Heuristics and Coalition Expectations

David Fortunato
University of California, Merced
dfortunato@ucmerced.edu

Randolph T. Stevenson
Rice University
stevenso@rice.edu

Abstract

A new, decidedly empirical body of research has shown that voters in parliamentary systems cast their ballots with an eye toward post-electoral bargaining, seemingly expressing their coalition preferences. But in order for voters to behave this way, they must have well-formed expectations about which coalitions are likely. Whether and how voters actually form such expectations remains an open question. Building on recent multidisciplinary work on heuristics, we provide a theory of how voters formulate expectations of coalition formation outcomes. We then test the implications of our theory with new data from four pre-electoral surveys administered in Germany, the Netherlands, New Zealand, and Norway. We find that voters in general expect the largest party to provide the prime minister and coalesce with a set of ideologically compatible parties to form a legislative majority. These cues, however, are context dependent: voters use them when they are, on average, accurate and not otherwise.

A great deal of recent work in comparative political behavior is aimed at providing an answer to an old question: do voters in coalitional systems (i.e., systems which are typically governed by coalition cabinets) use their votes to try to influence which coalitions form (Blais et al. 2006; Bowler, Karp and Donovan 2010; Duch, May and Armstrong 2010; Kedar 2005; Meffert and Gschwend 2010)? The scholars who have produced this work, unlike many of their predecessors, have taken a decidedly empirical approach to answering it — employing new individual-level survey data and using aggregate electoral data in new ways — and have produced an unambiguous answer: "coalition-directed voting" is widespread in coalitional systems. That is, many voters appear to condition their votes on both their preferences for which coalitions form and their expectations about the likelihood of different coalition possibilities.

This new understanding of coalitional voters marks a sea change in the literature on voting behavior in these systems, which has long been influenced by Downs' (1957) famous skepticism of the ability of such voters to navigate the complexities of their political systems.[1] The empirical case for this change, however, has (in one important respect) run ahead of theoretical work. Specifically, while work on coalition-directed voting often finds that voters take complex coalition possibilities into account in their vote choices, it seldom actually explains how voters manage to accomplish this feat. That is, most work has sidestepped the source of Downs' concerns — how can voters possibly make sense of the myriad coalitional possibilities and come to the kind of reasonable expectations about likely coalition outcomes that are required to cast a rational, outcome-oriented vote?

For decades the consensus answer to this question (asserted or implied without empirical evidence) was simply that they could not. Indeed, to Downs and others it seemed self-evident that in these complex systems, "voters do not vote as though elections were government-selection mechanisms" (Downs 1957, 300). Further, these pessimistic assumptions had very real consequences for

[1]This argument is not confined to prospective voting. Concerns similar to Downs' were raised in the retrospective literature by Powell and Whitten (1993), who suggested that voters would be unable to make the kinds of complicated attributions of responsibility required to hold parties accountable in complex coalitional systems.

how voters in these systems have been studied. For decades, scholars constructed models of voting in complex multi-party systems that were nearly indistinguishable from those applied to majoritarian systems. This theoretical agenda was also reflected (and is perhaps most evident in) the electoral surveys constructed to test it. Up until the late 1990's, one would hardly be able to tell the difference between electoral surveys of voters in majoritarian systems like the United States or the United Kingdom and multi-party coalitional systems like Germany or the Netherlands. While survey designers in all of these systems — reflecting the theoretical consensus in the literature — filled their surveys with questions probing partisan affiliations and attachment, perceptions of party policy positions, and other party-centered variables, they often failed to include any questions probing preferences over governing coalitions, or, indeed, almost any other questions touching on multi-party governance. As late as 1994, Iversen's evaluation of three models of voters in Western Europe never even mentions how voters might move beyond consideration of the policy positions of parties to concern themselves with how these positions get translated into actual policies through the formation of governing coalitions.[2] More generally, before the late 1990's nearly the entire empirical literature on prospective voting in coalitional systems neglected to consider how voters might account for the way in which post-electoral bargaining determines policy outcomes. This is highlighted nicely in Kedar's (2005) summary of how the debate over proximity and directional voting models lost sight of outcome-oriented voting.[3]

[2]It is interesting that despite the near universal acceptance of Downs' pessimistic view among behavioral scholars, institutional formal theorists moved in a quite different direction, assuming a level of voter sophistication in coalitional systems far greater than empirical research on voter knowledge would support. See Austen-Smith and Banks (1988), Baron and Diermeier (2001), and Indridason (2011) for relevant examples.

[3]We are in no way arguing that comparativists in general were unconcerned with how party policy positions get translated into policy in the coalition formation process. Of course, the large literature on coalition formation is testimony to that. Our observation is only that students of voting behavior in these systems were relatively unconcerned with this problem.

However, even as Kedar (2005) highlights Downs' argument that voters must take expected policy outcomes into account when casting their vote, she does not address his doubts over their ability to do so in coalitional systems. In other words, the conversation leapt from, "voters are incapable of forming these expectations" to "we assume voters can and do form them," without ever investigating whether voters are, indeed, able to form reasonable expectations about which coalition may form and, more importantly, how such expectations are born.

Kedar is by no means alone in bypassing this theoretical difficulty. Other studies make a similar leap (e.g. Bowler, Karp and Donovan 2010; Carman and Johns 2010; Meffert and Gschwend 2010). For example, Duch, May and Armstrong (2010) present a theoretical model of "coalition-directed voting" in which voters' preferences over individual parties are weighted by the expected policy output and formation probability of each coalition the party could potentially enter after the election - and test this model by actually estimating and including those formation probabilities and assigning them to voters (using Martin and Stevenson's (2001) empirical model of cabinet formation). In effect, Duch et al. assume that voters are just as capable of predicting which cabinets form as so as political scientists' best quantitative models. Despite such strong assumptions, however, we lack any direct empirical evidence for how reasonable they are; though, common sense suggests that they are implausible.

Thus, in this paper, one of our goals is to test the assumption that voters have well-formed and accurate beliefs about the probabilities different coalitions will form. This goal, however, is not the main point of the paper, which goes well beyond this simple question to ask *how* voters form their coalition expectations. To answer this question, we draw on recent advances in the general study of heuristics to build a simple theory of how voters in most political contexts structure their expectations of which coalitions will form and then test the implications of this theory on new data from four pre-electoral surveys administered in Germany, the Netherlands, New Zealand, and Norway over the period from 2008 to 2012.

While the focus of this paper is exclusively on expectations about coalition formation, we hope that our theory and evidence will contribute to the more general effort to understand how voters

in coalitional systems deal with the complexity of those systems. Our view that they use relatively simple heuristics to do so is consistent with other work in this emerging literature, which has used the same approach to examine how voters make responsibility attributions (Duch, Przepiorka and Stevenson 2014) or update their beliefs about the positions of parties (Fortunato and Stevenson 2013*a*). This work, along with the contribution of this paper, offers the very real possibility of mapping out the most important heuristic tools voters in complex coalitional systems use to form the various complex cognitions that they need to navigate these systems, to cast informed outcome-oriented votes, and ultimately to fill their role in the process of democratic accountability.

Before we lay out our theory of how voters form expectations about cabinet formation, we first briefly summarize how scholars have thought about the problem of expectation formation more generally in order place our theory in its larger context as well as to consider potential alternative explanations.

General Models of Expectation Formation

While a general review of the vast literature on expectations formation is not possible here, our reading of this literature suggests that most work can be sorted into three broad categories:[4]

- Adaptive expectations, which rely on historical trends, occasionally incorporating new information or discounting more distant histories (e.g., if it was sunny yesterday and the day before, I expect it to be sunny today).

- Direct and mediated message-based expectations, which outsource expectation formation — usually to experts or media (e.g., if the weatherman has predicted it, I expect it to be sunny today).

[4]We do not clutter this section with extensive citations to this well-known literature. General reviews can be found in many places. For those new to the literature, Evans and Honkapohja (2001) is a useful starting point.

- Rational or model-based expectations, which utilize some model of how the world works to transform a set of informational inputs into expectations (e.g., my model is that the color of the sky predicts rain or sun, so if there was a red sky last night, I expect it to be sunny today).

It is worth taking a moment to consider what a model of coalition expectations built from each of these perspectives might look like. This will help us to motivate our own approach (which we argue fits into the "model-based expectations" group) and alert us to general alternative approaches that we will want to consider while constructing our empirical tests.

The adaptive expectations model depends on the ability of the individual to observe outcomes and update expectations based on those outcomes. For example, in the canonical application of adaptive expectations to prices, one observes previous prices (perhaps noisily and with greater or lesser time lags) and updates expectations by weighting those observations in some way. Importantly, there are two prerequisites for this kind of expectation formation to be accurate: the target of one's expectation (in this example, prices) must be persistent over time so that the past is a reasonable guide to the future and must occur frequently enough over time for one to recognize this persistence.

Applied to expectations about which coalitions form, it is certainly the case that similar coalitions tend to occur over time (e.g., Armstrong and Duch 2010). However, unlike the persistence of prices, predicting the cabinet that forms after this election based on which one formed after the last election is a highly uncertain and conditional exercise — i.e., the identity of the incumbent cabinet will sometimes be a guide to the identity of the next cabinet and sometimes not, depending on the circumstances (see the discussion below). In addition, and again unlike prices, individuals only observe a limited number of cabinets over their life. Still, the general notion that history may guide expectations about future cabinets is certainly one we need to consider and we will do so in the empirical analysis that follows.

The second major theoretical framework implies that individuals' expectations are formed by aggregating relevant media messages. Models like Zaller's (1992) argue that individual beliefs and attitudes reflect the content of media messages the individual has received and accepted, and these models have been adapted to the generation of expectations in political science (e.g, Duch and

Stevenson's 2011 work on inflation expectations). The specific aggregation mechanism varies in different theories (e.g., online processing versus memory based models) but in all such explanations it is the frequency with which individuals receive and accept specific messages that drives their beliefs. Thus, applied to expectations about coalition formation, we simply expect individual's expectations to reflect the relative frequency with which different coalition options are mentioned in the media. Consequently, while we propose a different expectations model below, in the empirical analysis we do try to account for the distribution of media message about different cabinet possibilities.

Finally, rational or "model-based" explanations of expectation formation simply posit that individuals have a more or less elaborate model in their heads that they use, along with appropriate informational inputs, to form expectations about what will happen given those inputs and the model.[5] In the case of expectations about cabinet composition, this would mean voters have some model of the coalition formation process in their heads that maps observable informational inputs (e.g., the sizes of the parties, the ideological positions of the parties, etc.) to expectations about formation probabilities. One appealing implication of such theories, as applied to cabinet formation, is that they do not require the kind of often repeated events that adaptive expectations models seem to require and so can be readily applied to novel situations (i.e., when a new party enters the arena, or gets much bigger or smaller, or changes their policy positions).[6] Likewise, an empirically useful

[5]Much ink has been spilled about the nature of the information set that goes into such models and what exactly the properties of that information set (and the model) must be to be called "rational;" but, none of that need concern us here, since the relevant point of this literature for us is not really whether expectations are rational (in the sense of meeting whatever requirements one wants to impose for that to be true) but whether they are *model-based*.

