

WHAT YOU SEE IS NOT ALWAYS WHAT YOU GET:
BARGAINING BEFORE AN AUDIENCE UNDER MULTIPARTY GOVERNMENT

LANNY W. MARTIN *

GEORG VANBERG **

ABSTRACT. Party elites in coalition governments are acutely aware that the deals they strike will be critically evaluated by their supporters, and that they risk losing support if they are perceived as ineffective negotiators. This has a powerful influence on the bargains parties strike. Because most supporters are unaware of the complex aspects of bargains, and instead rely on simple heuristics to evaluate their most visible features, parties have incentives to meet supporter expectations primarily on easily observable outcomes. To do so, they make tradeoffs on less observable outcomes. This implies that the more visible features of a bargain typically do not accurately reflect the relative success of parties in coalition negotiations. We evaluate our argument using original data on the office rewards and policy risks of portfolio allocation in 16 parliamentary democracies. Our findings support our argument, and they have important implications for the nature of representation under multiparty government.

* Professor, Department of Social and Political Sciences, Dondena Centre for Research on Social Dynamics and Public Policy, Bocconi University. lanny.martin@unibocconi.it.

** Professor, Department of Political Science, Duke University. georg.vanberg@duke.edu.

We wish to thank Christoph Vanberg, as well as seminar participants at Aarhus University, Emory University, Humboldt University, King's College London, the University of Bologna, the University of Essex, the University of Kentucky, the University of Mannheim, the University of Milan, and the University of Zurich for helpful comments on previous versions of this article.

Word count: 12,000.

INTRODUCTION

Coalition government is a ubiquitous feature of politics in most parliamentary democracies.¹ One of its central characteristics is that it requires bargaining and compromise among parties with distinct electoral platforms. Coalition partners must agree on a common policy, and must distribute important ministries and other perquisites of office. What is the nature of these bargains? How successful are parties in securing what they value, and what factors shape their ability to do so? These questions are important not only for academics and political analysts, but also for citizens as they make voting decisions, gauge the performance of their elected representatives, and evaluate the quality of governance.

We offer an answer to these questions by proposing a theory of coalition bargaining that focuses on two features of negotiations that have, to date, not been explicitly addressed. The first is that party elites worry about how the bargain they agree to will be perceived by their supporters. Making this aspect of bargaining more explicit matters because there are different types of supporters on which parties rely. Some follow politics closely, and are aware of the details of coalition negotiations. Others (including the majority of voters) are politically less knowledgeable, and are only aware of the most salient and easily observable aspects of coalition bargains. The second feature we focus on is that parties bargain over a range of “goods,” including policy and offices. Importantly, not all of these goods are equally “visible” to politically less knowledgeable supporters. As we show, the combination of these two features has significant implications for the nature of the bargains parties strike.

Specifically, to secure the continued backing of politically less knowledgeable supporters, parties face incentives to negotiate outcomes that reflect supporter expectations on the most salient, easily observable dimensions of the bargain. At the same time, outcomes that reflect those expectations will typically *not* align with the underlying bargaining positions of the parties. As a result, party elites use side payments of less observable goods to reconcile the overall coalition bargain with relative bargaining power. Parties that are “undercompensated” (relative to their bargaining strength) with respect to the *more visible* aspects of the bargain will be remunerated by their “overcompensated” partners through concessions on more complex, *less visible* aspects of the bargain. Thus, a central implication of our

¹ For reviews of the coalition literature, see, e.g., Laver and Schofield (1998) and Martin and Vanberg (2015).

argument is that the more easily observable features of a bargain typically do not accurately reflect the influence of the various parties on the overall outcome of negotiations. In other words, looks can be deceiving—a conclusion that has far-reaching consequences for the nature of representation under multiparty government.

Empirically, we provide initial support for our argument with a fresh examination of how parties distribute ministerial portfolios. A long-standing tension in coalition research concerns the contrast between a theoretically-grounded expectation about portfolio allocation and a well-known empirical regularity. Descriptively, coalition members tend to receive a share of cabinet ministries that is roughly proportional to the share of legislative seats they contribute to the coalition, a pattern referred to as “Gamson’s Law” (Gamson 1961; Browne and Franklin 1973; Warwick and Druckman 2006; Falcó-Gimeno and Indridason 2013). In contrast, standard bargaining models lead to the expectation that the distribution of spoils during coalition negotiations should reflect the bargaining power of the parties, which typically does not correspond to their seat shares. Our argument reconciles this tension. Specifically, we contend that coalition parties—anticipating voter expectations—have strong reason to favor a “Gamsonian” allocation of the easily observable numerical distribution of portfolios. They align this allocation with the bargaining power of the parties by making side payments with less visible goods. Our analysis of one less observable feature of portfolio allocation—namely, the “policy risk” inherent in the distribution of cabinet ministries—provides clear support for this expectation and, by implication, for our broader theoretical argument.

COALITION BARGAINING AND SUPPORTER EVALUATIONS

Multiparty government requires bargaining and compromise. This poses a challenge for parties. On the one hand, they must come to an agreement and cooperate with their partners on a common agenda. To maintain support, on the other hand, they must convince supporters that they are not overly accommodating, and that they represent their interests effectively (Martin and Vanberg 2011; Fortunato 2019). As a result, the types of bargains elites will agree to depend partly on their expectations of how those bargains will be evaluated by their supporters.

Existing accounts of coalition politics do not explicitly focus on this feature of bargaining (though the assumption that the preferences of party elites are induced by the underlying concerns of supporters (see, e.g., Strøm 1990) does suggest an implicit connection). Our argument puts supporter evaluations at center stage. Doing so allows us to explore the implications of an important fact of life in modern democracies: Not all citizens are equally sophisticated or knowledgeable about politics (see, e.g., Fortunato, Stevenson and Vonnahme 2016). While some supporters, such as highly involved political activists or interest groups, are sufficiently motivated and resource-endowed to follow politics closely, this is not true for all of them. In particular, a long tradition of research in political psychology suggests that many voters—a key source of party support—lack deep political knowledge. Rational ignorance and lack of interest mitigate against detailed assessments of complex political information (Downs 1957; Tullock 1967). Instead, many voters employ readily available information and informational shortcuts—“heuristics,” in the language of this literature—to evaluate the performance of policymakers (see, e.g., Healy and Malhotra 2013; Fortunato and Stevenson 2019).²

Consider the implications for coalition bargaining. These findings suggest that in evaluating how well their party has represented their interests in coalition negotiations, less sophisticated supporters—including most voters—focus on results they can more easily assess using informational shortcuts. Empirical work finds clear support for this idea. For example, Lin et al. (2017) demonstrate that voters employ a “proportionality heuristic”: Voters expect that, in coalition governments, parties will secure shares of cabinet ministries that correspond to their relative size.³ Duch and Stevenson (2013) show that voters rely on a “responsibility attribution heuristic” to draw inferences about policy influence based on the agenda power of particular cabinet members (such as the prime minister). Similarly, Bowler, Gschwend and Indridason (N.d.) show that voters rely on party size as well as ideological centrality as heuristics to estimate the policy influence of coalition parties. These findings have an important implication: The approval of politically less knowledgeable supporters—primarily voters—is likely to depend on how well a party “measures up” against a yardstick supplied by the informational shortcuts they employ.

