to compare across party-system size and countries, we can use the estimates in Table 1 to test H1-3. Recall that H1 predicts subjects will expect the plurality party to provide the PM (positive estimate); H2 predicts subjects will expect coalition parties to be ideologically compatible (negative estimate on total distance from PM); and H3 predicts subjects will expect the coalition to be winning (positive estimate). Both the VaPS and VaPS-lite model clearly reject the null of no association for H1-3.

Table 2: Comparing cue-weights across party system size in Denmark, Germany, and the Netherlands.

	Party system size				
	Five	Six	Seven	Eight	Nine
Largest party is PM	2.106	2.180	2.317	2.402	2.541
	(0.040)	(0.039)	(0.038)	(0.037)	(0.036)
Largest party is a junior partner	−0.730	−0.838	−0.723	−0.825	−0.808
	(0.060)	(0.060)	(0.058)	(0.059)	(0.059)
Median party is PM	0.642	0.681	0.676	0.796	0.704
	(0.045)	(0.045)	(0.045)	(0.046)	(0.046)
Median party is a junior partner	0.723	0.714	0.743	0.861	0.801
	(0.040)	(0.039)	(0.037)	(0.037)	(0.037)
Minimal-winning coalition	0.515	0.577	0.648	0.726	0.825
	(0.032)	(0.031)	(0.030)	(0.030)	(0.030)
Minority	−0.787	−1.006	−1.150	−1.161	−1.282
	(0.062)	(0.058)	(0.055)	(0.052)	(0.051)
Total distance from PM	−1.240	−1.353	−1.537	−1.664	−1.739
	(0.029)	(0.029)	(0.030)	(0.030)	(0.031)
Opposition span	−0.186	−0.315	−0.414	−0.489	−0.556
	(0.019)	(0.019)	(0.020)	(0.020)	(0.021)
Opposition spread × minority coalition	−0.232	−0.203	−0.221	−0.381	−0.303
	(0.040)	(0.042)	(0.044)	(0.044)	(0.045)
Number of parties in coalition (logged)	−0.300	−0.465	−0.559	−0.643	−0.704
	(0.023)	(0.020)	(0.018)	(0.016)	(0.015)
Alternatives per episode	80	192	448	1,024	2,304
Episodes	8,084	8,052	8,063	8,087	8,069
N	646,720	1,545,984	3,612,224	8,281,088	18,590,976
Log Likelihood	−24,025	−28,324	−32,338	−36,148	−39,865
$E[\hat{y}\|y]$	0.279	0.229	0.202	0.170	0.160

We now ask whether the correlates of our subjects' coalition expectations are consistent as we change the size of the party system in our experiments from five to nine parties, which increases total potential alternatives from 80 to 2,304. Pooling data from our directed surveys of Denmark, Germany and the Netherlands, we estimate the VaPS model for each party system size, and record the results in Table 2. The correlations and their relative magnitudes are consistent across party system size. Estimates for the core cues—largest party as PM, total distance from PM, and minimal winning—show remarkable stability. Indeed, they increase in size and efficiency as the party system grows, while estimates on most ancillary cues do not. This confirms that our results for H1-H3 in the Table 1 were not artifacts of the smaller party system and suggests external validity to large party systems.

Table 3: Conditional logit regressions of cabinet formation and prime ministerial selection for individual countries.

	Country of sample									
	Denmark	France	Germany	Hungary	Italy	Netherlands	Poland	Spain	Sweden	UK
Largest party is PM	1.963	2.202	2.136	2.553	2.370	2.397	2.356	2.425	2.088	2.881
	(0.068)	(0.045)	(0.046)	(0.054)	(0.050)	(0.043)	(0.048)	(0.048)	(0.049)	(0.060)
Largest party is a junior partner	−0.472	0.134	−0.838	−0.086	−0.120	−0.493	−0.253	−0.502	−0.437	−0.147
	(0.100)	(0.064)	(0.068)	(0.075)	(0.070)	(0.061)	(0.071)	(0.080)	(0.067)	(0.089)
Median party is PM	0.710	0.032	0.510	−0.135	−0.123	0.408	0.024	0.257	0.393	0.015
	(0.077)	(0.051)	(0.051)	(0.059)	(0.056)	(0.046)	(0.054)	(0.055)	(0.055)	(0.066)
Median party is a junior partner	0.601	0.252	0.697	0.208	0.241	0.590	0.351	0.380	0.535	0.417
	(0.075)	(0.047)	(0.042)	(0.047)	(0.046)	(0.037)	(0.046)	(0.048)	(0.049)	(0.053)
Minimal-winning coalition	0.293	0.711	0.822	0.788	0.798	0.460	0.773	0.817	0.518	0.949
	(0.059)	(0.050)	(0.035)	(0.047)	(0.047)	(0.033)	(0.048)	(0.049)	(0.041)	(0.048)
Minority	−0.915	0.732	−0.893	0.169	0.423	−0.520	0.630	0.639	−0.510	0.383
	(0.111)	(0.066)	(0.073)	(0.072)	(0.069)	(0.060)	(0.067)	(0.069)	(0.073)	(0.077)
Total distance from PM	−1.508	−0.826	−0.940	−1.125	−1.026	−0.766	−0.822	−1.160	−1.099	−0.864
	(0.058)	(0.036)	(0.031)	(0.037)	(0.035)	(0.026)	(0.034)	(0.037)	(0.034)	(0.040)
Opposition span	−0.212	0.083	−0.069	−0.060	−0.088	−0.147	−0.076	−0.125	−0.300	0.016
	(0.035)	(0.027)	(0.020)	(0.025)	(0.025)	(0.018)	(0.025)	(0.026)	(0.023)	(0.027)
Opposition spread × minority coalition	−0.056	−0.475	−0.417	−0.707	−0.570	−0.451	−0.500	−0.442	−0.301	−0.471
	(0.068)	(0.041)	(0.048)	(0.047)	(0.044)	(0.041)	(0.043)	(0.043)	(0.048)	(0.048)
Number of parties in coalition (logged)	−0.310	−0.160	−0.447	−0.080	−0.053	−0.301	−0.094	0.011	−0.269	−0.349
	(0.041)	(0.027)	(0.026)	(0.030)	(0.029)	(0.024)	(0.028)	(0.028)	(0.029)	(0.031)
Alternatives per episode	80	80	80	80	80	80	80	80	80	80
Episodes	2,375	4,432	7,013	4,346	4,432	7,542	4,480	4,422	5,206	4,354
N	190,000	354,560	561,040	347,680	354,560	603,360	358,400	353,760	416,480	348,320
Log Likelihood	−7,160	−15,433	−19,763	−13,989	−14,971	−24,482	−15,392	−14,383	−16,694	−12,470
$E[\hat{y}\|y]$	0.265	0.139	0.322	0.214	0.178	0.210	0.168	0.191	0.236	0.237

Table 3 compares subject expectations across contexts by estimating the VaPS specification individually for each country, limited to 5-party systems for comparability. In all countries, the largest positive estimate is for the largest party providing the PM, and the largest negative estimate is for total distance from the PM, allowing conclusive rejection of the nulls for H1 and H2 across all contexts.