[6]This advantage is one of the principal reasons behind Knight's (1992) argument that social institutions more generally — thought of as shared expectations about appropriate behaviors in social situations — should not be regarded as stemming from empirical regularities (i.e., adaptive processes) but rather as "models" that people carry with them about how others will behave in given situations.

aspect of "model-based" theories is that the proposed model is usually distinct from its informational inputs. Thus, one can get particular empirical leverage on identifying whether people are using a particular mental model when one knows that the informational inputs they are using are *wrong* — in that case one can look to see if individuals' expectations are also wrong "in the right way" — i.e., in the way that would be predicted if they are using one's proposed mental model to process that mistaken information into expectations. In the empirical analysis below, we will take advantage of the relatively frequent disconnect between objective information relevant to our proposed mental models of coalition expectations and people's subjective perceptions to get exactly this kind of empirical leverage.

Model-based theories can obviously have heavy informational and cognitive requirements. As we noted above, for example, the Duch, May and Armstrong (2010) model of coalition-directed voting implicitly assumes a model of expectation formation that requires voters to gather all the information that a political scientist might use to predict coalition formation (sizes, ideology, minimal winning status, median inclusion, and much more) and then combine it optimally (via appropriate "regression" weights) to produce expectations. However, it is exactly this kind of explanation of expectations that Downs (1957) rejected. Or more accurately, because such models seemed implausible, he rejected the whole idea that coalitional voters could form accurate expectations (by any means).

A model-based approach, however, need not be so demanding. It is also possible that individuals apply simple heuristics, with limited informational inputs, to form expectations in ways that mimic the accuracy of fully rational models of expectation formation (Gigerenzer et al. 1999). Indeed, a great deal of recent work (some of which is discussed below) has taken up the challenge of rigorously developing explanations of the conditions under which a given heuristic will be used. Such explanations begin by identifying the information sets and rules that constitute plausible heuristics in a given situation and then evaluate their likely usefulness, in a given context, under the principal that the most useful heuristics are those that are cheap, simple, and accurate. This provides a theoretical road map for both identifying plausible heuristics and specifying a priori which ones are

most likely to be used (which, of course, allows one to specify falsifiable hypotheses) — this is the task we undertake below.

A Heuristic Theory of Expectations of Coalition Formation

Our model for how voters form expectations about which coalitions will form relies on the general idea that voters use heuristics to substitute for the complex calculations and high information demands this task would otherwise require. In general, we follow the approach to exploring heuristics developed by Gerd Gigerenzer and the Adaptive Behavior and Cognition (ABC) group at the Max Plank Institute in Berlin. Gigerenzer and Gaissmaier (2011, 455) explain that a heuristic is a decision-making "strategy that ignores part of the [relevant] information, with the goal of making decisions more quickly, frugally, and/or accurately than more complex methods." More concretely, we define a heuristic as a rule that maps a specific information set into a target cognition, in our case, an expectation about the chances a given cabinet coalition will form after an election.

This definition is consistent with Gigerenzer and Gaissmaier's (2011) notion as long as the heuristic relies on a limited information set and a rule that is simple and easy to apply. Such heuristics are cheap (they do not require extensive information search relative to alternative strategies) and simple (the rules map information to outputs without convoluted or complex reasoning — like applying regression weights). In addition, Gigerenzer and Gaissmaier argue a heuristic will only get used if it is (in addition to being cheap and simple) also, on average, accurate in the context in which individuals learn the heuristic (or at least accurate enough relative to the costs of alternative strategies).

This idea of average accuracy in context is important to understand, since it will play a large role in our discussion of potential heuristics voters may use to form expectations about coalition formation. Heuristics, by their nature, are short-cuts that allow an individual to produce a target cognition in a given circumstance without having to know all the details of that circumstance. Consequently, if the accuracy of a heuristic in a given context changes dramatically based on regular fluctuations in the circumstances in which an individual needs to apply it, the heuristic will not

be useful. For example, voters must often form coalition expectations in the context of a coalition formation opportunity (say after an election). Suppose that to do this they adopt a simple rule: "give a high probability of formation to the incumbent cabinet and a low probability to everyone else." As Martin and Stevenson (2010) showed, even within the context of a given country, the accuracy of this rule will be highly dependent on the circumstances that led to a particular formation opportunity (whether the previous cabinet fell due to conflict, among other things). But, of course, the whole idea of the heuristic is to free the voter from the need for such circumstantial knowledge and so this kind of "circumstantial accuracy" is not what we mean by the idea that heuristics need to be accurate on average within the context in which they are learned. Finally, it should be clear that "the context in which a heuristic is learned" can vary (in our case from country to country or perhaps slowly over time within countries) and that accuracy must be judged within this context. Practically, this means that while we are interested in identifying heuristics that apply across countries, we recognize that voters in different national contexts may come to use different heuristics. In what follows we will try to distinguish between the context in which heuristics are learned and the specific circumstances in which a target cognition is produced by referring to the former as "national context" and the later as a specific circumstance.

With these conceptual distinctions in hand, our search for the heuristics individuals use to form expectations about cabinet composition will focus on identifying informational inputs that are cheap, rules that are simple to apply, and combinations of inputs and rules that result in expectations that are reasonably accurate predictors of the coalitions that actually form in national contexts.

The rule: "take the best"

A great deal of research has been conducted on a handful of simple rules that seem to be part of a large number of different heuristics. Examples, of these rules include the "recognition" rule, Dawes' Rule, the "take the last" rule, and the "fluency" or "speed" rule (Gigerenzer and Gaissmaier 2011; Gigerenzer and Goldstein 1996).

The "*take the best*" rule is perhaps the most researched decision rule and is also one of the

most pervasive in actual applications. It is appropriate in contexts in which one must evaluate a number of alternatives on some target criterion where the alternatives have a set of observable characteristics ("cues") correlated with the target criterion. For example, suppose that a consumer wants to purchase some apples and her target criterion is taste. In order to select a set of apples, she forms expectations about the "taste value" of each available alternative. Assuming the shopper is unwilling to take a bite out of each apple, taste is unobservable, though it is likely correlated with a set of cues that *are* observable (such as color, firmness, and aroma). The *take the best* rule simply says to (A) rank order the available cues in terms of their accuracy at predicting the target criterion, and (B) assign higher criterion values (e.g., an expected taste value) to alternatives with higher values for the first cue (e.g., apples that are deeper red), (C) stop at this point if the criterion value of the items are sufficiently differentiated for one's purposes (e.g., if one is choosing three apples, the criterion differentiates three alternatives clearly from the others), and (D) repeat step (C) for each subsequent criterion until meeting the stopping rule.[7]

In our apple choosing example, where, say, color is the best predictor of taste, followed by aroma, and then firmness, the shopper would assign apples with the brightest colors the highest expected taste values. If there were more apples with bright colors than she needed, she would attempt to further differentiate the apples in the brightly colored group by assigning higher or lower expected taste values based on aroma, and so on. At each cue, the chooser "takes the best" until the universe of alternatives is sufficiently differentiated for the task at hand (which may require a single choice, the choice of several alternatives, ranking all of the alternatives, or ranking a subset).

[7]This formulation of the *take the best* heuristic is more general than the one usually found in the literature. Specifically, our formulation applies when the goal is not necessarily to make a single choice among alternatives (choose one apple) but to make more complicated evaluations (choose more than one apple, rank some of the apples, etc.) In the case of coalition expectations, we suspect that the coalitional voter's task is closer to the latter: to talk about politics and/or vote in a satisfying way, she likely needs to rank some subset of alternatives rather than simply decide which alternative is most likely.

Gigerenzer and Gaissmaier (2011) argue that, in many cases, a small number of cues, or even a single cue, is sufficient to adequately differentiate alternatives. Further, it is important to note that such processes are likely to occur subconsciously, with the crucial information about the relative accuracy of cues absorbed in a variety of ways from the environment.

Applied to the question of how likely different coalition cabinets are to form, the rule says one should rank order the cues describing potential cabinets (i.e., characteristics like their size, ideological compactness, minimum winning status, etc.) according to their predictive accuracy and then evaluate each potential cabinet on the first cue. If this cue results in expectations about formation that are sufficient for one's purposes (e.g., to cast a vote or talk sensibly about politics) then stop. Otherwise, use the next most predictive cue to modify (and differentiate) expectations further.

The concept of how well cues differentiate between alternatives deserves some attention here because even if a cue is generally accurate, it will not be useful if it does not discriminate amongst alternatives. What it means to discriminate, however, depends on the task to which the *take the best* criterion is being applied. There are two different cases: one in which the task is to select one or several alternatives from a set and one in which the task is to assign criterion values to all (or some subset of) the alternatives. In the first case, we would think of a cue as having high discriminatory power the smaller the number of alternatives that have the highest value of the cue (as long as there was at least one such alternative). In the second case, however, we should think of cues as having more discriminatory power if the proportion of alternatives that they differentiate from one another is larger. In the best case, such a cue will assign different expected criterion values to each alternative. In contrast, a cue with poor discrimination will leave many alternatives with the same value — which may be no value at all. For the problem of coalition expectations, incumbency and ideological similarity are ideal examples of this difference: incumbency identifies a single alternative, whereas many alternatives will have distinct values for ideological similarity.

There is also a relationship between the costs of collecting information about and applying a cue and its discrimination. Specifically, when information about a cue is costly to collect (or to apply),

this will result in its value not being known for some or all alternatives (since that information is too costly). Because of this, a costly cue will fail to discriminate and, even if otherwise accurate, will be of little use. For example, minimal winning status is certainly a potential cue for predicting which cabinets form and may also be accurate. However, many voters will find it costly to "collect" (i.e., even if they know the sizes of the parties, they may not know the minimal wining status of potential coalitions, or perhaps even what this concept means) and so it will not discriminate for them. As such, the usefulness of different cues in the *take the best* heuristic are determined by the cost of collecting the necessary information about each alternative and of applying that information in the right way, the accuracy of the cue, and the extent to which the cue discriminates among alternatives.

In most published applications of the *take the best* heuristic, very few cues (often just one) are actually used and we have no reason to expect our application to be different. Indeed, we argue below that many potential cues (i.e., characteristics of potential coalitions) are poor predictors of which cabinets accurately form, are too costly to collect or apply for all alternatives, or lack discriminatory power.

Take the best and the sequential process of coalition formation

Before we turn to identifying the factors that predict coalition formation in general, we need to first consider the question of whether we want to model expectations over potential coalitions directly (as a single choice) or, following much of the empirical and theoretical literature on coalition formation, consider the process as having two sequential parts: a Prime Minister is identified and then that PM takes the lead in identifying a set of cabinet partners. Given the general idea that voters are using *take the best*, the later strategy supposes that the voter forms expectations by applying a "nested" set of *take the best* rules. That is, she first applies a *take the best* rule to the question of which party is likely to be the PM and then, having identified a PM, she applies a second round of *take the best* to potential cabinets that include this PM.