² While heuristics are often “ecologically rational” (Goldstein and Gigerenzer 2012; Fortunato, Stevenson and Vonnahme 2016), they can also lead to systematic bias and mistaken inferences, incentivizing policymakers to focus on outcomes that voters’ heuristics tap into (Healy and Malhotra 2013, 299).

³ They also show that voters assess fairly accurately party seat shares and portfolio distributions.

Critically, not all the “goods” over which partners bargain are equally susceptible to informational shortcuts. Some are highly salient, information about them is readily available, and the relative success of parties in obtaining them can be evaluated by voters using simple heuristics. For example, voters can easily perceive whether the share of cabinet portfolios received by a party corresponds to its relative size (Lin et al. 2017). Most voters are also likely aware of whether the most prestigious cabinet ministries (such as the Prime Ministership) are awarded to the largest coalition party. Consistent with this expectation, Browne and Franklin (1973) find that the largest party is most likely to secure the PM position. Similarly, for parties focusing on a small set of issues, even less sophisticated supporters may be aware of whether the party obtains the relevant ministry (e.g., whether a Green party secures the Environment ministry). Consistent with this expectation, Bäck, Debus and Dumont (2011) demonstrate that niche parties are typically successful in securing the ministry most salient for them. In contrast, other aspects of coalition bargains are more complex. For instance, evaluating policy compromises on technically challenging subjects, as well as assessing the relative success of parties in influencing these agreements, requires considerable time and effort in gathering and processing detailed information. Consequently, bargaining outcomes on such issues are less susceptible to accurate assessments using informational shortcuts, and are therefore less likely to be “picked up” by the evaluations of politically less knowledgeable supporters.⁴

A simple model. To analyze these implications more formally, we consider a simple model.⁵ We assume that there are two coalition parties negotiating over the allocation of two types of goods. The first good is “simple” in that information about the relative distribution of the good is easy to obtain and evaluate against a heuristic. The second good is “complex,” indicating that the agreement reached by the parties about the relative distribution of the good is far more difficult for outside actors to observe (in a way we make more precise below). A “coalition bargain” consists of an allocation of a share of the simple and complex goods to party A, with the remainder going to party B. We denote a

⁴ Bowler, Gschwend and Indridason (N.d.) show that voters rely on party size and ideological centrality as shortcuts to estimate policy influence. Importantly, as they note, their findings do not speak to the connection between voter expectations and the actual policy compromises coalitions reach. The critical point for our argument is that it is more difficult for supporters to evaluate complex policy against their heuristic, i.e., to determine whether the actual coalition bargain corresponds to supporter expectations.

⁵ We provide the intuition behind the model here; for technical results, see Appendix A in the Supplementary Information (SI).

bargain by $\Omega = \{s, c\}$, where $s \in (0, 1)$ is the share of the simple good going to party A and $c \in (0, 1)$ is the share of the complex good going to party A.⁶

Second, we assume that each party has a target constituency whose support for the party is a function of the party’s ability to secure a favorable bargain Ω . These supporters are divided into two types. Some supporters are sophisticated, and form accurate evaluations of their party’s performance. These supporters might be politically interested voters or, more likely, organized interests with policy expertise and particularly high stakes in the policy process. Backing for a party (electorally, monetarily, or through organizing efforts) from these supporters increases as the party obtains a “better” deal. That is, support for a party by sophisticated supporters is given by a “support function” that increases (at a diminishing rate) as the party secures a greater share of the simple and complex goods.

The second—and more numerous—class of supporters do not evaluate party performance by assessing the overall coalition bargain. These supporters’ evaluations focus on salient information that is easily assessed using informational shortcuts. We incorporate this feature by assuming that unsophisticated supporters are unaware of the distribution of the complex good. Instead, they make evaluations only on the basis of the simple good. Moreover, we assume that the supporters’ heuristic establishes some “yardstick” of what their party ought to secure. They continue to support parties that do so, but support from this group drops significantly if parties fail to meet these expectations. Supporters may shift their backing to other parties, or become less likely to turn out (or lend other forms of support). In other words, the behavior of unsophisticated supporters does not respond “smoothly” to what the party accomplishes. Instead, support depends critically on whether the coalition bargain meets—in the terminology used by Kahneman and Tversky (1979)—the reference point established by the supporters’ heuristic for the simple good. We denote this reference point by p_i .⁷

⁶ These goods can be genuinely distributive, but can also refer to policy agreements, with the proviso that parties bargain over where in the Pareto set policy will be located.

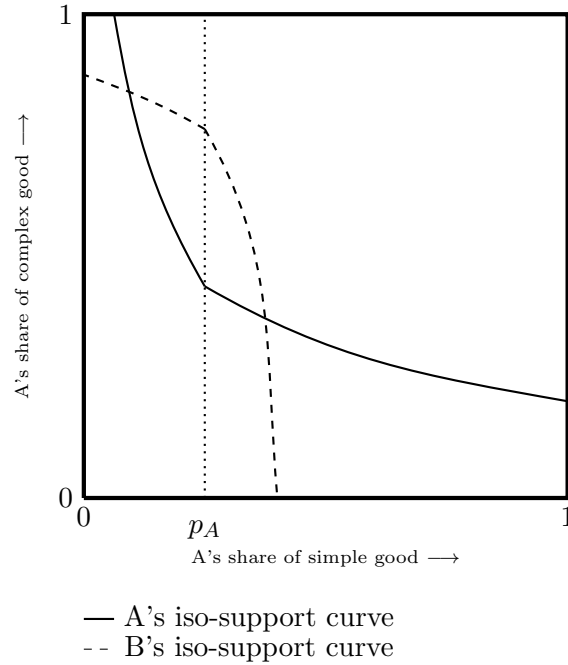
⁷ Informally, this implies that the support function for unsophisticated supporters does not depend on the complex good. For the simple good, it has a steeper slope below the reference point implied by the heuristic than the support function for sophisticated supporters, as well as a more shallow slope above the reference point. Of course, *simple* and *complex* goods are merely analytically convenient ends of a spectrum—many of the goods that parties bargain over, whether cabinet portfolios, budgets, or policy agreements, can be more or less complex and difficult to assess. What is critical for our argument is merely that there is a meaningful distinction in the readiness with which unsophisticated supporters can use their heuristics to assess coalition bargains for different types of goods.