Estimates for minimal-winning status are consistently large and positive, but less stable cross-nationally.

In Poland and Spain, the minimal-winning estimate is only slightly (but significantly) larger than the minority size estimate ($p = 0.009$ and $p = 0.002$, respectively). In France, the estimates for minimal-winning and minority status are statistically indistinguishable ($p = 0.698$). Thus, while we cannot reject the null for H3 in France using the minimal-winning criterion, we can reject it everywhere else. Furthermore, majority cabinets (combining minimal-winning and surplus) are more likely to be expected than minority cabinets in every country.

We also observe clear contextual differences in ancillary cue estimates (e.g., median is PM) which we will explore more in testing H4. For now, we emphasize that estimates for our focal heuristic cues are remarkably consistent across contexts, supporting H1 and H2 globally, and H3 in all countries but France.

Heuristic model fit

The previous section demonstrated that coalition expectations are broadly consistent with the VaPS perspective, though three features (largest party as PM, ideological compatibility, and minimal-winning status) provide most of the predictive power in that statistical model. Here, we assess various versions of our heuristic model in a different way. Instead of using a statistical model, we directly calculate the deterministic predictions of the theory for each respondent-formation episode and examine how often these predictions are correct (among other metrics).[25]

Identifying the heuristic prediction requires defining "ideological compatibility" and "winning." Given a lack of consensus in the literature, we test eight definitions of compatibility and four of winning. The compatibility measures are:

1. *Connected*: parties in coalition are contiguous
2. *Coalition span*: total range of the coalition
3. *Average mean distance*: average of coalition parties' distances to unweighted coalition mean
4. *Average weighted mean distance*: average of coalition parties' distances to seat-weighted coalition mean
5. *Average PM distance*: average of partners' distances to PM
6. *Average weighted PM distance*: average of partners' seat-weighted distances to PM
7. *Max PM distance*: the maximum partner distance to PM

[25]This approach to theory assessment, common in the other sciences, is often discussed under the label "logical model" evaluation in political science (e.g., Atsusaka 2021).

8. *Total PM distance*: the sum of partner distances to PM

While correlated, these definitions imply somewhat different formation processes. For example, minimizing distance to the prime minister implies the PM is central to government formation, whereas minimizing total dissimilarity implies that no one party is more important than others. We see definitions 1-3 as holistic or egalitarian, 4 as ambiguous, and 5-8 as clearly PM-driven.

These definitions differ in two other important ways. First, some measures are more complex, or, difficult to calculate or intuit. For example, *average weighted PM distance* is substantially more difficult to calculate than *total PM distance*. Second, some measures discriminate, or identify a unique prediction, more often than others. For example, there will be fewer "ties" when evaluating *average weighted PM distance* than *total PM distance*, as the seat weights induce variation discriminating between close sets. As such, we will want to consider how many predictions our heuristic model makes under different ideological compatibility measures when we compare them below.

Our definitions of winning, from least to most restrictive, are:

1. Majority inclusive of ties
2. Majority
3. Minimal-winning inclusive of ties
4. Minimal-winning

"Majority" refers to any majority, minimal-winning or surplus. By "ties" we mean coalitions that have precisely 50% of seats rather than a proper majority. As with compatibility measures, there is little guidance in the literature as to how voters think about winning or majority status, whether or not they understand that 50% is not a majority, etc., so we investigate each of these possibilities.[26]

We produce predictions under each combination of these definitions and those for compatibility and examine how well each predicts our respondents expectations. Specifically, for each episode, we identify all coalitions that could form, differentiating the PM (80 potential cabinets in 5-party systems, up to 2,304 in 9-party systems, with a sample average of 444). We then identify subsets in which the plurality party is PM and the winning criterion is met. Within that subset, the heuristic prediction ($\hat{y}$) is the cabinet that (meets or)

[26]While some voters may conflate large pluralities with majorities (Duch, Przepiorka and Stevenson 2015), establishing a "close enough" threshold is arbitrary and will systematically make our model look more accurate (because it will make some expected minority cabinets look like majorities) so we avoid it. Oversized "minimum safe" definitions (Provins, Monroe and Fortunato 2022) similarly require arbitrary thresholds so we avoid those as well.

minimizes the compatibility criterion. All definition combinations yield at least one prediction per episode, except for the "connected" compatibility criterion when coupled with minimal-winning or minimal-winning inclusive of ties, which is impossible to satisfy in 12-18% of episodes.

Table 4: Describing heuristic model fit across definitions of ideological compatibility and winning. In each column, the best performing measure is **bolded** and worst is *italicized*. There is an average of 444 potential coalitions per episode.

	Majority + ties			Majority		
Definition	Average # of predictions	Correct classification	Predictive efficiency	Average predictions	Correct classification	Predictive efficiency
Connected	*8.508*	**0.361**	*0.042*	*8.216*	**0.333**	*0.041*
Coalition span	1.578	0.246	0.156	1.658	0.251	0.151
Average mean distance	1.236	0.206	0.166	1.254	0.204	0.163
Average weighted mean distance	1.135	0.202	0.178	1.140	0.201	0.176
Average PM distance	1.299	0.205	0.158	1.300	0.200	0.154
Average weighted PM distance	**1.063**	*0.173*	0.162	**1.067**	*0.175*	0.164
Max PM distance	2.022	0.254	0.125	2.151	0.255	0.119
Total PM distance	1.201	0.231	**0.192**	1.208	0.235	**0.194**
Ensemble Model	2.110	0.405	0.192	1.885	0.383	0.203