There is much to recommend segmenting expectations about potential cabinets in this way. First,

this sequential process appears to be close to the actual process by which many cabinets form: A PM-designate is identified, who then seeks to form a cabinet that he or she will lead (e.g. Warwick 1996; Müller and Strom 2003). Further, making this distinction simplifies the cognitive process of mapping information about parties into expectations (as we will see below) because it vastly decreases the number of coalitions to consider (only those that are led by a viable PM candidate) and allows questions about who will get into a cabinet as a partner to be made conditional on the identity of a PM (i.e., conditional on the *relationship* between potential partners and the PM).

In addition, there is evidence that voters actually do think about cabinet formation as a two-step process, or at least get information about potential cabinets in this way. Perhaps the best evidence lies in media coverage. Popular media cover cabinet formation episodes almost exclusively in terms of formateurs (sometimes referred to as "cabinet negotiators") and potential partners. This language is pervasive in both domestic coverage (e.g. Hankel 2012) and foreign coverage, even in countries that are unfamiliar with coalition politics (e.g. Donadio and Kitsantonis 2012). Given this, the sections below first ask which cues will be employed in a *take the best* model of prime minister selection and then ask which cues will shape expectations of which parties join the PM in cabinet.

Expectations about who will become Prime Minister

Given our idea that voters use a *take the best* strategy to form expectations about which party will be the PM, we need to identify a set of potential cues they might use and consider whether each of these cues is likely to be accurate, discriminatory, and/or relatively easy to know for each party. Fortunately, there is no shortage of theoretical and empirical work exploring which parties are most likely to win the premiership (Bäck and Dumont 2008; Glasgow, Golder and Golder 2011; Stevenson 1997; Warwick 1996). Nearly every article investigating the allocation of the premiership tests some hypothesis regarding size, median status, or incumbency. Below, we assess the accuracy, simplicity,

costs, and discrimination of each of these cues.[8]

Which party will be the largest?

The largest party cue is obviously highly discriminatory in the context of choosing a single party that is most the likely to be PM. It is also extremely accurate both within and across national contexts. In our extension of the Strøm, Müller and Bergman (2008) data, the largest party becomes PM in nearly 77% of selection opportunities.[9] This number is even greater in post-election selection opportunities. Its accuracy (unrivaled among other cues) is matched by its frugality: voters need only have expectations for which party is most likely to win a plurality of seats — no other information on any other party is necessary to discriminate a single alternative. Further, media coverage of the projected plurality winner tends to dominate election news, making it very easy for voters to inexpensively gather (or be inadvertently exposed to) this information. Across the four surveys we use in the empirical analysis later in this paper, the overwhelming majority of voters were able to correctly predict which party would win the most votes. Over 70% correctly predicted the largest party in New Zealand (2008), over 84% in Norway (2009), and a robust 95% in Germany (2009). Even in the 2012 Netherlands election — where the plurality was hotly contested to the last minute — over 50% of respondents made the correct prediction. Indeed, an additional 37% believed that the PVDA would be the largest — the party that turned out to have the next most seats, with less than a 1% difference in the two parties' final seat shares. In sum, the largest party cue is discriminatory in the context of identifying a single most likely PM, accurate, simple to apply, and requires a readily available and easy to understand piece of information.

[8]There are several other (less common) hypotheses regarding selection of the prime minister that have been explored in the literature. (e.g, Glasgow, Golder and Golder 2011). We do not consider these here because they are either rare in general or are characteristics of parties or contexts that do not exist in our sample.

[9]We extend their data on prime minister selection to 2010 for: Austria, Belgium, Denmark, Finland, France, Germany, Iceland, Ireland, Italy, Netherlands, New Zealand, Norway, and Sweden.

Which party is median?

The party of the median legislator obtains the Prime Ministry a little over half the time. While this cue is relatively accurate and discriminatory (i.e., isolates a single alternative in one step), it is much more difficult to apply (and requires more information) than a cue like which party is the largest. This is because the cue requires the synthesis of two types of information: the ideological location of the parties as well as their sizes. In order to identify the party of the median legislator, voters must line up the parties in the ideological space, and then use their sizes to determine the location of the median legislator. Though we know that voters are adept at establishing an ideological rank-ordering of the parties (Fortunato, Stevenson and Vonnahme 2014), as well as a rank-ordering of size (Fortunato, Lin and Stevenson 2014), it is reasonable to wonder if the average voter understands the concept of the median party and can apply it successfully. Thus, while it is, or course, an empirical question whether they use median status as a cue in choosing prime ministers, our expectation is that they will not, preferring instead the simpler largest party cue.

Which party is incumbent?

The incumbency cue correctly predicts the prime minister nearly 67% of time — more accurate than the median party cue and just as discriminatory (i.e. it discriminates to a single alternative), but still less accurate than the largest party cue. This figure, however, over-simplifies the complex, conditional nature of incumbency. That is, previous research has provided evidence that the incumbency advantage is highly dependent on the specific circumstances of a given formation opportunity (such as economic performance and changes in seat share), as well as the institutional or national context in which the formation episode occurs (Martin and Stevenson 2010; Glasgow, Golder and Golder 2011). Thus, incumbency is likely to vary in its accuracy quite a bit over formation episodes within a given country (and also across countries) and so will either require a lot of circumstantial information for voters to apply accurately if they take the conditionality of the cue into account, or will be relatively inaccurate if they do not take its conditionality into account and just assume incumbent prime ministers will be returned.

Taken together, our brief review of the relevant literature suggests that the largest party cue is the least costly in addition to being the most discriminatory and objectively accurate. Because it is the simplest, cheapest, and most accurate of our candidates, we expect that it will be the most widely applied.

Expectations about who joins the prime minister in cabinet

After voters have identified a PM, their next task is to ask which parties will join that PM in cabinet (or how likely different cabinets containing the PM will be). Again, we posit that they use a *take the best* strategy in making these judgments and so we need to identify a set of possible cues they might use and evaluate their accuracy, discriminatory power, and costs. Just as there is no shortage of scholarship studying which party will become the prime minister, there also no shortage of research investigating which parties will join the prime minister. Drawing on the deep literature on coalition formation, we can identify several candidate cues that voters may use in formulating their expectations.

Ideology: prime ministers select sets of ideologically similar partners

Ideological compatibility is a highly accurate predictor of which parties will join the prime minister in cabinet in both partner selection models (Warwick 1996; Martin and Stevenson 2001) and single-step cabinet formation models (Martin and Stevenson 2001, 2010). Further, while there appear to be some differences in the importance of ideological compactness across countries, no one has argued (or shown) that its impact is conditional on the specific circumstances of a formation opportunity. Further, almost unique among the cues we will discuss below, this cue discriminates among all coalition possibilities.

It is also the case that the informational inputs required by the cue, while certainly not trivial, appear to be widely known. Fortunato, Stevenson and Vonnahme (2014) is only the most recent study to show that voters are remarkably well informed about the relative left-right ideological positions of parties. Indeed, there is a very active sub-literature in comparative politics devoted

entirely to tracing the connection between party ideology and voters' perceptions of these ideologies (e.g., Adams et al. 2004). Importantly, this literature makes the point that left-right ideology itself is a short-cut that voters can use to simplify a variety of otherwise demanding tasks they face as citizens (e.g., predicting the specific policy positions of parties and orienting themselves affectively to the range of parties in multiparty systems — Knutsen 1995). Consequently, voters may well have invested in understanding the relative left-right positions of parties for a number of reasons and so have that information cheaply available for the task of forming coalition expectations. Taken together, this suggests that the ideology cue is likely to be simple, cheap (at least relative to its other merits and the alternatives described below), discriminatory, and accurate.

History: prime ministers select sets of partners they have governed with previously

A history of co-governance has a positive effect on the probability of coalition formation. Martin and Stevenson (2010) provide evidence that the greater the history of co-governance among parties within a potential coalition, the more likely that coalition is to form, all else equal. Likewise Armstrong and Duch (2010) argue that historical patterns of formation are the principal way voters predict which future coalitions form. Further, this cue is accurate and simple to apply; voters need only compare the histories of co-governance among the potential partners (including the PM) and assign formation probabilities accordingly. It can also can be discriminatory (depending on the specific pattern of historical cabinet formation). What is less clear, however, is whether voters actually possess historical information of this kind. While electoral surveys have seldom probed respondents' knowledge of historical patterns of co-governance, the evidence that does exist is not encouraging for the idea that voters use history as an important cue in identifying likely cabinet partners. For example, Fortunato, Lin and Stevenson (2014) asked voters in the UK, France, and the Netherlands how frequently given parties had served in cabinets over the last 20 years; and, respondents could essentially make only a binary distinction between parties that *never* get into cabinets and all others. The situation was even worse when voters were asked to describe histories of co-governance (which parties have been in cabinet *together*). The implication of these findings jibes

with the general impression one gets from broader surveys of the historical knowledge of western publics (e.g., Buckley 2011): this kind of historical information is costly to collect and retain and so it is unlikely that large numbers of voters will have it on hand when they form their expectations about which cabinets are likely. This, in turn, will make this cue a poor discriminator, even though it is a fairly accurate predictor.

Median: prime ministers will select sets of partners that include the median party

The idea that the median party will be included in the cabinet appears frequently in the theoretical literature on cabinet formation, but a careful review of the empirical record reveals that in fact it is not actually an accurate predictor of which cabinets form in many national contexts (Glasgow, Golder and Golder 2012; Martin and Stevenson 2001, 2010). Further, similar to the median cue for PM expectations, it is clear that the level of information and sophistication needed to apply this cue is high compared to other cues, in that one needs to mix information about size and ideology as well as understand the concept itself. Finally, the cue lacks discriminatory power since, even though all potential coalitions will either contain the median or not, the cue does not discriminate within these groups. Our expectation then is that this cue is unlikely to be useful to most voters using *take the best* strategies to form their expectations.

Minimal winning: prime ministers select sets of partners that will be minimal winning

The minimal winning status of a potential cabinet is reasonably good predictor of the chances a cabinet will form both within and between national contexts (Martin and Stevenson 2001, 2010). However, it is also one of the most informationally demanding cues, as it requires that voters know (or can predict) the seat shares for each party in the legislature and that they understand and can calculate whether each coalition is minimal winning. Recent survey work by Duch and Tyran (2013) on Germany and Denmark suggests that voters lack these tools. This is unsurprising, since this is exactly the type of burdensome calculation that Downs (1957) had in mind when he dismissed the idea of coalition-directed voting altogether. Additionally, since the cue assigns potential coalitions

to only one of two sets (as with median inclusion), its discriminatory power over the whole set of alternatives is necessarily limited. Thus, our expectation is that few voters will incorporate minimal wining status in forming their expectations.