In the context of coalition governance, there are numerous heuristics supporters might rely upon. For example, they may expect that the largest coalition party should receive the Prime Ministership. Or supporters of niche parties may expect their party to secure the most salient ministry associated with the party’s core issue (Bäck, Debus and Dumont 2011). The relative size of coalition partners may also be a natural and easily observed focal point (Schelling 1960) indicating what a party is “entitled to,” especially with respect to goods that can be assessed in relative terms. Of course, other heuristics are possible as well. The key point for our argument is simply that supporters employ “shortcuts” that define expectations over what a party should secure—and that the likelihood that they will support a party drops sharply if it fails to secure its “fair share.”⁸

We assume that parties are motivated by a desire to maximize support—which depends on the backing they receive from sophisticated and unsophisticated supporters (and the relative size of these groups). All else equal, support for a party increases as it is able to secure a greater share of simple and complex goods in coalition negotiations. However, given the fact that the vast number of supporters are unsophisticated, backing for a party is especially sensitive to whether the bargain provides a share of the simple good that satisfies its supporters’ reference-point expectations. Bargains that fail to do so lead to significant losses among the group of less sophisticated supporters, while bargains that provide at least a “fair share” of the simple good retain their support.

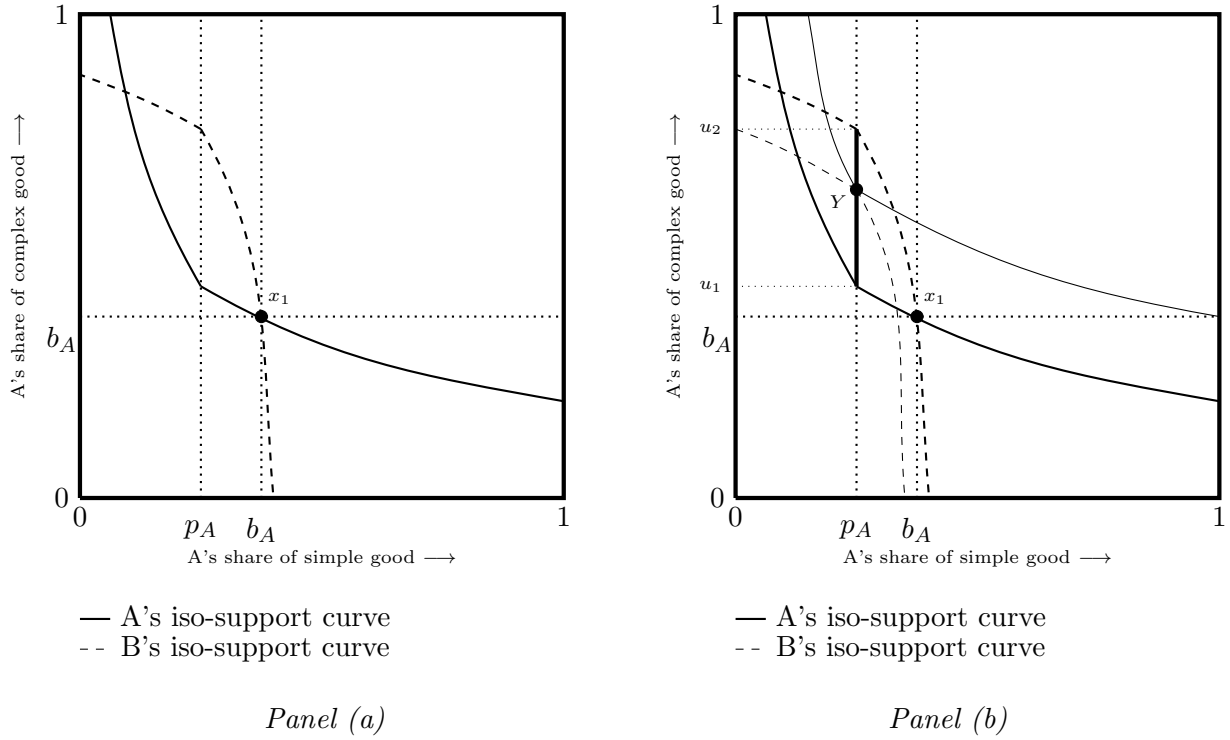
We can explore the consequences for coalition bargaining in the familiar Edgeworth Box framework, displayed in Figure 1. The x -axis indicates the allocation of the simple good between the two parties. The y -axis shows the allocation of the complex good. Thus, at the origin, party A receives no allocation of either good, with party B fully controlling both, while party A’s share of the goods increases as we move right and/or up. Along the x -axis, p_A marks the “fair share” reference point implied by the informational shortcut of less sophisticated supporters (e.g., if supporters employ a proportionality heuristic, then the parties’ relative sizes establish the reference expectations, and p_A corresponds to party A’s seat share). (Since less sophisticated supporters are assumed not to be able to assess the complex good, p_A is omitted from the y -axis.)

⁸ We focus here on heuristics that generate compatible expectations across parties. Relative size—which we focus on below—meets this condition. Exploring the implications of non-compatible heuristics is an obvious avenue for future work.

FIGURE 1. *Coalition Bargaining Space*

Graphically, we can indicate the parties' payoffs from the various allocations of the two goods by “iso-support curves,” which depict combinations of the two goods that result in the *same* level of support for the party. These curves are shown by solid lines for party A, and dashed lines for party B. Because parties are motivated to maximize support, they are indifferent among bundles that lie on the same iso-support curve, and prefer bundles on “higher” iso-support curves to those on lower ones. Sensibly, this means that Party A prefers allocations towards the northeast of the quadrant while party B prefers allocations towards the southwest.