	Minimal-winning + ties			Minimal-winning		
Definition	Average # of predictions	Correct classification	Predictive efficiency	Average predictions	Correct classification	Predictive efficiency
Connected	*1.758*	**0.248**	*0.141*	*1.466*	0.220	*0.150*
Coalition span	1.326	0.236	0.178	1.298	**0.237**	0.182
Average mean distance	1.213	0.222	0.183	1.218	0.225	0.184
Average weighted mean distance	1.133	0.218	**0.193**	1.136	0.221	**0.194**
Average PM distance	1.183	0.216	0.183	1.171	0.217	0.185
Average weighted PM distance	**1.059**	*0.175*	0.165	**1.060**	*0.179*	0.168
Max PM distance	1.340	0.230	0.172	1.263	0.227	0.180
Total PM distance	1.201	0.231	0.192	1.208	0.235	**0.194**
Ensemble Model	1.142	0.290	0.254	1.097	0.268	0.244

Given an average of 444 potential cabinets per episode, the baseline probability of randomly guessing a respondent's expectation is $E[y] = 0.002$. Table 4 details model fit under all compatibility and winning definition combinations using three statistics that average over all episodes (regardless of number of parties included): the average number of heuristic model predictions per episode, the heuristic model's rate of correct classification, $E[\hat{y}|y]$, which is the number of times the respondent's expected coalition was also a heuristic model prediction, divided by the total number of respondent expectations, and predictive efficiency, $E[y|\hat{y}]$,

which is the number of times a heuristic model prediction was also the respondent's expected coalition, divided by the total number of heuristic model predictions. High rates of classification are desirable, but so too is discrimination or specificity. The efficiency statistic captures the desire for specificity by penalizing multiple predictions per episode. In each column, the most desirable performance measure is **bolded** and least desirable is *italicized*. Recall that these are logical classifications, rather than statistical model estimates, so there are no confidence intervals to report.[27]

We also calculate the performance of the heuristic model by aggregating across all of these definitions of ideological compatibility separately for each winning criteria (labeled the "Ensemble Model" in the last row in each column). In other words, the Ensemble Model counts a prediction as correct if the respondent expected a cabinet to form that included the largest party, was winning by the indicated criterion, and was most ideologically compatible by at least one of these definitions. Since our model anticipates heterogeneity in how voters think about ideological compatibility and wining status, this ensemble approach provides the best measure of the extent to which voters actually expect ideologically compatible, winning coalitions to form.

The ensemble model reveal that up to 41% of all expectations are consistent with the heuristic model—an excellent fit.[28] Indeed, given the large literature in political science suggesting that roughly half of the citizens in modern democracies are uninterested and unknowledgeable about politics (Kinder and Kalmoe

[27] Another fit statistic often used to test the predictions of theoretical models is Selten's (1991) measure of predictive efficiency. Unlike the measure of predictive efficiency commonly used in the coalition formation literature (and reported in Table 4), Selten's measure recognizes that a theory that makes, for example, 2 predictions, one of which is correct, out of 500 alternatives, should be judged as more accurate than one that make 2 predictions, one of which is correct, out of 5 alternatives. Formally, the measure is the rate of correct classification (r) minus the proportion of available alternatives the model predicts (a). Given the small number of predictions (relative to the large set of alternatives) that our theory makes under each definition in Table 4, the value of a is quite small and so Selten's measure is nearly identical to the classification rate, so we do not report it separately.

[28] This is especially true given the size of the alternative space over which we are making predictions. To see this consider the following comparison: in the previous footnote, we explain that Selten's (1991) measure of model fit in our cases is essentially the same as the correct classification rate, for example 0.405 for the ensemble prediction for majority + ties. This gives us a way to compare this level of accuracy to the predictive accuracy of well understood behavioral models in political science that make predictions over much smaller choice sets (e.g., vote choice, where choice sets in the western democracies are usually between 2 and 12). For example, applied to a partisan model of voting in the United States that predicts 80% of votes correctly (e.g., Bankert 2021), Selten's fit measure is only $0.8 - 0.5 = 0.3$. Similarly, a vote choice model in a multi-party democracy that predicts 50% of choices in an 8 party system (which is fairly typical for such models) seems to predict less well than the U.S. model, but when the size of the choice set is accounted for, it compares much more favorably, with a Selten score of 0.375. Given this, our predictive accuracy (and Selten score) of between $\sim 0.21 - 0.41$ for the various ensemble models in Table 4 is impressive.

2017) this rate of predictive success is, in our view, remarkable.

Comparing definitions, the *connected* compatibility definition generates the most correct classifications but discriminates poorly (making the most predictions). Conversely, *average weighted PM distance* discriminates best but classifies poorly. *Total PM distance* is the most predictively efficient compatibility definition under three of four winning criteria and achieves the highest overall efficiency score (0.194). This definition has two interesting and appealing properties. The first is that, because it is additive in parties, all winning coalitions that minimize *total PM distance* are minimal-winning—hence identical fit statistics across majority and minimal-winning definitions (both inclusive and exclusive of ties). The second is that this measure captures an intuitive process for applying our heuristic model: identify the plurality party as PM and, if it is not a single-party majority, add its nearest neighbor. If this produces a majority then stop. If not, then add the next nearest neighbor, and so on.

Using *total PM distance* and minimal-winning, the model yields a correct classification rate of 0.235 and predictive efficiency of 0.194. While this rate of correct prediction is impressive, it may understates the usage of the heuristic model because it ignores "near misses," expectations closely mirroring the model prediction but failing strict classification. For instance, in a legislature $\{A_1^{40}, B_3^4, C_5^{12}, D_7^8, E_9^{36}\}$, a respondent might identify A as PM, add nearest neighbor B (44%, not majority), then add C to secure a majority (56%). This intuitive cabinet is in keeping with the spirit of our model but would not be counted as a correct classification because B is superfluous to the majority and increases the sum of distances.

In Appendix 10, we describe an approach that accounts for such near-misses in our predictions. In short, for each alternative, we use the seat shares and ideologies of the included parties to calculate the implied left-right position of all possible alternatives (using the seat-weighted mean position, following Martin and Vanberg 2011 and Fortunato 2021). We then create a measure of deviance from the heuristic model prediction by subtracting the model-predicted cabinet position from the position of all potential cabinets. We classify each potential coalition within one-unit of the heuristic model's prediction a "near miss" and calculate how many of our respondents' expected coalitions are near misses. This is recorded in Table 5, which explores near misses in episodes with a unique model prediction, so we need not calculate near misses across multiple predictions. It records potential alternatives, baseline accuracy rate ($\frac{1}{\#\text{ of alternatives}}$), the strict classification rate, and classification inclusive of near misses across party system sizes. Note that data for 6+ parties come solely from Denmark, Germany, and the Netherlands (DGN).