Majority: prime ministers select sets of partners that will be a majority

Empirical research on coalition formation is essentially unanimous in the finding that minority cabinets are less likely to form than majority cabinets, though the effect seems to be conditional on national context (e.g. minority coalitions are less likely where there are investiture rules and institutionally weak oppositions, see Martin and Stevenson 2001, 2010; Strøm 1990). For example, there are clearly cases in which majority status is a very poor predictor of the composition of governing cabinets (e.g., Norway). That said, where it is accurate it is a likely cue for voters to use. Of all the size-related cues voters might use to identify cabinet partners, it seems clear that this is the easiest to apply. Certainty, majority status, as a concept, is easier to grasp and apply than minimal-winning status and median status. Of course, it is also important to consider the fact that majority status is a relatively blunt instrument in that it will not discriminate among alternatives other than putting them into two groups: majority and minority. Overall, we expect that, given the cue's low informational cost, it will play some role in expectations formation, but only in those contexts in which it is relatively accurate.

Incumbency: prime ministers select sets of partners that are incumbent

Incumbent coalitions re-form fairly often; but, as Martin and Stevenson (2010) emphasize, the chance of re-formation is highly conditional on both the national context and the specific circumstances of the formation opportunity in question. Specifically, they show that an incumbent cabinet is more likely to form only when it competently managed the economy and did not terminate under internal conflict. Further, even these effects of incumbency are conditioned on national institutional contexts (e.g., the existence of continuation rules). Consequently, understanding when to apply the cue is difficult, requiring detailed knowledge of the specific circumstances of a given cabinet forma-

tion. Additionally, to apply an incumbency cue to the cabinet as a whole voters must know the identity of all incumbent parties. However, in the few studies that have probed whether voters can identify the membership of the incumbent cabinet (e.g. Fortunato and Stevenson 2013*b*; Fortunato, Lin and Stevenson 2014) the results show that many voters have trouble doing so.[10]

Summary of cues for selecting the set of cabinet partners

All but two of the cues we reviewed above suffered from problems that make us skeptical that they would find widespread use. Many voters simply do not know cabinet histories, or even the current identify of incumbent partners, and it seems unlikely they employ complex concepts like minimal winning status. Finally, the accuracy of other cues, like incumbency, are highly conditional on the specific circumstances of a given formation opportunity and so require knowledge of those circumstances to apply productively.

In contrast, two cues, majority status and ideological compactness, are more simple to apply, require informational inputs that appear to be comparatively frugal, and are, in many national contexts, among the most accurate cues. Thus, we expect many individuals to rely on these cues. That said, we expect the accuracy of the majority status cue, and to a lesser extent the ideological compactness cue, to be conditioned on national context. Examining the empirical record shows that majority status is not an accurate predictor of formation probabilities in all countries (e.g., Denmark and Norway have a long history of minority cabinets), nor is it the case for all countries that ideological compactness strongly increases a coalition's chances for formation (e.g., Belgium has a history of ideologically disparate coalitions and the Netherlands, in recent years, has had few ideologically connected coalitions). We expect that these sorts of differences in the accuracy of cues in different national context will be reflected in voters' coalition expectations — with majority status and ideological compactness playing a smaller role where they are less accurate predictors of coalition

[10]Netherlands: 1998 25%, 2012: 40%; UK 2012: 60%; New Zealand 1998: 25% and 10 Eastern European cases ranging from 55% to 70% when DKs are omitted (so these are likely closer to the other reported cases when DKs are included).

formation. While the empirical analysis below is based on surveys from only four countries (and so we can not perform a rigorous analysis of differences in heuristic use across national contexts), our sample countries do give us some purchase on these expectations (i.e., both Norway and the Netherlands are in our sample).

Data and Analysis

We are interested in identifying the heuristics people use to form expectations about cabinet formation. Our theory posits that individuals use the *take the best* framework both to identify the likely PM and then the set of parties that will make up the whole cabinet. Further, our review of the potential cues that could be used in in the *take the best* heuristic pushes us toward the a priori expectation that the cues that will be used most frequently (or that will have the best chance of being used) across all national contexts are the largest party cue (for the expectations about who will become PM) and the ideological compatibility cue (for expectations about which partners will join a PM). In addition, we expect the majority cue to be used in contexts in which it is accurate. Thus, the empirical models we evaluate below include measurements of these as well as measurements for the alternative cues we identified above

To test our expectations we have been able to gather original data from four surveys administered to random samples of adults in the week leading up to parliamentary elections in New Zealand (2008), Germany (2009), Norway (2009), and the Netherlands (2012). These surveys contain data previously unavailable to political scientists. Among these data are respondents' expectations for which party will become prime minister, which parties will join the prime minister in coalition (if any), and which party will win the plurality of seats in parliament. In addition to these new items, respondents were also asked to locate each of the parties on the general left-right policy dimension. Using these data, we can directly test our expectations.

Empirical specification

Identifying the degree to which our expected cues are used has two components: first, we must account for other potential model-based cues, precisely the set of cues discussed above. Second, we must account for the possibility that expectations are being driven not by a model-based process, as we argue, but by an adaptive expectations or a direct messaging process. Here we discuss the variables included to account for these processes.

Adaptive expectations

Accounting for adaptive expectations is fairly simple. As discussed above, voters who form expectations in this way use previous cabinet compositions to infer future compositions. Thus, in order to control for this possibility, we need only include measures of previous outcomes. In the prime minister selection model, we include the proportion of previous PM selection opportunities that each candidate party has provided the prime minister.[11] In the partner selection model, we borrow Martin and Stevenson's (2010) measure of cabinet participation history, "familiarity," the discounted, weighted proportion of months that the potential partners have governed with the expected prime minister in the post war period. Further, we include incumbency indicators in both models. Together, these variables should be sufficient to "block" any direct and "backdoor" effects of this type of expectation formation on the estimation of our focal variables (Pearl 2000).

Direct messaging

Direct messaging is not quite as straightforward a process to account for adaptive expectations, but it is still feasible to do so. When we think about all of the potential sources of direct messaging that may influence voter expectations of which coalitions are likely to form, there are effectively two: direct communications from parties about their intentions to coalesce with a given party (in the form of pre-electoral pacts or anti-pacts) and media messages. It is important to note that

[11] We provide description of all of our variables as well as descriptive statistics in the online appendix.

the second should subsume the first; however, none of our elections had formalized pre-electoral coalition pacts.

The most straightforward way to capture the media message would be to quantify the frequency with which all the different coalition possibilities are mentioned in the media in the period leading up to the elections we surveyed. However, doing so is costly and since the role this variable plays in our analysis is only that of a control, we did not immediately attempt to collect these data for all of our cases. Instead, we first built this measure for two countries and examined if, in these two cases, it made any difference to the results of our focal variables. As it turns out, it did not (we report these results in the online appendix). In addition, we pursued a secondary (less costly) strategy to proxy for the aggregate media message about which coalitions would from in each of our cases. Specifically, if we assume that the aggregate media message about coalition possibilities reflects the collective wisdom political elites (e.g., journalist, academics, commentators, and politicians) and this collective wisdom is in turn driven by the set of objective variables that are known to predict cabinet formation (and are also likely to be known by political elites), then including measures of these objective determinants of coalition formation can effectively block any path through which a connection between the aggregate media message and expectations could perturb the estimates of our focal variables (Pearl 2000).

In the theoretical section above, we reviewed many of the objective factors that political scientists have shown to impact coalition formation and so we look to these to provide the set of controls we need here. Thus, in the PM expectation model, we include the objective forecasts of which party will be the largest, which party will be the median, the history of PM selection, and the expected number of seats each party will have. In the partner selection model, for each set of potential partners, we include the average objective distance between the PM and each party in the set, whether the objective median party is in the set, the history of co-governance between the PM and each party in the set, measures of the incumbency status of the set, and measures of the size of the set. We provide details about the measurement of these variables in the online appendix.

Expectations about which party will become PM

Table 1 shows the raw proportion of survey respondents who predicted that the party they expected to win the most seats in the coming election would also become the prime minister. Across the four surveys, roughly 81% of the respondents made this prediction. This is a very large portion of the sample and is consistent with our expectations.

Table 1 here

Of course, Table 1 provides only an uncontrolled look at the data. We also estimate statistical models with a wide range of control variables, including measures meant to capture the alternative cues and the alternative theoretical frameworks we discussed above. The dependent variable in these models indicates whether or not a survey respondent believes a particular party will become the prime minister. As such, we use statistical models appropriate for unordered choice situations, specifically conditional logit models. [12]

The covariates we include in these models are: an indicator for the party the respondent believes will be the largest following the election; an indicator for the party the respondent believes will be median; an indicator for the party of the incumbent PM; the proportion of times a party has been selected as prime minister over the last thirty years (replacing this measure with the proportion of time served as prime minister does not substantively change any of the estimates); and an objective estimate of each party's seat share taken from the poll administered most proximate to our survey. In addition, we include objective versions of the largest party variable (also taken from the most proximate poll) and which party is forecast to control the median legislator (where the spatial placements are calculated from the Comparative Manifestos Project according to Lowe et al. 2011 and the seat shares are calculated using opinion polls).[13] Finally, it is possible that voters expectations

[12]Results from mixed logit models, reported in the online appendix, are the same.

[13]Using the actual election results, rather than the most proximate poll information does not change our conclusions. Also, in each election, the party forecast to win a plurality of seats actually went on to win it.

are colored by their preferences, e.g., Meffert et al.'s (2011) "wishful thinking" mechanism. To account for this possibility, we include a variable indicating which party the respondent intended to vote for and a measure of the distance each respondent perceived between herself and her choices.

Table 2 here

The results of our estimation is shown in Table 2.[14] The estimates show very strong support for our expectation that voters rely on the largest party cue. The coefficient on expected largest party status is positive, statistically significant, and exerts the largest effects of every subjective variable in the model. This is all the more impressive when we consider that the objective value of this variable is also included in the model. As we indicated in our discussion of model-based theories of expectation formation more generally, we can get leverage on identifying the models (or rules) people use to form expectations by examining what happens when the informational inputs to these models are wrong. In the current case, our variable measuring respondents' beliefs about which party will be largest remains positive and significant when the objective largest party variable is included along side it. For this to happen, a significant number of voters whose perceptions about which party would become the largest were incorrect must have conditioned their expectation on these incorrect perceptions.

Perhaps even more instructive is the comparison of the predictive power of the statistical model compared to the uncontrolled analysis in Table 2. Across all four surveys, we can correctly predict who respondents think will be prime minister in 81% of cases, using only their beliefs about which party would be largest. Adding eight more covariates, our statistical model only improves this to 83%. This lack of predictive power for other variables is consistent with respondents using *take the best* with only the largest party cue to from expectations about which party will be PM.

[14]We present individual country analyses and other robustness checks in the online appendix. These all show that the estimates of our focal variables are robust all specifications.

Who joins the prime minister in cabinet?