Consider the shape of the iso-support curves. Allocations that provide less than a party's “fair share” of the simple good result in a significant loss of support among less sophisticated supporters. This implies that below the reference point p_i , to maintain the same level of support (i.e., to stay on the same iso-support curve), reductions in the simple good need to be compensated by increases in the allocation of the complex good—i.e., the loss of support by less knowledgeable supporters must be balanced by increased support by those who are sophisticated. Given the relative sizes of the two groups, this requires a substantial increase in the allocation of the complex good. The same is not true above the reference point, where the backing of unsophisticated supporters is largely assured, and additional allocations of the simple good primarily affect support only among the sophisticated. Put

FIGURE 2. *Coalition Bargaining Space*

differently, the parties' iso-support curves have a "kink" at the reference point, and have a steeper slope below the "fair share" point than above it. For party A, this means that the curve has a steeper slope to the left of p_A ; for party B, it means that the curve drops off more sharply to the right of p_A .

We are now in position to consider the implications for the nature of the agreed-upon coalition bargain. Without loss of generality, suppose that the relative bargaining strength of party A (under whatever bargaining protocol the parties employ) is given by b_A . Given this, based on the parties' bargaining positions, we would expect party A to secure a share b_A of both goods. As indicated in Figure 2, Panel (a), this implies a coalition compromise at x_1 . Note that at x_1 , party A's share of the simple good—based on its bargaining power—exceeds the expectations of its unsophisticated supporters (e.g., the party's relative size under a proportionality heuristic), while party B's allocation of the simple good falls below the expectations of its unsophisticated supporters. Put differently, there is some divergence between the bargaining strength of the parties, and the share of the spoils they are expected to secure under the heuristic employed by unsophisticated supporters.

Critically, the premise that parties are concerned to maximize support, which depends on how supporters perceive the bargain that is struck, implies that the allocation at x_1 is inefficient: there are

bargains that *both* parties prefer to x_1 . Because x_1 provides Party B with less of the simple good than its unsophisticated supporters expect (resulting in a significant loss of support), Party B is willing to offer substantial concessions to increase its share of this good. Trading away the complex good results in some losses among the sophisticated group, but these are more than compensated by retaining the support of the (larger) group of unsophisticated supporters. For Party A, such a trade is similarly attractive. As long as it retains a share of the simple good that meets supporter expectations, it can maintain the support of its unsophisticated group while increasing support among its sophisticated group by obtaining a larger share of the complex good. This is highlighted by the fact that moving from x_1 towards the northwest will place *both* parties on higher iso-support curves. As illustrated in Panel (b), such trades only become impossible when the iso-support curves are just tangent—which occurs at allocations, such as Y , that split the simple good exactly at p_A . The bold part of the line at p_A between u_1 and u_2 represents the contract curve of (efficient) bargains that are preferred by both parties to x_1 .

Note that for all of these bargains, Party A (the party in the strong bargaining position) accepts a bargain that “undercompensates” it—relative to what the party could extract on the basis of its bargaining power—with respect to the simple good. In exchange, the party receives a greater amount of the complex good than it could expect on the basis of its bargaining strength. The reverse is true for Party B (which is in a weaker bargaining position). In short, if parties bargain efficiently, they will allocate the simple good in proportion to p_A , and use the complex good as a side payment that reconciles the bargain with the parties’ bargaining power.

Implication 1. *The distribution of simple goods in a coalition bargain will correspond closely to the distribution implied by the heuristic employed by unsophisticated supporters.*

Implication 2. *Parties that are undercompensated relative to their bargaining power in terms of simple goods will be overcompensated relative to their bargaining power in terms of complex goods.*

In addition, the argument implies that as the degree of undercompensation with respect to the simple good increases for a party, the contract curve over the complex good shifts in that party’s favor. To see this, note that shifting b_A to the right while holding p_A fixed causes x_1 to slide up the 45° line.

This induces the contract curve for efficient bargains (whose lower bound is marked by Party A’s iso-support curve through x_1 , and whose upper bound is marked by Party B’s iso-support curve through x_1) to move up. Of course, in the absence of assumptions regarding a specific bargaining process, the theory yields no point prediction about precisely where on the contract curve a bargain will be struck. Nevertheless, a reasonable expectation is that for a wide range of bargaining protocols, as a party is more severely undercompensated with respect to simple good, it will secure an increasingly favorable allocation of the complex goods as the contract curve shifts in its favor. This would be true, for example, for any bargaining protocol that implies bargains that are proportional to a constant, e.g., protocols under which party A receives some proportion $s \in (0, 1)$ of the surplus and party B receives proportion $1 - s$.⁹

Expectation 1. *As the extent of a party’s undercompensation relative to its bargaining power in terms of simple goods increases, the party is compensated with an increasing allocation of complex goods.*

The intuition underlying these results is clear: Party elites anticipate that accepting less than the share of easily observable “spoils” that unsophisticated supporters expect them to secure threatens significant losses, as the party is likely to be perceived as an ineffective negotiator. The same is not true of less observable (complex) payoffs, which these supporters are much less likely to be aware of. Consequently, party elites face incentives to exchange side payments of less observable goods to obtain a share of simple goods that meets supporter expectations. Importantly, this conclusion depends only on the nature of party preferences, and is independent of the specific bargaining protocol that determines the overall distribution of spoils. In other words, this result derives from the nature of parties’ preferences—induced by a concern for how supporters evaluate a coalition bargain—not the structure of the bargaining situation.

Of course, our simple model presents the logic in stark form—namely, a precise correspondence of simple goods to the distribution implied by the supporter heuristic. Embedding this logic in real-world bargaining will attenuate this precision. Goods may be “lumpy” in ways that prevent precise

⁹ To see this, denote the initial contract curve by $[l_1, u_1]$ and the new contract curve by $[l_2, u_2]$, with $l_2 > l_1$ and $u_2 > u_1$. Then the initial bargain is given by $Y_1 = l_1 + s(u_1 - l_1)$ and the new bargain is given by $Y_2 = l_2 + s(u_2 - l_2)$. Then $Y_2 > Y_1$ if $s(u_2 - u_1) + (1 - s)(l_2 - l_1) > 0$, which must be true.

division. Similarly, getting “close” to the allocation expected by unsophisticated supporters may be sufficient. Moreover, many aspects of coalition bargaining are not simple or complex, but fall along a continuum. Finally, supporters are not sophisticated or unsophisticated, but vary in political awareness. Nevertheless, the underlying logic applies: Party elites have strong incentives to work towards bargains that approximate supporter expectations for goods that supporters can more easily assess, using those aspects that they are less likely to be aware of as side payments. As we discuss in the conclusion, this fact has significant implications for the quality of representation—most importantly, for the ability of citizens to hold their parties accountable.