Strikingly, the heuristic model's predictive performance remains stable even as the choice set expands massively. In the most complex 9-party systems, the model classifies expectations at a rate of 0.58 inclusive of near misses and has a strict classification rate of 0.19 which is roughly 437 times better than baseline

Table 5: Heuristic model strict classification and near misses across party system sizes. Compatibility definition is *total PM distance* and winning criterion is minimal-winning. Only episodes in which heuristic model makes a unique prediction are included (83% of episodes), therefore correct classification and predictive efficiency are equivalent. Episodes with 6 or more parties are only available from Denmark, Germany, and Netherlands, identified as "DGN" below.

Parties	Alternatives	Baseline hit rate	Strict classification	Near misses	Sample
Five	80	0.0125	0.2340	0.5723	all countries
Six	192	0.0052	0.2578	0.6043	DGN
Seven	448	0.0022	0.2322	0.6016	DGN
Eight	1024	0.0010	0.2048	0.6038	DGN
Nine	2304	0.0004	0.1898	0.5816	DGN
Total (average)	417	0.0086	0.2295	0.5822	

chance. Across the full sample, the strict classification rate is 0.23 (over 26 times baseline chance). This is robust evidence for the widespread application of an abstract mental model consistent with our theory.

Finally, using only five-party episodes to ease comparison, Table 6 details cross-national variation in heuristic model fit, again, using *total PM distance* and minimal-winning.[29] Fit varies substantially across contexts. Indeed, the model correctly classifies subject expectations in Germany at 180% the rate at which it correctly classifies French expectations. Much of this variation comports with the heuristic model's alignment with local empirical regularities. For example, at the time of administration, every post-election cabinet formed in Germany in the post-war era had been a minimal-winning coalition that minimized partner distance from the plurality PM—precisely the contextual conditions that should induce good local fit to our expectations model. Similarly, the empirical regularities that emerge from government formation in France, Poland, Italy, and elsewhere often diverge from these patterns. Our ecologically rational heuristic model predicts that such deviations will ultimately be reflected in differences in expectations across contexts. We examine the consequences of this kind of contextual variation for coalition expectations in more detail in the next section.

[29] A fair comparison of correct classification rates across countries must account for variation in the distribution of party systems (treatments) used in each country. For 6 countries in the multi-country survey, these distributions are quite similar, using equal proportions from the Laver and Benoit (2015) classification of party systems, but for the three cases that were part of the pilot and our directed surveys, these distributions are different. To adjust for these different distributions of party systems, we regress the classification of each episode $\{0, 1\}$ on country while absorbing party system induced variation by estimating party-system fixed effects. Each classification estimate in Table 6 has a standard error of ~ 0.01.

Table 6: Heuristic model fit across country samples. Compatibility definition is *total PM distance* and winning criterion is minimal-winning. Only five-party episodes.

Sample	Strict classification	Near misses
Germany	0.328	0.581
Denmark	0.260	0.602
United Kingdom	0.251	0.593
Hungary	0.240	0.562
Netherlands	0.234	0.545
Sweden	0.221	0.560
Italy	0.211	0.566
Spain	0.210	0.570
Poland	0.196	0.555
France	0.178	0.536
France	0.178	0.536

Assessing contextual correspondence

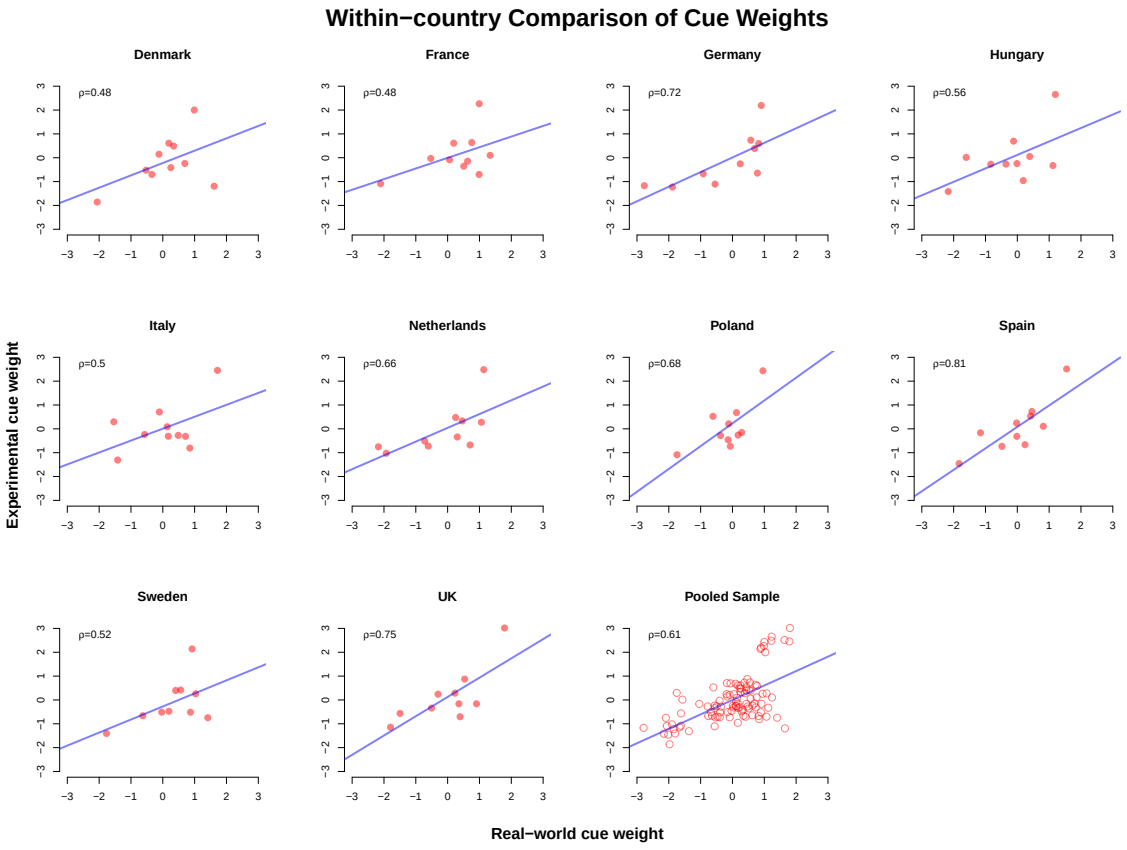
We test H4 by assessing whether the cabinet characteristics predicting real-world outcomes in a specific context also predict our subjects' experimental expectations in that context. Following the theory of ecologically rational heuristics (Fortunato and Stevenson 2019; Fortunato et al. 2021), we expect a strong, positive within-country correspondence between experimental and real-world cue weights. Recovering such correlations would support H4, demonstrating that voters' mental models of coalition formation are ecologically rational, or, well-adapted to their specific context.