Our expectation from the theoretical discussion of partner selection is that voters will, in most contexts, expect majority coalitions to form and will otherwise rank formation probabilities of sets of coalition partners according to their ideological compatibility with the expected PM. As above, the first step in evaluating this expectation is to have a look at the raw data on expectations about partner selection, which we glean from questions asking which parties the respondent thinks will join their expected PM in cabinet (the question wording is given in the online appendix). Figure 1 plots the average perceived ideological distance between each potential partner party and expected PMs against the percentage of respondents who say the potential partner will join the expected PM in cabinet. For example, the highlighted triangle in the German plot represents (along the y-axis) the proportion of respondents who said that the CDU would provide the PM and also indicated that the FDP would join the CDU in cabinet, plotted against the average distance these respondents perceived between the CDU and FDP (along the x-axis). In each plot, we include a dot for each expected PM paired with each potential partner. The highlighted triangles are the most widely predicted PM and it most commonly predicted partner.

Figure 1 here.

The message of the figure for three of our countries is clear: expectations about whether a given party will join an expected PM in cabinet are strongly conditioned on the perceived distance between that party and the expected PM, with the percentage of respondents naming a party as a partner projected to a given PM to be almost 100% when the ideological distance between them goes to zero. Further, in our theoretical discussion, we pointed out that we should only expect ideological compactness to be a useful cue for forming coalition expectations to the extent that it is an accurate predictor of cabinet composition in a given national context; and that it is clear from work on coalition formation that ideological compactness plays a smaller role in coalition formation in the Netherlands than other countries in our sample. We could also add that, with the recent re-emergence of grand coalitions in Germany, the general accuracy of the ideological compactness

cue has declined in that country. And, as it turns out, Figure 1 shows that coalition expectations are much less responsive to this cue in the Netherlands than our other cases (with Germany the second least responsive case). To add just a bit of rigor to this observation (without putting to much weight on an analysis of only four cases), we can quantify variation across our cases in the accuracy of the ideological compactness cue by looking at the average ideological distance between the parties that actually formed governments together (that formed since 1980).[15] When we do this, we find that the average ideological distance between coalition cabinet members has been the greatest in the Netherlands (by a substantial margin), next greatest in Germany, and smallest in New Zealand and Norway. This, of course, corresponds directly to the strength of the relationships depicted in Figure 1.

With respect to majority status, Norway is the only one of our four cases in which minority cabinets regularly form and so we would not expect the majority cue to be as helpful to expectation formation in this case relative to the others.[16] Given the various ways in which majority status is related to other variables we need to control for in our model, however, it is more illuminating to examine these expectations in the context of a full statistical model rather than with raw data.

We model a respondent's beliefs about the likelihood of different sets of cabinet partners as a discrete choice model in which the alternatives are all the potential sets of cabinet compositions that include the party he or she predicted would become prime minister. The dependent variable takes on a value of 1 for the constellation of potential partner parties that the respondent identified as the likely partners of their chosen PM and 0 for every other potential coalition. Our key independent variable is the average distance the respondent perceives between each potential partner party and

[15]We calculated the ideological distance for each coalition that formed using the government division variable reported in Martin and Stevenson (2010) (we replicated their measure for New Zealand). The general concept can, of course, be measured in a variety of other ways, including looking at the extent of connected coalitions and the like, but the ordering of our four cases is robust to these kind of alternatives.

[16]Norway was governed almost exclusively by minority cabinets from 1985 to 2005

their predicted PM.[17] As above, we include control variables to account for other potential cues and alternative explanations. Specifically, we use Martin and Stevenson's (2010) "familiarity" measure, the discounted, mean weighted proportion of time the members of the potential coalition have governed with the expected prime minister in the postwar period to account for the history cue (as well as adaptive expectations). We also include the CMP calculated median legislator party and an indicator for the respondents' perceived median party. We calculate, using the polls most proximate to our survey administration, various measures of coalition size and include these in separate specifications. These are: indicators for minority and surplus majority coalitions, where "minimum winning" serves as the baseline category, number of parties and number parties squared, and cabinet size (as a proportion of seats) and cabinet size squared.[18] We include an indicator for the current incumbent coalition to account for the incumbency cue as well as an objective measure of ideological distance between the set of potential partners and the prime minister. Taken together, these objective measurements not only account for other cues, but should also proxy for the aggregate media message as we discussed above.[19] Finally, we calculate the mean ideological distance of the partners from the respondent to account for the possible tendency of respondents to simply choose their most preferred parties to join their predicted PM in coalition.

Table 3 here.

The results of our models are given in Table 3. With respect to ideological compactness, the estimates are again consistent with our expectations. In each country, the relationship is in the

[17]To calculate this variable we first measure the respondent's perceived distance between each partner in a given potential coalition and her expected PM. Next, we calculate the mean of those distances over all partner parties in each potential coalition. Our results are robust to other measures, like the maximum distance.

[18]These size controls must be used in separate specifications since there is a great deal of collinearity between them.

[19]Again, we test the robustness of this approach to a more direct measure of the aggregate media message in the online appendix.

expected direction, significant, and robust to different specifications (including those reported in the online appendix). Further, the estimates are ordered as we expected, with the estimate for the Netherlands significantly smaller than the others (and Germany also somewhat less that the other two). Specifically, an increase of two standard deviations (from the mean) of average perceived ideological distance from the prime minister reduces the log-odds of reporting that the coalition would form by only -0.36 in the Dutch context, but by -0.81, -2.25, and -1.44 in the German, New Zealand, and Norwegian models, respectively.[20] With respect two majority status, the results show that in Germany, the Netherlands, and New Zealand, when our respondents think a potential coalition will have a minority of seats, their expectation that it will form is significantly depressed. However, this is not the case is Norway, just as we expected. It is important to emphasize here that there is nothing in how we collect our data which necessarily creates a connection between the importance of a variable (like majority status) in *actual* coalition formation and its (or its perception's) importance in *expectations* about coalition formation. That link comes from the process of expectation formation itself.

Finally, the lack of systematic, robust estimates on many of the other other control variables included in the model is consistent with the heuristic approach to coalition expectations that we posit here: a *take the best* rule with a small number of cues, of which ideology and majority status are the most useful.

Discussion

The literature on prospective, outcome oriented voting had for many years assumed that voters in complex, coalitional systems could not manage the complexity of these systems. However, starting in the late 1990's that view began to change as theoretical arguments backed by careful empirical work demonstrated that "coalition-directed" voting is in fact widespread. Theoretical models of coalition-directed voting, however, tend to make herculean assumptions about the ability of voters

[20]Calculated using the surplus majority-minority models.

to accurately forecast the outcomes of post-election coalition formation (as well as the policy-making process after a cabinet forms). This is usually done with the (typically implicit) caveat that voters act "as if" they can make the theoretically required forecasts; but, so far there has been little effort to explore the actual process that voters might use to produce these "as if" projections. In this paper, we have tried to fill this gap with an explanation rooted in the application of a few simple heuristics and so hope to put the larger theoretical literature on coalition-directed voting on firmer footing (at least with respect to the formation of coalition expectations).

Our theory is rooted in an established social-psychological literature that argues that people are able to accomplish complex cognitive tasks, quickly and accurately, by employing simple heuristics. These rules of thumb may be used to leverage a small existing pool of information — such as the relative ideological locations of political parties — into expected outcomes of highly complex processes — like the formation of coalition cabinets. Drawing on the theoretical and empirical literature on prime minister selection, partner selection, and cabinet formation more generally, we argued that three cues are particularly cheap, simple, and accurate in many national contexts and so expected these cues to impact expectations more consistently than alternatives (and other non-heuristical explanations as well). Specifically, we expected voters in all of the national contexts we explored to think the largest party would becomes prime minister (with other variables adding little additional explanatory power). Next, we expected voters in most national contexts to think majority coalitions were more likely than minority coalitions and that the parties in these coalitions would be those that are ideologically close to the expected PM. Using original survey data, we found robust support for these expectations. Further, where these simple cues have been less accurate predictors of cabinet membership in the past, we expected voters would not consistently reach for them to shape their expectations. Again, though we have a limited number of cases, the data is consistent with this expectation.[21]

[21]This later effect — that individuals manage (probably sub-consciously) to use the appropriate heuristics in the appropriate contexts (where they are accurate) — has been dubbed "ecological rationality" in the wider literature on heuristics (Gigerenzer et al. 1999; Gigerenzer and Gaissmaier

We hope that these results will not only reassure scholars of coalition-directed voting that their theoretical assumptions may in fact be defensible from an "as if" viewpoint, but also encourage models of coalition-directed voting that actually build in formal descriptions of expectation formation that look like the more realistic one described here. Finally, we see this work as only a first step. Specifically, the four surveys reported here are the first to ask respondents a wide range of questions about their coalition expectations as well as their perceptions of party sizes, ideological positions, and roles in the current cabinet and legislature. Information about these kinds of expectations and perceptions is essential to any effort to understand how individuals make sense of complex political systems and we hope that the initial efforts reported here will encourage others to collect similar data across a much wider set of countries and electoral contexts.

2011) and to our knowledge this is the first evidence of the ecological rationality of political heuristics across national contexts.

References

Adams, J., M. Clark, L. Ezrow and G. Glasgow. 2004. "Understanding change and stability in party ideologies: Do parties respond to public opinion or to past election results?" *British Journal of Political Science* 34(4):589–610.

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

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." *Public Choice* 135(3-4):353–373.

Baron, David P and Daniel Diermeier. 2001. "Elections, governments, and parliaments in proportional representation systems." *The Quarterly Journal of Economics* 116(3):933–967.

Blais, André, John H Aldrich, Indridi H Indridason and Renan Levine. 2006. "Do voters vote for government coalitions? Testing Downs' pessimistic conclusion." *Party Politics* 12(6):691–705.

Bowler, Shaun, Jeffrey A Karp and Todd Donovan. 2010. "Strategic coalition voting: evidence from New Zealand." *Electoral Studies* 29(3):350–357.

Buckley, Jack. 2011. "National Assessment of Educational Progress: The Nation's Report Card: Civics 2010." *Retrieved May* 20:2012.

Carman, Christopher J and Robert Johns. 2010. "Linking coalition attitudes and split-ticket voting: The Scottish Parliament elections of 2007." *Electoral Studies* 29(3):381–391.

Donadio, Rachel and Niki Kitsantonis. 2012. "Before Greek Coalition Forms, Questions About Its Resolve.".

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

Duch, Raymond and Jean-Robert Tyran. 2013. "Coalition Context, Voter Heuristics and the Coalition Directed Vote." *Typescript, Nuffield College, University of Oxford* .

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. 2011. "Context and Economic Expectations: When Do Voters Get It Right?" *British Journal of Political Science* 41(1):1–31.

Duch, Raymond, Wojtek Przepiorka and Randolph Stevenson. 2014. "Responsibility Attribution for Collective Decision Makers." *American Journal of Political Science* Forthcoming.

Evans, George W and Seppo Honkapohja. 2001. *Learning and expectations in macroeconomics*. Princeton University Press.

Fortunato, David, Nick Lin and Randolph T. Stevenson. 2014. "Political Knowledge in Coalition Democracies." *Presented at the MPSA Annual Conference 2014* .