AN APPLICATION TO PORTFOLIO ALLOCATION

There are several possible ways to evaluate our argument empirically. One is to examine our underlying assumptions about supporter behavior: Do supporters (especially voters) employ heuristics that establish “yardsticks” by which party performance is judged, and does the use of such heuristics vary with political sophistication? Does party failure to meet the expectations implied by the heuristics result in a loss of support? Another way is to examine our assumptions about elite behavior: Is there evidence that party leaders involved in coalition negotiations are concerned about how supporters will perceive bargaining outcomes, and do these concerns shape what leaders are willing to agree to? Finally, an empirical evaluation could investigate the implications of these behavioral assumptions for the actual bargains that are struck by party elites: Do these bargains reflect the tradeoffs across more and less observable goods that are predicted by the argument?

Given the limited scope of the current study, we focus on this last avenue of investigation here, leaving the first two for future work. We recognize, of course, that examining these behavioral implications provides only partial evidence. At the same time, it is important to stress that (as we show below) these implications discriminate between our argument and the most prominent alternative explanations of bargaining in the coalition literature, thus increasing confidence that the results we report offer initial support for the theory.

We focus on a salient aspect of coalition bargaining: the allocation of ministerial portfolios. Coalition research, almost since its inception, has confronted a tension between a well-known empirical regularity

and the predictions of most theoretical models of coalition bargaining. The regularity is referred to as “Gamson’s Law”—the robust finding, confirming the intuition of Gamson (1961), that coalition members tend to receive a share of cabinet ministries that is roughly proportional to the share of legislative seats they contribute to the government (Browne and Franklin 1973; Warwick and Druckman 2006). From a theoretical perspective, Gamson’s Law poses a puzzle because the vast majority of bargaining models do not predict it. Instead, most models lead to the expectation that the distribution of spoils should correspond to the *bargaining power* of the parties, which often does not coincide with their seat shares (see, e.g., Morelli 1999; Montero and Vidal-Puga 2011).¹⁰

Our argument provides an avenue for reconciling this tension. The relative size of parties provides a readily available yardstick that supporters may use to assess their party’s performance for outcomes that are easily measured against it. The numerical distribution of cabinet portfolios is an obvious example. It is salient, widely reported, easily understood, and readily assessed using a “relative party size” heuristic. If less sophisticated supporters make use of this heuristic, and party elites are concerned about their perceptions, one would expect these elites to construct bargains that ensure that the numerical allocation of portfolios follows relative size closely.¹¹

Of course, the Gamsonian regularity is also consistent with an alternative explanation: a simple commitment by party elites to proportionality; i.e., it is possible that party leaders do not worry about supporter perceptions, but simply act in accordance with a fairness norm that focuses on party size rather than bargaining power. Indeed, Verzichelli (2008, 239) argues that such a proportionality norm underlies coalition bargaining since it “means that each coalition party is rewarded in proportion to its contribution to the parliamentary strength of the government, which corresponds to many people’s

¹⁰ Bassi (2013) presents a model that predicts proportionality of portfolios to seats, providing a theoretical foundation for Gamson’s Law. However, this result requires a specific bargaining protocol, possibly limiting its applicability to “real world” bargaining, which, as Laver, de Marchi and Mutlu (2011) point out, is relatively unstructured and free-flowing.

¹¹ Most supporters are unlikely to retain detailed information regarding the initial portfolio distribution, and to make active use of such information in later voting decisions. Rather, as the “online tally” model of supporter decision-making suggests (Lodge, Steenbergen and Brau 1995), exposure to such information affects the *summary evaluations* of a party. Such summary evaluations have been found to be long-lasting, even if the particular information underlying them is quickly forgotten (Lodge, McGraw and Stroh 1989).

ideas of fairness.” Another alternative explanation holds that party leaders might simply use proportionality as a focal solution in structuring the coalition bargain irrespective of voter perceptions (Bäck, Meier and Persson 2009).

The key aspect that distinguishes our argument from these alternatives is that it implies that the distribution of goods is *conditional* on their observability. If the proportional allocation of spoils is a consequence of an intrinsic commitment to “fair” distribution in accordance with relative size or its focal qualities, then the observability of goods by less sophisticated supporters is irrelevant for their distribution. In contrast, under our argument, it is supporter perceptions that are critical. Easily observable goods are allocated proportionally because those party leaders whose bargaining power is below their seat share—wary of losing the backing of unsophisticated supporters by failing to secure a proportional share of an easily observable good—are *trading away* less observable goods to parties in a stronger bargaining position. This leads directly to an implication that discriminates between our argument and the alternative accounts. Specifically, our argument implies that—by accepting a Gamsonian distribution of portfolios—parties whose bargaining power exceeds their seat share agree to be “undercompensated” (relative to their bargaining power) in terms of the easily observable share of portfolios, but are remunerated by side payments in the form of complex goods. Moreover, these side payments increase the more significant the party’s undercompensation. Neither of the alternative accounts leads to this expectation.

There are numerous complex goods that can potentially serve as side payments. Most obviously, parties might demand policy concessions (which, especially for more complex issues, are less likely to be observable for less sophisticated supporters). One attractive and important “complex” good in this context concerns the policy implications that flow from the particular assignment of specific ministries to specific parties. Ministries, headed by a particular cabinet member, play a key role in translating abstract policy agreements into concrete proposals. Legislative initiatives are drafted under the direction of a minister, presented by that minister to the cabinet, and then introduced in parliament. While the precise degree of ministerial influence, and the conditions under which ministers can exercise discretion, continues to be a rich subject of academic debate (Laver and Shepsle 1996; Martin and Vanberg 2011), there exists a general consensus that ministers can exercise policy influence

unless coalition partners expend resources to monitor and constrain ministerial discretion. This implies that a key consideration for party leaders in evaluating a portfolio allocation is the extent to which it grants power to ministers who pose a potential threat.

The extent of this “policy risk” inherent in a portfolio distribution depends on two features. First, given the issues under the jurisdiction of a given ministry, how significant are the policy differences between a coalition partner and the party controlling the portfolio? If the parties have few policy differences, ministerial discretion poses little danger. Second, how important are the issues under the jurisdiction of the ministry to the party and its constituents? If the issues relevant for a particular ministry are not salient for the party, then ministerial discretion is of less concern, even if there are policy differences between the parties. In short, the “policy risk payoffs” of a portfolio allocation for a coalition party depend on the extent to which it confers influence to partners with significantly different preferences on issues the party cares about.