To estimate country-specific cue weights, we modify the VaPS models in Table 1. First, for experimental expectations, we limit the data to five-party episodes and interact each covariate with country to derive the country-specific cue weights that will be become our dependent variable in the coming analysis. For the real-world data, following Fortunato, Stevenson and Vonnahme (2016), we estimate country-level random coefficients for each cue to derive the country-specific cue weights that will be our central covariate in the coming analysis. We use random effects because of separation, or perfect prediction issues (e.g., all post-war German cabinets are minimal-winning with a plurality PM). Random coefficients estimators produce efficient, comparable estimates under separation without imposing strict ex-ante constraints (Crisman-Cox 2021). With these country-level estimates, we can assess the correspondence between the real-world empirical regularities of cabinet formation and the correlates of our subjects' experimental coalition expectations in each country.

Figure 1 plots the correspondence of standardized real-world and experimental cue weights for each country. Each point represents a cue (e.g., plurality PM inclusion), with the x-axis indicating its real-world

correlation to cabinet formation and the y-axis indicating its correlation to cabinet formation in our subjects' expectations. In every country there is a clear, strong, positive relationship, with a pooled correlation of $\rho = 0.61; p < 0.01$. The data are quite conclusive that the weights voters place on observable characteristics in generating their coalition expectations are strongly and positively correlated with the actual predictiveness of those characteristics to cabinet formation in their country. This is strong evidence of ecologically rational expectations.

Figure 1: Comparison of relative weights of cues within countries.



While overall within-country correspondence is strong, the plots reveal it is disproportionately driven by a few cues. The pooled plot shows distinct clusters of highly positive and highly negative cues, alongside a larger group that is clustered near the experimental (y-axis) zero-point but scattered across the real-world axis. We probe this using full- and restricted-sample regressions.

Table 7 regresses the experimental cue weights on their real-world counterparts, presenting a pooled model, models with fixed effects (country or cue), and restricted regressions that include only the three cues that are focal to our heuristic model, or omit those three. The within-country results, which essentially

replicate Figure 1, conclusively reject the null for H4.

Table 7: Relationship between real-world and experimental cue-weights in coalition expectations. All countries, 5-party systems only.

Covariate	Pooled	Within country	Within cue	Within country, only 3 focal cues	Within country, omit 3 focal cues
Real-world weight	0.242	0.246	0.019	0.345	0.070
	(0.026)	(0.027)	(0.028)	(0.022)	(0.033)
Country FEs		✓		✓	✓
Cue FEs			✓		
Observations	100	100	100	30	70
R^2	0.461	0.480	0.906	0.931	0.112

Focal cues are: largest party is PM, ideological compatibility, and minimal-winning status.

Comparing the results of the focal cues and non-focal cues shows that, even though both sets are positively and significantly associated, the strength of association is much greater for the focal cues in magnitude and efficiency. Again, this is in keeping with an ecologically rational model, where just a few cues can be leveraged into efficient inferences. The association of the non-focal cues to their real-world counterparts *is* significant, but strength of association is weaker. This suggests that subjects applying the heuristic model seem to do so in a manner calibrated to their local context, and, that when our subjects' expectations deviate from the strict predictions of the heuristic model, they tend to do so in a way that is consistent with the empirical regularities of real-world coalition formation in their country. However, the explanatory power of the deviations is quite modest compared to the explanatory power of the focal cues in the core heuristic model.

The within-cue model reveals an insignificant correlation, but this is driven entirely by two outlying deviations: the minority cue estimates in Denmark and Sweden. Danish and Swedish subjects largely expect majority cabinets (see Table 3), even though their countries exclusively (Denmark) or routinely (Sweden) form minority governments. Both place the second-largest negative weight on the minority cue, despite its strong positive real-world weight. Omitting these two observations reveals a significant within-cue correspondence of $\beta = 0.120$ (0.028). We speculate that Danes and Swedes intuitively understand that policymaking requires majorities and intentionally divorce their localized acceptance of minority governance from the experimental task.

Finally, we note that, on average, there is substantially less within-cue variability in expectations

($\bar{sd} = 0.251$) than there is in real-world coalition formation ($\bar{sd} = 0.827$). This means that our subjects' mental models of cabinet formation are more similar across countries than are the patterns of actual cabinet formation. A small part of this may be due to the compression of potential alternatives in these five-party systems (more on this below), but this is also in keeping with an efficient heuristic model in which the majority of the predictive power in expectations is derived from just a few focal cues that are widely applicable across countries.

Table 8: Comparing relationship between real-world and experimental cue-weights in coalition expectations across party system size in Denmark, Germany, and the Netherlands.

	Party system size				
Covariate	Five	Six	Seven	Eight	Nine
Real-world weight	0.261	0.286	0.314	0.331	0.342
	(0.044)	(0.046)	(0.050)	(0.053)	(0.056)
Country FEs	✓	✓	✓	✓	✓
Observations	30	30	30	30	30
R^2	0.577	0.595	0.605	0.604	0.592

Before concluding, we use the data from our directed surveys of Denmark, Germany, and the Netherlands to assess how the correspondence between real-world and experimental cue-weights may be influenced by party-system size. To generate the country-specific cue-weights, we estimate fully interacted conditional logit models of expectation formation for each party system size. Table 8 regresses these estimates on the real-world cue-weights for Denmark, Germany, and the Netherlands to recover the within-country association for legislative compositions of five, six, seven, eight, and nine parties. The comparison is revealing as it shows closer correspondence as the party system grows. Why? One reason is that the expansion of the party system increases the number and variability of potential alternatives which, in turn, allows our subjects more latitude to build coalitions in keeping with the empirical regularities of real-world cabinets in their country. This implies that the more constrained choice set utilized in our most of experimental instruments perhaps subdued our ability to recover evidence for our empirical predictions in general and H4 in particular, due to the compression of alternatives.