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

Fortunato, David and Randolph T Stevenson. 2013*b*. "Performance voting and knowledge of cabinet composition." *Electoral Studies* 32(3):517–523.

Fortunato, David, Randolph T. Stevenson and Greg Vonnahme. 2014. "Context, Heuristics, and Political Knowledge: Explaining Cross-National Variation in Citizens' Left-Right Knowledge." *Presented at the APSA Annual Conference 2013* .

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

Gigerenzer, Gerd, Peter M Todd, ABC Research Group et al. 1999. "Simple heuristics that make us smart.".

Gigerenzer, Gerd and Wolfgang Gaissmaier. 2011. "Heuristic decision making." *Annual review of psychology* 62:451–482.

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

Glasgow, Garrett, Matt Golder and Sona N Golder. 2012. "New empirical strategies for the study of parliamentary government formation." *Political Analysis* 20(2):248–270.

Hankel, Arne. 2012. "Opstelten ontvangt direct leiders VVD, PVV en CDA.".

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

Iversen, Torben. 1994. "Political leadership and representation in West European democracies: A test of three models of voting." *American Journal of Political Science* 38(1):45–74.

Kedar, Orit. 2005. "When moderate voters prefer extreme parties: Policy balancing in parliamentary elections." *American Political Science Review* 99(02):185–199.

Knight, Jack. 1992. *Institutions and social conflict*. Cambridge University Press.

Knutsen, Oddbjørn. 1995. "Value orientations, political conflicts and left-right identification: A comparative study." *European Journal of Political Research* 28(1):63–93.

Lowe, Will, Kenneth Benoit, Slava Mikhaylov and Michael Laver. 2011. "Scaling policy preferences from coded political texts." *Legislative studies quarterly* 36(1):123–155.

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." *American Political Science Review* 104(03):503–518.

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.

Meffert, Michael F and Thomas Gschwend. 2010. "Strategic coalition voting: Evidence from Austria." *Electoral Studies* 29(3):339–349.

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

Pearl, Judea. 2000. *Causality: models, reasoning and inference*. Vol. 29 Cambridge Univ Press.

Powell, G. Bingham and Guy D Whitten. 1993. "A cross-national analysis of economic voting: taking account of the political context." *American Journal of Political Science* pp. 391–414.

Stevenson, Randolph T. 1997. How Parties Compete: Electoral Performance and Government Participation in Parliamentary Democracies PhD thesis University of Rochester.

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

Strøm, Kaare, Wolfgang C Müller and Torbjörn Bergman. 2008. "Cabinets and coalition bargaining: the democractic life cycle in Western Europe.".

Warwick, Paul V. 1996. "Coalition Government Membership in West European Parliamentary Democracies." *British Journal of Political Science* 26(4):471–99.

Zaller, John. 1992. *The nature and origins of mass opinion*. Cambridge university press.

Table 1: Respondents Choosing Expected Largest Party as PM

Survey	Proportion
New Zealand 2008	0.84
Germany 2009	0.91
Norway 2009	0.72
Netherlands 2012	0.87

Table 2: Conditional Logit Model of PM Expectations

Covariate	Parameter	(se)
Expected Largest	2.022	(0.068)
Party Supporter	1.619	(0.098)
Perceived Distance	-1.079	(0.170)
Objective Largest	-0.381	(0.098)
Objective Median	-1.057	(0.121)
Perceived Median	-0.263	(0.085)
Incumbent	0.370	(0.081)
PM History in Selections	1.514	(0.303)
Expected Seats (polls)	9.713	(0.618)
N alternatives	26446	
$ln(likelihood)$	-6789.95	

Table 3: Conditional Logit Models of Partner Selection

Covariate	Germany			Netherlands			New Zealand			Norway		
Average Subjective Distance from PM	-0.386	-0.578	-0.419	-0.131	-0.173	-0.110	-0.763	-1.090	-0.893	-0.617	-0.645	-0.640
	(0.088)	(0.089)	(0.094)	(0.050)	(0.046)	(0.051)	(0.062)	(0.071)	(0.067)	(0.063)	(0.063)	(0.065)
Average CMP Distance from PM	-4.190	-1.847	-3.148	-1.738	0.563	-2.284	-1.920	-2.052	-2.222	-0.433	-0.354	-1.244
	(0.510)	(0.478)	(0.566)	(0.415)	(0.349)	(0.442)	(0.221)	(0.230)	(0.233)	(0.271)	(0.274)	(0.318)
Familiarity	9.866	12.670	15.798	0.460	-4.175	0.400	2.374	-0.613	-9.914	-1.507	-1.773	-0.928
	(1.335)	(1.479)	(1.507)	(0.395)	(0.431)	(0.415)	(0.597)	(0.533)	(1.338)	(0.538)	(0.460)	(0.538)
Incumbent Coalition	2.041	0.745	0.202	0.205	-0.244	0.111	-2.872	-2.948	-1.670	5.289	4.752	4.924
	(0.242)	(0.194)	(0.189)	(0.314)	(0.311)	(0.314)	(1.004)	(1.004)	(1.009)	(0.196)	(0.106)	(0.106)
Contains CMP Median Legislator	-1.556	-0.619	-5.051	0.610	1.462	0.504	-3.102	-2.274	1.673	0.401	0.935	0.414
	(0.260)	(0.288)	(0.370)	(0.082)	(0.089)	(0.083)	(0.377)	(0.364)	(0.740)	(0.110)	(0.108)	(0.108)
Contains Subjective Median Party	1.018	1.154	0.945	-1.026	-0.323	-1.097	-0.248	-0.244	-0.526	-0.412	-0.141	-0.477
	(0.112)	(0.113)	(0.123)	(0.125)	(0.124)	(0.126)	(0.152)	(0.152)	(0.157)	(0.115)	(0.112)	(0.116)
Average Distance from Respondent	-0.544	-0.467	-0.975	-2.597	-0.712	-2.750	-0.605	0.218	-1.026	-1.886	-0.733	-2.048
	(0.126)	(0.138)	(0.138)	(0.109)	(0.120)	(0.114)	(0.140)	(0.177)	(0.151)	(0.150)	(0.169)	(0.153)
Surplus Majority	3.198			-0.040			-0.407			-1.046		
	(0.192)			(0.134)			(0.135)			(0.242)		
Minority	-1.909			-1.495			-1.149			0.826		
	(0.166)			(0.143)			(0.144)			(0.211)		
Number of Parties		6.981			-1.587			2.266			0.845	
		(0.376)			(0.129)			(0.216)			(0.246)	
Number of Parties Squared		-0.986			0.068			-0.406			-0.328	
		(0.058)			(0.014)			(0.033)			(0.042)	
Cabinet Size			50.264			6.145			31.738			24.736
			(2.431)			(1.387)			(3.797)			(2.289)
Cabinet Size Squared			-26.683			-1.443			-33.097			-28.639
			(1.643)			(1.167)			(2.865)			(2.615)
N alternatives	13712	13712	13712	404991	404991	404991	86272	86272	86272	233728	233728	233728
$ln(likelihood)$	-1716.365	-1712.620	-1577.477	-4103.503	-3984.797	-4130.282	-2488.265	-2385.569	-2330.736	-2306.395	-2203.434	-2267.122

Figure 1: Raw Partner Selection Plot

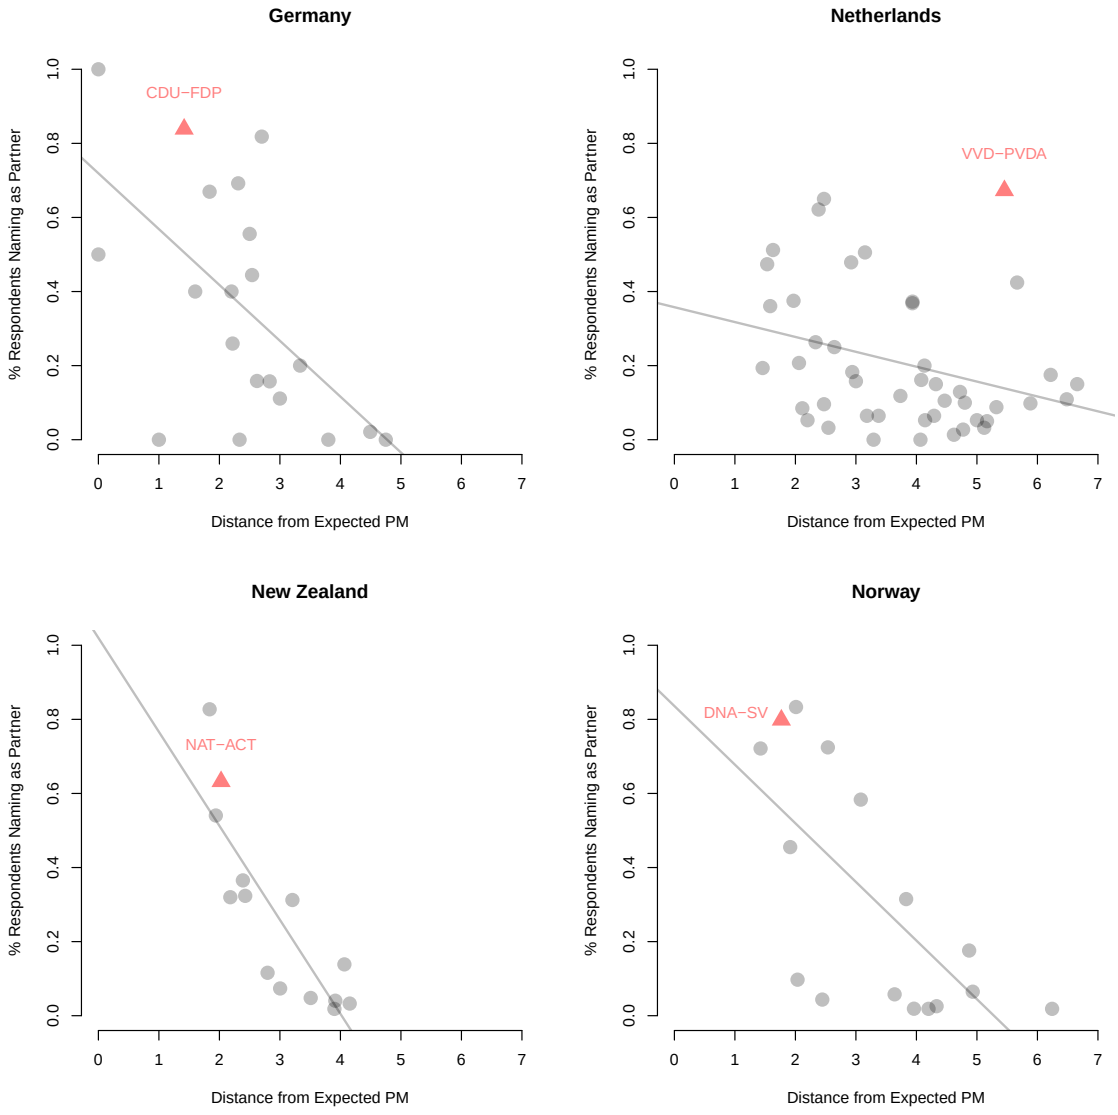

Online Appendix for:

Heuristics and Coalition Expectations

Here we provide information on our survey questions and variables as well as robustness checks for our statistical models.