Assessing policy risk is a challenging task for unsophisticated supporters. Doing so requires an understanding of the jurisdictions and responsibilities of ministries, as well as of the specific policy goals of the parties, and the extent to which these goals are in conflict. In other words, the policy risk aspect of the portfolio allocation is a complex good—and thus provides one way to deliver side payments to parties that are undercompensated with respect to the numerical allocation of portfolios. Moreover, this side payment may be particularly attractive because it is “paid” directly alongside the Gamsonian distribution of portfolios, reducing concerns about enforcement of the compensation agreement. Thus, the theoretical implications outlined earlier lead to the following hypotheses:

Hypothesis 1. *The numerical allocation of cabinet portfolios should be highly proportional to parties’ relative size.*

Hypothesis 2. *As a party is more severely undercompensated (relative to its bargaining power) in the number of portfolios it receives, it will be compensated with a more favorable policy risk payoff from the portfolio allocation.*

ANALYSIS AND FINDINGS

To test these hypotheses, we examine original data on the allocation of 5,910 ministerial posts in 308 coalition bargaining situations in 16 European democracies. In SI Appendix B, we provide more detail on the set of governments included, and we also list the archival sources used to create the data on ministry assignments. Across the 308 governments, our sample consists of 910 observations on government parties (the units of analysis).

Hypothesis 1 predicts that parties will receive a share of ministries that corresponds closely to their seat contributions to the government. Numerous studies have examined this relationship, typically via a simple linear regression of the portfolio shares of parties on their relative seat shares. We conduct a similar analysis for our sample, with results shown in Table 1.¹²

TABLE 1. *Effect of Relative Legislative Size on Share of Ministerial Posts*

Explanatory Variables	Estimates
Seat Contribution	0.790 *** (0.011)
Intercept	0.068 *** (0.004)

Displayed are OLS coefficient estimates, with standard errors in parentheses. Dependent variable: Numerical share of cabinet ministries. N: 602 coalition parties (after 308 parties, one per bargaining situation, are randomly dropped from the analysis—see footnote 12). $R^2 = 0.89$. Significance levels : * : 10% ** : 5% *** : 1%.

The parameter estimates are close to those produced by previous studies (see, e.g., Warwick and Druckman 2006, 648). The share of seats a party contributes to the government has a strong positive effect on the share of ministries it receives. Notably, the match is not *perfectly* proportional. The coefficient on *Seat Contribution* is statistically below unity and the intercept is statistically above zero, meaning that smaller parties receive more than their “fair share” of portfolios, a finding that is in line with previous results (Browne and Franklin 1973). This tendency can be seen in Figure 3, which

¹² Because the dependent variable is compositional, the error terms are correlated across parties. Following the approach of Carroll and Cox (2007), we randomly drop one party from each bargaining situation. We also estimate multilevel models on this subsample (shown in SI Appendix C) to account for the possibility that the mean level of proportionality, and the effect of seat share on portfolio share, may differ by bargaining situation. Likelihood ratio tests reveal no significant improvement over the model shown in Table 1.

plots the fitted prediction from Table 1 (with 95% confidence bounds) against a 45° line. Most small parties lie above the “perfect proportionality” line, while most large parties lie below it. Nonetheless, on average, for every 5% increase in a party’s seat share, its expected share of ministries increases by approximately 4%.




FIGURE 3. *Effect of Relative Legislative Size on Share of Ministerial Posts*

Of course, as noted earlier, such proportionality is also consistent with some models of coalition bargaining (e.g., Bassi 2013), or a general commitment by party elites to proportionality as a “fair” or “focal” bargaining solution. The distinguishing feature of our argument is that the proportional allocation of portfolios is induced by party elites’ concerns for the perceptions of their supporters, leading us to the expectation (captured in Hypothesis 2) that elites will reconcile the proportional allocation of the number of portfolios with the underlying bargaining strength of the parties by using less observable aspects of the bargain as side payments.

As we discussed, there are numerous complex goods that can serve as side payments. We focus on a complex good that derives from the initial allocation of ministries. Since cabinet ministers are

responsible for supervising the formulation and implementation of policies in their domains, gaining control of a ministry can give a coalition party a significant advantage in shaping policy outcomes in its purview. By the same token, a party can incur significant costs from handing over control of a ministry to one of its partners, especially if the ministry deals with salient, divisive issues. Again, we refer to these costs as the “policy risk payoffs” of the allocation. To measure them, we must perform two tasks. First, we have to measure the *positions* of parties on a set of policy dimensions, and assess the *importance* they place on those dimensions. Second, we must associate the dimensions with specific ministerial posts.

To measure party positions and issue salience, we must first define the nature of the policy space. This is an inherently subjective task. As Benoit and Laver (2006, 69) point out, “there is no such thing as the ‘true’ underlying dimensionality of any policy space,” which means ultimately that, in terms of both the *number* and the *substance* of policy dimensions, “[the] type of space we want to measure depends upon the type of model we want to analyze.” In defining the policy space for our purposes, we attempt to satisfy two criteria. First, we want the number of dimensions to be sufficiently high as to adequately capture both the major areas in which ministers make decisions and the areas in which their parties compete. Second, in terms of substance, we want our dimensions to correspond roughly to the set of dimensions other researchers have highlighted as important in party competition, and to produce similar patterns of party contestation.

Based on these criteria, we define a policy space of *six* conceptually distinct dimensions: economic regulation and state services, traditional morality, environmental protection, internationalism and European integration, social cohesion and national identity, and conflict and military power. Using the MARPOR quasi-sentence coding categories, and the scaling approach developed by Lowe et al. (2011), we calculate issue saliency scores and policy positions for parties on these dimensions, described in SI Appendix D. We also present there a confirmatory factor analysis to demonstrate that the chosen MARPOR categories load on the predefined policy dimensions in reasonable ways, and we investigate the face validity of our measures by plotting their averages for party families in postwar Western Europe. While we do not expect our definition of the policy space to be the “last word” on the matter,

we are encouraged by these results, and believe that future efforts in this area should entail similar construct and face validity checks.

Matching these policy dimensions to specific ministries obviously requires knowledge of ministerial competences. We have assessed these based on information from a variety of sources, including historical descriptions provided by the ministries, formal legislative documents detailing the policy competences of each department, and the types of government bills and directives introduced by ministers. Most of the time, multiple ministers share responsibility within a given policy area; furthermore, any individual minister typically deals with issues in several policy areas. Accordingly, in our coding scheme, ministries that have significant policymaking responsibility in different policy areas are assigned to multiple policy dimensions.¹³

To test Hypothesis 2, the key concept we need to capture with the dependent variable is the *level of policy risk* a government party is exposed to as a result of the agreed-upon distribution of ministries. Is it the case, as we have argued, that a party that “gives away” a number of ministries it “deserves” (based on its bargaining strength) to achieve numerical proportionality for all coalition parties does so in such a way as to lower its policy risk? That is, does an “undercompensated” party generally avoid ceding ministries to ideologically distant partners on dimensions the party cares a lot about? On the other side of the bargain, is an “overcompensated” party generally willing to tolerate greater policy risk to attain numerical proportionality? That is, in exchange for receiving a greater number of ministries than warranted by its bargaining strength, is such a party inclined to allow its undercompensated partner to control ministries on important dimensions on which the preferences of the two parties significantly diverge?