External validity

The preceding analyses demonstrate that many voters possess a mental model capable of generating reasonable coalition expectations. By isolating the influence of size and position, the experiments yield robust,

internally valid causal estimates. However, there are two potential hurdles to drawing strong conclusions about how voters generate coalition expectations in the real world from this positive result. First, most subjects evaluated five-party systems, which are arguably less complex than real-world environments with larger party-systems. We believe the stability of cue weights (Table 2) and model fit (Table 5) across increasing system sizes effectively addresses this concern.

Second, invoking a mental model in the real world might be unnecessary or more effortful than relying on elite signals or governance histories, so that when these cues are available (in the real-world, but not in our experiments with fictional parties) voters may turn to them instead. In Appendix 9, we report results from two follow-up experiments in Denmark and Germany that address this possibility. Specifically, they use the same experimental design as above but substitute real parliamentary parties for fictional ones. This allows respondents to bring to the experiment their pre-existing knowledge, including coalition histories and media narratives about likely coalitions, and use that knowledge, instead of size and ideology, to form expectations within the context of our experiment.

If our main experiments are not generalizable to the real world—only demonstrating that voters *can* use size and ideology to form sensible expectations, not that they actually do, because in the real world they use other cues we hold constant—we would expect to see much smaller effects of size and ideology in these new experiments. This is not the case. We find the same strong effects of size and ideology for experiments with real parties as with fictionalized ones (see Tables 9-11 in the Appendix).

Another benefit of using real parties is that it allows us to measure additional characteristics of potential coalitions and estimate their effects on expectations within our experiment. While a full discussion is in the Appendix, the most important of these variables captured historical patterns of coalition formation (and perceptions of those patterns) and, for Germany, a direct measure of the media message about the likelihood of each potential cabinet at the time of the survey. Thus, we can compare the estimated effects of these variables to the effects of size and ideology in the same model.

The results reveal that our subjects do incorporate coalition histories and the media message as expected. They are more likely to expect cabinets to form if they perceive longer histories of co-governance and positive media narratives. Crucially, however, subjects remain highly responsive to size and position characteristics. In fact, the relative weights placed on plurality PM inclusion, ideological compatibility, and winning status are substantially larger than those placed on history, media narratives, and the other added covariates discussed in the Appendix. This provides strong evidence that voters continue to apply our proposed mental model even in contexts more closely mirroring the real world.

Finally, it is worth emphasizing that our main (fictionalized) experiments *could have yielded null results.*

Had subjects failed to use a sensible mental model in our controlled setting, we could reasonably conclude that voters are unlikely to do so in the "in the wild." It is only relevant to ask if our subjects employ our proposed heuristic model in real-world settings because we now know that up to 41% of them use it to generate coalition expectations in our experiments. We think that the experiments using real-world parties, along with the existing correlational evidence from observational studies in real-world settings, which universally produce strong correlations between size and ideology and expectations, should increase our confidence that voters really do use such a mental model.

Discussion

In this paper we argue that many voters can (and do) produce sensible expectations for cabinet composition using a simple mental model of coalition formation in which the largest party provides the PM, the cabinet is of a "winning" size, and partners are "ideologically compatible." Our ensemble theoretical models—which allow for some heterogeneity in voters' use of these concepts—deterministically predict the exact cabinet that survey respondents think is most likely in up to 41% of cases. This is strikingly close to the rate at which the same set of variables can predict real-world cabinets, as we would expect if respondents' heuristic inferences (coalition expectations) combine cues using weights that reflect the real-world associations between cues and the inferential target (coalition compositions) as theories of ecologically rational heuristics suggest and our subsequent analyses show. Indeed, we also find evidence that coalition expectations vary with the empirical regularities of cabinet formation in particular contexts. These local adjustments, however, are small when compared to more universal drivers of expectations—party sizes and ideologies—that are important everywhere.

This is an impressive rate of prediction and local adaptation relative to other models of mass political attitudes and behavior in multi-party systems, especially those that predict choices over large sets of alternatives. Indeed, if we think that voters produce coalition expectations to think about and understand multi-party politics, to engage in political discourse, and to ultimately act on those ruminations by casting a coalition-directed vote, our results imply that most politically engaged or otherwise interested voters will have little trouble generating sensible coalition expectations, even for novel partisan combinations.

These experiments also move the small empirical literature on coalition expectations forward. While previous correlational work has included some size and ideological variables, showing that they correlate with expectations as much or more than alternatives like media messages and patterns of historical cabinet composition, no examination of a broad set of size and ideological variables, and no experimental work on

these influences, has ever been executed. Experimental research, however, is essential in this area because coalition expectations are embedded in a highly endogenous system in which parties communicate strategically about which coalitions are likely, previous coalition outcomes impact party sizes and ideologies in future periods, party sizes and ideologies impact which cabinets form, what the media says about them, and what voters expect, and, all of this runs through voters' noisy perceptions.

The consequences of this causal complexity, and the utility of experiments to sort it out, is perfectly illustrated by the two existing experimental studies of coalition expectations by Falcó-Gimeno and Muñoz (2017) and Bahnsen, Gschwend and Stoetzer (2020). Despite experimentally manipulating very direct elite/media messages, the causal impact of these manipulations on expectations was surprisingly mixed and, on balance, much weaker than indicated by previous correlational evidence. Indeed, taken together with our experiments, which verify and reinforce the correlational evidence that size and ideology matter to voters' coalition expectations, these results argue for a holistic theoretical model of coalition expectations in which coalition characteristics are causally prior to media messages and so impact coalition expectations both directly and through the media.

This is progress. Our experiments, understood in the context of both the correlational evidence and the experimental evidence on elite signals, paint a picture of a literature that is slowly zeroing-in on a causal structure driving coalition expectations, for some reasonable share of the population, that gives more weight to coalition characteristics than elite signals (with the impact of history not yet clarified).