Principal question wordings

The wordings for our principal questions are given below in English. The Dutch, German, and Norwegian wordings are available from the authors upon request.

Who will become Prime Minister? Single answer question:

> "Thinking about the upcoming election, in your opinion, from which political party will the next Prime Minister come?"

Who will join the PM in cabinet? Single answer question:

> "Following the upcoming election, which party, if any, is most likely join [INSERT EXPECTED PRIME MINISTER PARTY] as partner in the government?"

Who will join the PM in cabinet? Multiple answer question conditional on response other than "none" to above question:

> "Are there any other parties that are likely to join [INSERT EXPECTED PRIME MINISTER PARTY] in the government? If so, which ones?"

Who will be largest? Single answer question:

> "Following the upcoming election, which political party do you think will gain the largest share of seats in parliament?"

Measurement

In this section we present descriptive statistics for the variables in our models and formalize the measurement for the more complex variables.

PM selection model

The descriptive statistics are given in Table 1 and the variables are, in general, quite simple. We discuss them in order of their appearance in the main text table. Note that about half of the variables are subjective, describing the relationship between the respondent and their alternatives, and half are objective, describing characteristics of the alternatives.

Expected Largest: Binary variable indicating the party (choice alternative in the conditional/mixed logit model) that the respondent believed would win the most seats in the coming election.

Party Supporter: Binary variable indicating the party that the respondent intended to vote for in the coming election. Respondents were allowed to give "Don't Know" response or decline to respond.

Perceived Distance: The distance between each respondent's self placement and placement of each of their alternatives. This variable is normalized to a 0-1 scale

Objective Largest: Binary variable indicating the party that was predicted to win the most seats in the coming election according to the most proximate poll. In each case, this prediction was correct — though, in the Netherlands, the plurality winner and first runner up were separated by less than 1%.

Objective Median: Binary variable indicating the party that was predicted to control the

median legislator in parliament. This prediction was calculated by forecasting the election results using the most proximate poll results. Once we forecasted the seatshares, we used the CMP left-right party placements, calculated according to Lowe et al. (2011) to identify the party that would control the median legislator.

Perceived Median: Binary variable indicating the party (or parties) that each respondent placed as the median according to their own left-right placements.

Incumbent: Binary variable indicating the party that held the prime ministry leading to the election.

PM History in Selections: The number of times a particular party has provided the Prime Minister over the last thirty years, divided by the number of selection opportunities over that period.

Expected Seats: The proportion of seats each party is predicted to win according to the most proximate poll.

Table 1: Descriptive Variables for Pooled PM Selection Model

Variable	Mean	SD	Min	Max
Expected Largest	0.123	0.329	0.000	1.000
Party Supporter	0.108	0.311	0.000	1.000
Perceived Distance	0.273	0.228	0.000	0.909
Objective Largest	0.127	0.333	0.000	1.000
Objective Median	0.157	0.364	0.000	1.000
Perceived Median	0.342	0.474	0.000	1.000
Incumbent	0.127	0.333	0.000	1.000
PM History in Selections	0.125	0.214	0.000	0.692
Expected Seats (polls)	0.123	0.113	0.000	0.420

Partner selection model

The descriptive statistics are given in Table 2 and again the variables are generally quite simple.

Average Subjective Distance from PM: The average distance a respondent perceives be-

tween their selected PM and each partner in the potential coalition: $\frac{\sum_{i=1}^{n} |PM - partner_i|}{n}$. For example, say that a German respondent has placed the CDU as a 7, the SPD as a 3, the FDP as a 9, and has predicted the CDU to become PM. For the potential coalition where the SPD and FDP partner with the CDU, this variable would take on a value of $\frac{2+5}{2} = 3.5$.

Average CMP Distance from PM Choice: This measure is contracted identically to the measure above to generate an objective measure of ideological distance from the PM for each potential coalition using CMP left-right estimates of partisan ideology calculated according to Lowe et al. (2011).

Incumbent Coalition: Binary variable indicating the incumbent coalition leading to the election.

Average Distance from Respondent: The average distance the respondent perceives between each party in the potential coalition and *themselves* or: $\frac{\sum_{i=1}^{n} |Respondent - party_i|}{n}$.

Familiarity: This measure is borrowed from (Martin and Stevenson 2010, 509), who explain it as follows:

> "For any two parties, M and K (where M is not equal to K), the familiarity between the two parties at any point in time is equal to the percentage of days (since the formation date of the first democratic cabinet after 1945) that the two parties have participated in the same cabinet up until that point. Familiarity between parties M and K is always equal to 100% if $M = K$. That is, a party is always assumed to be completely "familiar" with itself. We then weight this measure by the expected portfolio shares belonging to the pair of parties. Specifically, the weight for any pair of parties, M and K (where M is not equal to K) is two times the product of their individual portfolio shares. Where M is equal to K (i.e., where the "pair" consists of a single party), the weight is simply the square of the party's portfolio share. Once we have the portfolio-weighted familiarity scores for each pair of parties in a potential government, we then sum across all such

pairs to create the potential government's aggregate familiarity score."

Our measure departs from theirs, however, in that "expected portfolio share" is calculated with forecasted seat shares from the most proximate electoral poll.

Contains CMP Median Legislator: Binary variable indicating that the party that was predicted to control the median legislator in parliament is included in the potential coalition. As above, this prediction was calculated by forecasting the election results using the most proximate poll results. Once we forecasted the seatshares, we used the CMP left-right party placements, calculated according to Lowe et al. (2011) to identify the party that would control the median legislator.

Contains Subjective Median Party: Binary variable indicating that a party that the respondent placed as the median according to their own left-right placements is included in the potential coalition.

Surplus Majority: Binary variable indicating that the potential coalition form which one party could removed without affecting the majority status of the cabinet. The seatshares are forecasted with the most proximate poll. The baseline variable is a minimal winning coalition.

Minority: Binary variable indicating that the potential coalition has less than 50% of the expected seatshare. The seatshares are forecasted with the most proximate poll. The baseline variable is a minimal winning coalition.

Robustness

In our explanation of our explanation of how we control for direct messaging expectations in our empirical analysis, we mention that, for two of our elections we produced a quantitative measure of media messages regarding the likelihood of different coalitions. Here, we present that measure and the analyses including it, and show that our method of controlling for media messages, by controlling for their objective determinants, is sufficient to get clean estimates of our focal variables.

Table 2: Descriptive Statistics for Partner Selection Models

	Germany		Netherlands		New Zealand		Norway		Pooled			
Variable	Mean	SD	Mean	SD	Mean	SD	Mean	SD	Mean	SD	Min	Max
Average Subjective Distance from PM	2.014	1.406	3.339	1.623	2.622	1.757	2.800	1.375	3.064	1.596	0.000	10.000
Average CMP Distance from PM Choice	0.718	0.324	0.412	0.146	0.668	0.245	0.619	0.258	0.508	0.231	0.000	2.231
Incumbent Coalition	0.032	0.177	0.001	0.031	0.004	0.062	0.002	0.044	0.002	0.046	0.000	1.000
Average Distance from Respondent	1.023	0.789	1.576	0.889	1.106	0.851	1.419	0.801	1.462	0.871	0.000	8.000
Familiarity	0.302	0.142	0.201	0.120	0.236	0.157	0.171	0.133	0.197	0.131	0.000	0.832
Contains CMP Median Legislator	0.516	0.500	0.500	0.500	0.502	0.500	0.501	0.500	0.501	0.500	0.000	1.000
Contains Subjective Median Party	0.631	0.482	0.384	0.486	0.474	0.499	0.557	0.497	0.451	0.498	0.000	1.000
Surplus Majority	0.387	0.487	0.440	0.496	0.463	0.499	0.438	0.496	0.441	0.496	0.000	1.000
Minority	0.484	0.500	0.525	0.499	0.498	0.500	0.499	0.500	0.513	0.500	0.000	1.000

An robustness check on media messages

Though we are confident in our strategy of blocking the effects of media messaging, we realize that it requires a bit of faith on behalf of the reader and that some may have lingering concerns. To assuage these concerns we have generated, for two of our elections, a measure of the media's messaging regarding the likelihood of various coalitions forming. For Germany 2009 and New Zealand 2008 we collected every article printed in the countries' highest circulation newspaper (Die Zeit and the New Zealand Herald, respectively) mentioning the names of at least two political parties in the two months leading up to the election (we chose Germany and New Zealand for the reasons of data availability and language familiarity). We then hand coded each article, counting up the number of times, over that two month period, each unique party-pair was mentioned as being either likely or unlikely to coalesce. If an article mentioned a three-party coalition as being likely, for example, the "traffic-light coalition" in Germany, we would code the individual party pairs therein. In this case, the SPD-FDP, SPD-Green, and FDP-Green pairs.

These counts of positive and negative mentions were then used to scale a media favorability rating for each party pair building on the Lowe et al. (2011) method of policy position scaling. More specifically, each unique pair's media favorability is given by $log(\frac{\#\text{ positive mentions } +0.5}{\#\text{ negative mentions } +0.5})$, where

higher values indicate that party pairs that are discussed as more likely to coalesce and negative values indicate that party pairs are discussed as less likely to coalesce in popular media coverage.

In the German case, for example, the highest rated pair was the CDU/CSU and the FDP — the two parties that actually went on to coalesce — and the lowest rated pair was the FDP and Die Linke, the two parties anchoring the right and left poles of the Bundestag, respectively. These values where then used to compute average media favorability ratings for each potential coalition in Germany and New Zealand and were added to the models we presented in the main text.

The results of the replications for Germany are in Table 3 and the New Zealand replication is in Table 4. Again, we estimate a model for each size control placing the estimates excluding and including our media favorability measure side-by-side. In general, there is very little difference between each pair of estimates. In each estimate, our focus variable, the average distance each respondent perceives between her selected prime minister and its potential partners is large, negative, and robust. This also holds for our objective measures of ideological compatibility. Indeed, there are only a handful of estimates that gain or lose significance with the inclusion of media favorability, and just two that change sign. The New Zealand model, in particular, is quite stable. Indeed, the media favorability measure is significant and in the predicted direction in each of the New Zealand models. This stability, we believe, is at least partially a function of the larger number of possible coalitions in New Zealand as compared to Germany. In Germany, where there is a relative small number of possible coalitions, there is a large degree of collinearity between some of the measures, particularly between the objective distance and size variables with or media message variable, which leads to some unstable estimates. Note, however, that subjective distance from the PM remains quite stable and robust throughout all of the estimations and that the cabinet size variables, less effected by collinearity with the media favorability measure, still support the majority cue.