As these questions suggest, and as noted earlier, a party’s policy risk payoff from an allocation depends on two factors: (1) *how distant*, on the relevant issues, the party is ideologically from the coalition partners that gain control of the ministries, and (2) *how important* those issues are to the party. As these factors increase, the negative policy consequences experienced by the party (and its supporters) will become more severe, and/or the party has to devote greater resources towards

¹³ In SI Appendix E, we list the major ministries in our sample of countries, the policy dimensions assigned to them, and the typical issues on these dimensions for which ministers are responsible.

constraining its partner’s ministers in an effort to avoid such consequences. We incorporate these factors into our dependent variable, which is constructed in the following steps:

- (1) we calculate a *dimension-specific policy risk payoff* for a coalition party c as the average saliency-weighted distance between c and all the ministries (including those controlled by c itself) that are responsible for that dimension (the ministries M_d),

$$\frac{1}{M_d} \sum_{m \in M_d} y_{cd} |i_{cd} - i_{md}|,$$

where y_{cd} is the relative saliency that party c attaches to dimension d , i_{cd} is c ’s policy position on the dimension, and i_{md} is the policy position of minister m on the dimension;¹⁴

- (2) we sum the dimension-specific payoffs across the set of (six) policy dimensions D to create a *multidimensional policy risk payoff*,

$$\sum_{d \in D} \frac{1}{M_d} \sum_{m \in M_d} y_{cd} |i_{cd} - i_{md}|;$$

- (3) finally, since we are interested in measuring the policy risk for a party *given* the number of ministries ceded to other parties (i.e., we are interested in the level of policy risk a party has accepted given the numerical constraints it has agreed to), we divide by the number of portfolios controlled by the party’s partners, p_{oc} , thus obtaining the average multidimensional policy risk payoff per partner-controlled ministry:

$$\frac{1}{p_{oc}} \sum_{d \in D} \frac{1}{M_d} \sum_{m \in M_d} y_{cd} |i_{cd} - i_{md}|.$$

In short, the dependent variable, which we refer to as a party’s *Portfolio-Adjusted Policy Risk Payoff*, taps into a feature of portfolio allocation that parties presumably value highly—the distribution of policymaking authority—but whose details most supporters will be largely unaware of because understanding it requires knowledge that goes well beyond knowing *how many* ministries a party controls. Even knowing *which ministries* a party receives would be insufficient. To evaluate the policy risks their party faces from a particular allocation, supporters need to know which ministries are responsible for which policy areas, how the ministries not received by their party are distributed

¹⁴ As noted in SI Appendix D, we scale the relative saliency score for a party such that a dimension of above-average importance to the party receives a score greater than 1.0, and a dimension of below-average importance receives a score less than 1.0. Each minister is assumed to have the same position as his or her party.

among its coalition partners, and the policy preferences of each of these parties in the areas under their control.

As stated in Hypothesis 2, we expect that parties that are, relative to their bargaining power, *undercompensated* in the numerical share of portfolios receive side payments in the form of a more favorable policy risk payoff. Thus, we require a measure of bargaining power. There are, of course, multiple approaches to assessing bargaining power. One alternative derives from the literature on power indices. These measures focus on the size of parties and assess each party’s “pivotality” across the set of winning coalitions. The downside of this approach is that it pays no attention to the policy preferences of actors, or other factors that potentially strengthen a party’s position in coalition negotiations. This may be sensible in a purely distributive setting, but is problematic in our context since coalition negotiations clearly depend on a host of factors beyond size, including the ideological positions of the potential coalition partners. Therefore, we take a different tack. An important source of leverage for parties in the process of government formation is their ability to join an alternative government to the one they are currently negotiating: Parties that have credible “exit options” from the current negotiations are more likely to be influential. Thus, an alternative approach to assessing bargaining power is to focus on parties’ ability to be part of (alternative) governments.¹⁵

There are multiple empirical models of coalition formation that would allow us to estimate the probability that a party can get into an alternative government. We rely on a recent model by Kayser, Orlowski and Rehmert (N.d.)—hereafter, KOR—that builds on work by Martin and Stevenson (2001, 2010). Like these earlier studies, KOR use a conditional logit model to calculate the probability that a particular government is chosen (out of the set of all potential governments), where the probability in their model is a function of 9 size-related, policy-related, and contextual factors. Some notable advantages of the KOR specification (versus the Martin and Stevenson specifications) are that it is

¹⁵ Of course, just as for the power index approach, there are limitations to using a party’s probability of being able to join an alternative coalition as a measure of bargaining power; for example, a party’s pivotality—which is central to power indices—may not be reflected in its “exit probability” if the party is a surplus member of many coalitions. Given these limitations, we also conduct our empirical analysis using a prominent power index—the normalized Banzhaf index—as an alternative measure of bargaining power. The results are very similar, as shown in SI Appendix J.

more parsimonious, excludes certain variables that are difficult to measure reliably, and has a better predictive performance.¹⁶

We construct our measure of bargaining power in the following steps:

- (1) using the KOR specification, we predict the *probability of formation* for each potential government in each of our 308 bargaining situations, including the coalition that was actually chosen;
- (2) for each government party in the chosen coalition, we calculate the party’s probability of being included in an *alternative* government by summing the formation probabilities of all *non-chosen* potential governments in which that party is a member (we refer to this quantity as a government party’s *exit probability*);
- (3) we *normalize* each government party’s exit probability by dividing it by the sum of exit probabilities across all government parties.

Thus, our measure of party bargaining power is simply its *relative exit probability*.¹⁷ Using this measure, we create our central covariate, *Undercompensation in Portfolio Share*, by *subtracting* the portfolio share received by a party from its bargaining power.

Again, we expect that any party that would receive—on the basis of its bargaining power—a share of portfolios that is less than its perceived “fair share” (i.e., its relative seat contribution) would be willing to trade off complex goods to its partners to achieve that fair-share allocation. What would such a tradeoff look like in practice? Before we move to a full analysis, it is useful to consider a single episode of coalition bargaining. We focus here on the last Dutch government in our sample: the second cabinet led by Mark Rutte, which lasted from 2012 to 2017, and consisted of two parties, the VVD and the PvdA.