That said, there is clearly more to be done. Most importantly, any experimental research program that asks subjects to engage in hypotheticals should be concerned with the generalizability of the results to the real-world. Following the other experiments in this area, we have tried to address that concern by asking whether voters continue to use size and ideological cues to form expectations when we replace the fictional parties in our experiment with real ones (see Appendix 9). While this approach pays for increased external validity with decreased internal validity, the fact that we get essentially the same result in both settings provides some assurance that voters are using simple mental models to form expectations in the real world and not just in our experiments. That said, it would be useful to extend these kinds of experiments to other settings and to find real-world situations that might give us leverage on these questions—e.g., episodes in which novel parties enter the legislature or elites make a concerted attempt to push false media narratives about formation probabilities not otherwise indicated by expected electoral outcomes or ideological relationships (e.g., the British Conservative party's 2015 warnings voters that a vote for Labour meant a Labour-Scottish National Party coalition). Further, it is important that future research is designed to better understand the potential influence of voter uncertainty—over parties' sizes and

ideologies, as well as history, elite signals, etc.—on how individual voters structure their inferential models, and, voters' capacity in the aggregate to coordinate on viable cabinet alternatives (Cox 1997; Powell 2000).

Finally, if we zoom out from our substantive focus on understanding coalition expectations, this work speaks to larger issues about how voters cope with the complexity of systems with multiple parties in which coalition cabinets are the norm. It adds to a growing progressive research program that has found that voters in these contexts use simple heuristics to make sensible inferences not only about which coalitions will form but to update their left-right images of parties, to forecast what policies different cabinets will enact, and to allocate policy-making and administrative responsibility over parties in (and out of) coalition governments (e.g., Fortunato and Stevenson 2013; Fortunato et al. 2021; Lee, Santoso and Stevenson In press). This stream of research is promising, not only because it is pushing forward our understanding of democratic processes in these systems, but because it is revealing perhaps a greater degree of democratic accountability and responsiveness than scholars may have predicted just a few decades ago.

DATA AVAILABILITY STATEMENT

Research documentation and data that support the findings of this study are openly available at the American Political Science Review Dataverse: https://doi.org/10.7910/DVN/E9PQM1

ACKNOWLEDGEMENTS

We are very grateful for feedback from Despina Alexiadou, Mariyana Angelova, Yuki Atsusaka, Ruth Dassonville, Alejandro Ecker, Laurenz Ennser-Jedenastik, Tim Hellwig, Raimondas Ibenskas, Lanny Martin, Stefan Müller, Wolfgang Müller, Carolina Plescia, Bing Powell, Oli Proksch, Yuki Shiraito, Zeynep Somer-Topcu, Jae-Jae Spoon, Markus Wagner, Daniel Weitzel, and especially Mik Laver. We are also grateful to conference participants at Aarhus Universitet, APSA, Central European University, Copenhagen Business School, EPSA, London School of Economics and Political Science, PolMeth, SPSA, Universität zu Köln, Universität Wien, University of Oxford, and University of Rochester. Any errors in the manuscript belong to the authors alone.

FUNDING STATEMENT

This study was supported by SFB 884 Political Economy of Reforms at the University of Mannheim. Fortunato is also grateful to the Institute for Humane Studies (IHS018110) and the Center for Advanced

Study in the Behavioral Sciences for generous support.

CONFLICT OF INTEREST

The authors declare no ethical issues or conflicts of interest in this research.

ETHICAL STANDARDS

The authors declare the human subjects research in this article was reviewed and approved by the University of Mannheim and the University of Essex, detailes numbers are provided in the appendix.

References

Armstrong, David A and Raymond M Duch. 2010. "Why can voters anticipate post-election coalition formation likelihoods?" *Electoral Studies* 29(3):308–315.

Atsusaka, Yuki. 2021. "A logical model for Predicting Minority representation: Application to redistricting and Voting Rights cases." *American Political Science Review* 115(4):1210–1225.

Austen-Smith, David and Jeffrey Banks. 1988. "Elections, Coalitions, and Legislative Outcomes." *American Political Science Review* 82(2):405–22.

Bäck, Hanna and Patrick Dumont. 2008. "Making the First Move: A Two-Stage Analysis of the Role of Formateurs in Parliamentary Government Formation." *Public Choice* 135(3-4):353–73.

Bahnsen, Oke, Thomas Gschwend and Lukas F Stoetzer. 2020. "How do coalition signals shape voting behavior? Revealing the mediating role of coalition expectations." *Electoral Studies* 66:102166.

Bankert, Alexa. 2021. "Negative and positive partisanship in the 2016 US presidential elections." *Political Behavior* 43(4):1467–1485.

Bargsted, Matias A and Orit Kedar. 2009. "Coalition-targeted Duvergerian voting: how expectations affect voter choice under proportional representation." *American Journal of Political Science* 53(2):307–323.

Bowler, Shaun, Gail McElroy and Stefan Müller. 2022. "Voter expectations of government formation in coalition systems: The importance of the information context." *European Journal of Political Research* 61(1):111–133.

Cox, Gary. 1997. *Making Votes Count.* Cambridge: Cambridge University Press.

Crisman-Cox, Casey. 2021. "Estimating substantive effects in binary outcome panel models: a comparison." *The Journal of Politics* 83(2):532–546.

Cutler, Josh, Scott De Marchi, Max Gallop, Florian M Hollenbach, Michael Laver and Matthias Orlowski. 2016. "Cabinet formation and portfolio distribution in European multiparty systems." *British Journal of Political Science* 46(1):31–43.

De Marchi, Scott and Michael Laver. 2020. "Government formation as logrolling in high-dimensional issue spaces." *The Journal of Politics* 82(2):543–558.

De Marchi, Scott and Michael Laver. 2023. *The governance cycle in parliamentary democracies.* Cambridge University Press.

Döring, Holger and Philip Manow. 2021. "Parliaments and governments database (ParlGov)." *An infrastructure for empirical information on parties, elections and governments in modern democracies. Version* 14.

Downs, Anthony. 1957. *An Economic Theory of Democracy.* Harper and Row, New York.

Duch, Raymond M and Jean-Robert Tyran. 2012. Coalition context, voter heuristics and the coalition directed vote. In *Prepared for the Annual meeting of the European Political Science Association, Berlin, Germany.*

Duch, Raymond M, Jeff May and David A Armstrong. 2010. "Coalition-directed voting in multiparty democracies." *American Political Science Review* 104(04):698–719.

Duch, Raymond M. and Randolph T. Stevenson. 2008. *The Economic Vote: How Political and Economic Institutions Condition Election Results.* Cambridge: Cambridge University Press.

Duch, Raymond, Wojtek Przepiorka and Randolph Stevenson. 2015. "Responsibility attribution for collective decision makers." *American Journal of Political Science* 59(2):372–389.