Table 3: Conditional Logit Models of Partner Selection in Germany with Media Controls. Light shading indicates parameters that have lost/gained significance at traditional levels, dark shading indicates parameters that have changed sign with media control inclusion.

Covariate	Surplus/Minority		Number of Parties		Seatshare	
Average Subjective Distance from PM	-0.386	-0.266	-0.578	-0.600	-0.419	-0.377
	(0.088)	(0.092)	(0.089)	(0.088)	(0.094)	(0.095)
Average CMP Distance from PM Choice	-4.190	-4.226	-1.847	-1.470	-3.148	-3.471
	(0.510)	(0.600)	(0.478)	(0.480)	(0.566)	(0.585)
Familiarity	9.866	1.098	12.670	2.253	15.798	20.591
	(1.335)	(1.166)	(1.479)	(1.857)	(1.507)	(2.297)
Incumbent Coalition	2.041	2.562	0.745	-0.310	0.202	0.431
	(0.242)	(0.275)	(0.194)	(0.231)	(0.189)	(0.206)
Contains CMP Median Legislator	-1.556	-0.334	-0.619	-0.577	-5.051	-5.714
	(0.260)	(0.361)	(0.288)	(0.325)	(0.370)	(0.435)
Contains Subjective Median Party	1.018	0.772	1.154	1.204	0.945	0.905
	(0.112)	(0.120)	(0.113)	(0.116)	(0.123)	(0.124)
Average Distance from Respondent	-0.544	-0.618	-0.467	-0.539	-0.975	-0.934
	(0.126)	(0.139)	(0.138)	(0.137)	(0.138)	(0.139)
Average Media Favorability		2.264		0.918		-0.428
		(0.115)		(0.116)		(0.146)
Surplus Majority	3.198	3.547				
	(0.192)	(0.199)				
Minority	-1.909	3.181				
	(0.166)	(0.311)				
Number of Parties			6.981	3.795		
			(0.376)	(0.549)		
Number of Parties Squared			-0.986	-0.591		
			(0.058)	(0.077)		
Cabinet Size					50.264	62.365
					(2.431)	(5.016)
Cabinet Size Squared					-26.683	-33.547
					(1.643)	(3.011)
N alternatives	13712	13712	13712	13712	13712	13712
$ln(likelihood)$	-1716.365	-1478.660	-1712.620	-1680.875	-1577.477	-1573.013

Table 4: Conditional Logit Models of Partner Selection in New Zealand with Media Controls. Light shading indicates parameters that have lost/gained significance at traditional levels with media control inclusion.

Covariate	Surplus/Minority		Number of Parties		Seatshare	
Average Subjective Distance from PM	-0.763	-0.952	-1.090	-1.056	-0.893	-1.014
	(0.062)	(0.069)	(0.071)	(0.072)	(0.067)	(0.071)
Average CMP Distance from PM Choice	-1.920	-1.460	-2.052	-1.082	-2.222	-1.735
	(0.221)	(0.226)	(0.230)	(0.231)	(0.233)	(0.236)
Familiarity	2.374	1.226	-0.613	-2.390	-9.914	-9.019
	(0.597)	(0.622)	(0.533)	(0.585)	(1.338)	(1.347)
Incumbent Coalition	-2.872	-3.650	-2.948	-3.492	-1.670	-2.385
	(1.004)	(1.006)	(1.004)	(1.005)	(1.009)	(1.011)
Contains CMP Median Legislator	-3.102	-2.071	-2.274	-1.378	1.673	1.824
	(0.377)	(0.387)	(0.364)	(0.372)	(0.740)	(0.712)
Contains Subjective Median Party	-0.248	-0.508	-0.244	-0.289	-0.526	-0.673
	(0.152)	(0.157)	(0.152)	(0.155)	(0.157)	(0.159)
Average Distance from Respondent	-0.605	-0.620	0.218	0.339	-1.026	-0.975
	(0.140)	(0.150)	(0.177)	(0.176)	(0.151)	(0.158)
Average Media Favorability		0.636		0.563		0.534
		(0.045)		(0.048)		(0.046)
Surplus Majority	-0.407	-0.338				
	(0.135)	(0.141)				
Minority	-1.149	-1.045				
	(0.144)	(0.146)				
Number of Parties			2.266	0.868		
			(0.216)	(0.238)		
Number of Parties Squared			-0.406	-0.225		
			(0.033)	(0.034)		
Cabinet Size					31.738	28.765
					(3.797)	(3.919)
Cabinet Size Squared					-33.097	-29.248
					(2.865)	(2.842)
N alternatives	86272	86272	86272	86272	86272	86272
$ln(likelihood)$	-2488.265	-2372.852	-2385.569	-2309.877	-2330.736	-2258.028

Robustness checks for main models

PM selection

In this section we present alternative analyses as robustness checks on the models presented in the main text. In the main text we present a stacked conditional logit model of PM selection, so that we may borrow power across surveys to compensate for within-survey collinearity of objective independent variables. Here, we estimate separate models for each survey and identify the collinearities for our readers. Note that our focal variable remains quite large and robust across all models.

Table 5: Conditional Logit Models of PM Selection by Country

Covariate	Germany	Netherlands	New Zealand	Norway
Expected Largest	1.531	2.944	2.003	1.113
	(0.177)	(0.129)	(0.174)	(0.122)
Party Supporter	0.657	1.148	0.907	1.297
	(0.180)	(0.186)	(0.172)	(0.122)
Distance from Respondent	-0.268	-0.540	-3.305	-2.327
	(0.589)	(0.385)	(0.515)	(0.291)
Objective Largest	6.402	0.492	2.548	-5.942
	(5.476)	(0.199)	(4.980)	(0.559)
Objective Median	3.909	-0.302	0.238	**E**
	(2.549)	(0.533)	(1.133)	
Perceived Median	-0.637	-0.062	0.195	-0.435
	(0.179)	(0.181)	(0.223)	(0.164)
Incumbent	**A**	**C**	3.132	**E**
			(3.995)	
PM History in Selections	**B**	1.215	**D**	7.471
		(0.837)		(0.627)
Expected Seats (polls)	-10.203	10.998	5.910	22.113
	(22.612)	(1.625)	(13.226)	(1.904)
N alternatives	4585	7606	5392	8739
$ln(likelihood)$	-282.147	-387.265	-249.129	-646.259

Notes: **A** CDU/CSU collinear with Objective Largest; **B** SPD collinear with Objective Median and CDU/CSU collinear with Objective Largest; **C** VVD, collinear with Objective Largest; **D** National and Labor collinear with Objective Largest and Incumbent, respectively; **E** DNA, Objective Largest, Objective Median and Incumbent all collinear.

The principal concern one could have with the analyses presented in the main text is the possible violation of the IIA assumption required for conditional logit estimation. Though, our diagnostic tests imply that this is not an issue, and Glasgow, Golder and Golder (2012) demonstrate that are very few statistically discernible differences between conditional logit and mixed logit estimations on cabinet formation data, we are, clearly, using very different data here. Thus, we present mixed logit models in this section to put our readers at ease. First, in Table 6 below, we present the same prime minister model from the main text, however, we allow for random coefficients on each included covariate. Comparing this model to the conditional logit in the main text shows that there is precious little difference between the model results, although, there is significant variation in the effects of the expected largest cue, particularly, as we discuss in the main text, Norwegians are less likely to apply this cue than respondents in our other countries.

Partner selection

Unlike the PM selection model, we cannot estimate a mixed logit of partner selection allowing for random coefficients on each parameter — the model is simply too complex. However, we can estimate a series of separate models where each parameter, in turn, is allowed random coefficients. In each iteration of this exercise our focal variables remain in the predicted direction, statistically significant, and substantively robust. The full results of this exercise are available from the authors on request. Here, we present the results of models where random coefficients are allowed on our focal variable, using surplus majority and minority size controls. We choose to present this model rather than one of the others to point out that there is an incredibly small amount of variation on the subjective distance parameters *within countries*.

Table 6: Mixed Logit Model of PM Expectations

	Estimates	
Covariate	Mean	SD
Expected Largest	3.637	2.424
	(0.322)	(0.433)
Party Supporter	1.863	0.964
	(0.152)	(0.497)
Distance from Respondent	-2.370	0.199
	(0.312)	(0.638)
Objective Largest	-0.449	0.311
	(0.160)	(0.415)
Objective Median	-2.314	3.233
	(0.282)	(0.314)
Perceived Median	-0.404	0.079
	(0.140)	(0.620)
Incumbent	0.957	1.010
	(0.149)	(0.374)
PM History in Selections	2.654	0.038
	(0.481)	(0.540)
Expected Seats (polls)	12.955	2.119
	(1.014)	(1.440)
N alternatives	26322	
$ln(likelihood)$	-1762.2205	

Table 7: Mixed Logit Models of Partner Selection: random coefficient on focal variable

Covariate	Germany	Netherlands	New Zealand	Norway	Pooled
Average Subjective Distance from PM (mean)	-0.537	-0.378	-0.889	-0.903	-0.808
	(0.085)	(0.043)	(0.059)	(0.057)	(0.027)
Average CMP Distance from PM Choice	-6.382	-3.756	-2.720	-2.500	-2.925
	(0.446)	(0.375)	(0.206)	(0.240)	(0.120)
Familiarity	4.442	1.835	5.355	1.286	1.719
	(1.061)	(0.268)	(0.446)	(0.311)	(0.166)
Incumbent Coalition	2.279	0.443	-2.901	5.209	2.232
	(0.245)	(0.307)	(1.004)	(0.194)	(0.065)
Contains CMP Median Legislator	-0.218	0.370	-1.620	0.396	-0.345
	(0.191)	(0.081)	(0.167)	(0.114)	(0.048)
Contains Subjective Median Party	1.050	-0.939	-0.038	-0.326	0.164
	(0.114)	(0.123)	(0.151)	(0.113)	(0.057)
Average Distance from Respondent	-0.509	-2.037	-0.229	-0.990	-1.059
	(0.123)	(0.095)	(0.126)	(0.129)	(0.055)
Surplus Majority	3.106	0.058	-0.257	-0.922	0.343
	(0.189)	(0.133)	(0.133)	(0.242)	(0.062)
Minority	-1.856	-2.297	-1.594	-0.183	-1.501
	(0.164)	(0.143)	(0.140)	(0.204)	(0.060)
Average Subjective Distance from PM (sd)	0.005	0.007	0.007	0.007	0.006
	(0.112)	(0.201)	(0.081)	(0.083)	(0.120)
N alternatives	26446	807947	171870	463563	1469826
$ln(likelihood)$	-1766.791	-4382.417	-2648.612	-2536.649	-13016.875

References

Glasgow, Garrett, Matt Golder and Sona N Golder. 2012. "New empirical strategies for the study of parliamentary government formation." *Political Analysis* 20(2):248–270.

Lowe, Will, Kenneth Benoit, Slava Mikhaylov and Michael Laver. 2011. "Scaling policy preferences from coded political texts." *Legislative studies quarterly* 36(1):123–155.

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