Eberl, Jakob-Moritz and Carolina Plescia. 2018. "Coalitions in the news: How saliency and tone in news coverage influence voters' preferences and expectations about coalitions." *Electoral Studies* 55:30–39.

Falcó-Gimeno, Albert and Jordi Muñoz. 2017. "Show Me Your Friends: A Survey Experiment on the Effect of Coalition Signals." *The Journal of Politics* 79(4):1454–1459.

Fortunato, David. 2021. *The cycle of coalition.* Cambridge University Press.

Fortunato, David, Nick CN Lin, Randolph T Stevenson and Mathias Wessel Tromborg. 2021. "Attributing policy influence under coalition governance." *American Political Science Review* 115(1):252–268.

Fortunato, David and Randolph T Stevenson. 2013. "Perceptions of partisan ideologies: The effect of coalition participation." *American Journal of Political Science* 57(2):459–477.

Fortunato, David and Randolph T Stevenson. 2019. "Heuristics in context." *Political Science Research and Methods* 7(2):311–330.

Fortunato, David and Randolph T Stevenson. 2021. "Party government and political information." *Legislative Studies Quarterly* 46(2):251–295.

Fortunato, David, Randolph T Stevenson and Greg Vonnahme. 2016. "Context and political knowledge: Explaining cross-national variation in partisan left-right knowledge." *The Journal of Politics* 78(4):1211–1228.

Gigerenzer, Gerd. 2019. "Axiomatic rationality and ecological rationality." *Synthese* pp. 1–18.

Gigerenzer, Gerd and Daniel G Goldstein. 1996. "Reasoning the fast and frugal way: models of bounded rationality." *Psychological review* 103(4):650.

Gigerenzer, Gerd and Daniel G Goldstein. 1999. Betting on one good reason: The take the best heuristic. In *Simple heuristics that make us smart.* Oxford University Press pp. 75–95.

Gigerenzer, Gerd and Thomas Sturm. 2012. "How (far) can rationality be naturalized?" *Synthese* 187(1):243–268.

Glasgow, Garrett, Matt Golder and Sona N Golder. 2011. "Who 'wins'? Determining the party of the prime minister." *American Journal of Political Science* 55(4):937–954.

Gschwend, Thomas. 2001. *Strategic voting in mixed electoral systems.* State University of New York at Stony Brook.

Gschwend, Thomas. 2007. "Ticket-splitting and strategic voting under mixed electoral rules: Evidence from Germany." *European Journal of Political Research* 46(1):1–23.

Hjermitslev, Ida B. 2024. "Punishing the pseudo-opposition: Accountability under a minority government." *European Journal of Political Research* 63(2):774–786.

Indridason, Indridi H. 2011. "Proportional representation, majoritarian legislatures, and coalitional voting." *American Journal of Political Science* 55(4):955–971.

Irwin, Galen A and Joop JM Van Holsteyn. 2012. "Strategic electoral considerations under proportional representation." *Electoral Studies* 31(1):184–191.

Kayser, Mark A, Matthias Orlowski and Jochen Rehmert. 2023. "Coalition inclusion probabilities: a party-strategic measure for predicting policy and politics." *Political Science Research and Methods* 11(2):328–346.

Kedar, Orit. 2005. "When moderate voters prefer extreme parties: Policy balancingin parliamentary elections." *American Political Science Review* 99(2):185–199.

Kinder, Donald R and Nathan P Kalmoe. 2017. *Neither liberal nor conservative: Ideological innocence in the American public*. University of Chicago Press.

Laver, Michael and Kenneth Benoit. 2015. "The basic arithmetic of legislative decisions." *American Journal of Political Science* 59(2):275–291.

Laver, Michael, Scott De Marchi and Hande Mutlu. 2011. "Negotiation in legislatures over government formation." *Public Choice* 147(3):285–304.

Lee, Seonghui, Agustina Haime and Randolph Stevenson. 2019. "Measuring knowledge of parties' legislative seat shares." *Political Science Research and Methods* 7(1):185–196.

Lee, Seonghui, Philip Santoso and Randolph Stevenson. In press. *Heuristic Inference and the Left-Right: An Experimental Analysis of How Voters Form Their Left-Right Images of Parties*. Cambridge.

Martin, Lanny W and Georg Vanberg. 2011. *Parliaments and coalitions: The role of legislative institutions in multiparty governance*. Oxford University Press.

Martin, Lanny W and Georg Vanberg. 2025. *Voter Sophistication and Coalition Politics: How Ideological Awareness Shapes Government Formation in Multiparty Democracies*. Oxford University Press.

Martin, Lanny W. and Randolph T. Stevenson. 2001. "Government Formation in Parliamentary Democracies." *American Journal of Political Science* 45(1):33–50.

Martin, Lanny W. and Randolph T. Stevenson. 2010. "The Conditional Impact of Incumbency on Government Formation in Parliamentary Democracies." Unpublished Manuscript, Rice University.

Meffert, Michael F, Sascha Huber, Thomas Gschwend and Franz Urban Pappi. 2011. "More than wishful thinking: Causes and consequences of voters' electoral expectations about parties and coalitions." *Electoral Studies* 30(4):804–815.

Müller, Wolfgang C and Kaare Strøm. 2003. *Coalition governments in western Europe*. Oxford University Press.

Powell, G. Bingham. 2000. *Elections as Instruments of Democracy: Majoritarian and Proportional Views*. New Haven: Yale University Press.

Provins, Tessa, Nathan W Monroe and David Fortunato. 2022. "Allocating costly influence in legislatures." *The Journal of Politics* 84(3):1697–1713.

Selten, Reinhard. 1991. "Properties of a measure of predictive success." *Mathematical social sciences* 21(2):153–167.

Strøm, Kaare. 1990. *Minority Government and Majority Rule*. Cambridge: Cambridge University Press.

Todd, Peter M and Gerd Ed Gigerenzer. 2012. *Ecological rationality: Intelligence in the world.* Oxford University Press.

Todd, Peter M and Gerd Gigerenzer. 2007. "Environments that make us smart: Ecological rationality." *Current directions in psychological science* 16(3):167–171.

Tromborg, Mathias Wessel, Randolph T Stevenson and David Fortunato. 2019. "Voters, responsibility attribution and support parties in parliamentary democracies." *British Journal of Political Science* 49(4):1591–1601.