These parties were roughly equal in size. Given 14 ministerial posts in the cabinet, each party’s “fair share” (as implied by the relative size heuristic) was 7 posts. Each party received exactly that number. However, while the parties were roughly equal in size, they were *not* equal in bargaining power: Given the composition of the legislature (which tilted towards parties of the right), the probability of VVD

¹⁶ In SI Appendix F, we present the results from our conditional logit model of government formation, using an updated version of the Martin and Stevenson (2010) coalition data set. We also describe the model variables in more detail.

¹⁷ In SI Appendix G, we illustrate the creation of this measure for a recent German government.

TABLE 2. *Portfolio Allocation in Rutte II Government*

Second Cabinet of Mark Rutte (2012–2017)					
VVD			PvdA		
	Prime Minister			Deputy Prime Minister	
	Economic Affairs			Finance	
	Security and Justice			Foreign Affairs	
	Defense			Education and Culture	
	Infrastructure and Environment			Social Affairs	
	Health and Welfare			Interior and Kingdom Relations	
	Housing and Civil Service			Foreign Development Cooperation	
	Seat Share	Implied Portfolios	Actual Portfolios	Bargaining Power	Implied Portfolios
VVD	0.52	7	7	0.64	9
PvdA	0.48	7	7	0.36	5

becoming part of an alternative coalition was almost twice as high as the exit probability for the PvdA. Thus, the Gamsonian allocation of portfolios left the PvdA, relative to its bargaining power, *overcompensated* with respect to the number of portfolios, while the VVD was *undercompensated*. Our argument implies that this numerical “bonus” received by the PvdA came at a price. Namely, since the most important issue for both parties in 2012 was how to improve the economy, a matter over which they were quite divided, we would expect the VVD to have taken most of the ministries with jurisdiction on economic matters, thereby lowering the risk that the PvdA could move policy too far away from the VVD’s preferred position.¹⁸

With two parties and 14 ministries, there are 3,432 possible ways to allocate portfolios such that each party receives seven ministries. In Figure 4, we display the policy risk payoffs for each of these potential allocations for both parties. One important point to note is that the relationship between payoffs for the two parties is negative, meaning that switching ministries would typically (though not always) produce a loss for one party and a gain for the other. A more crucial point is that the actual

¹⁸ The saliency-weighted divisions between the parties on the economic regulation and state services dimension was almost 8 times their average level on the other dimensions. The two parties disagreed over several important issues, such as how quickly and by how much the budget should be cut, and whether there should be an increase in health care premiums.

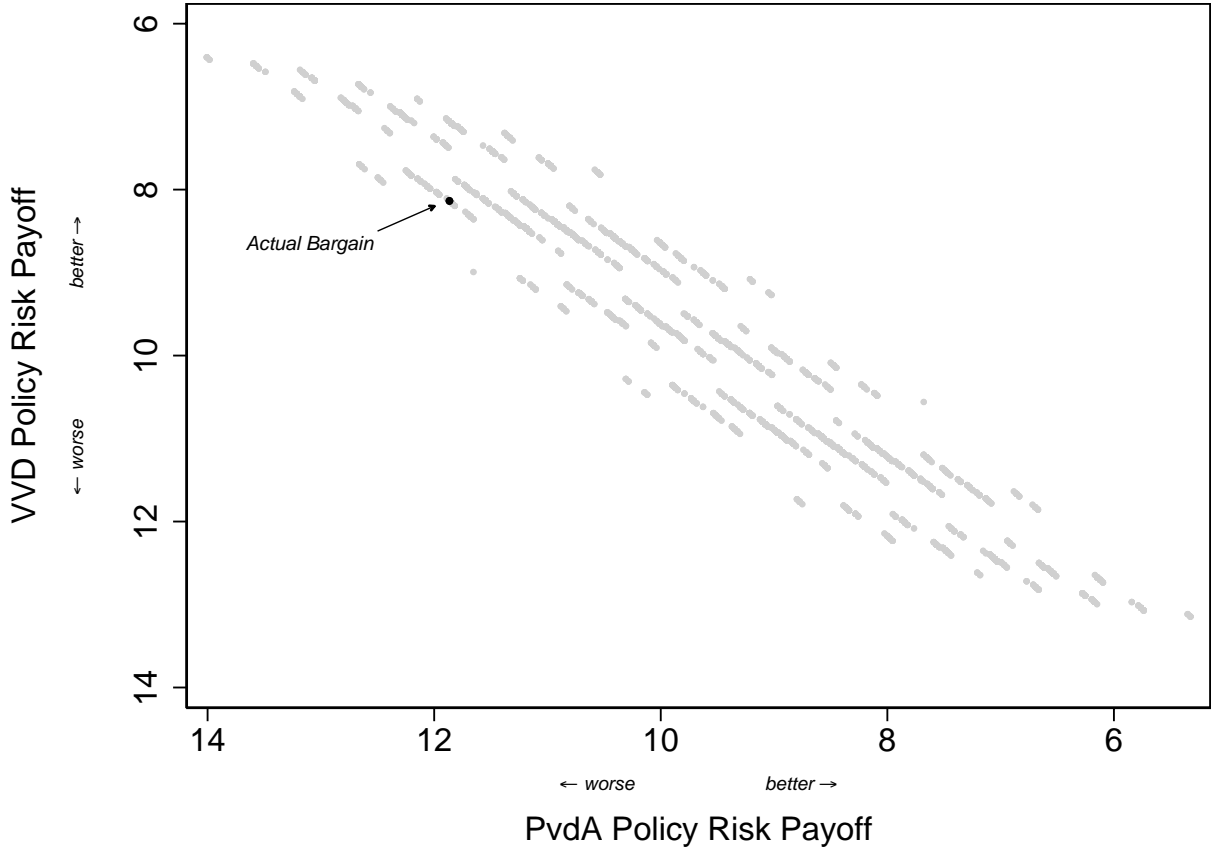


FIGURE 4. *Portfolio-Adjusted Policy Risk Payoffs Across Alternative Seat-Proportional Allocations in the Rutte II Government*

bargain that was struck—in line with the expectations of our theory—was much more favorable for the VVD than for the PvdA. Specifically, the VVD’s policy risk payoff was in the top 20% of its possible payoffs. Meanwhile, the payoff for the PvdA was in the bottom 15% of its possible payoffs: i.e., there were approximately 3,000 alternative ways to distribute the ministries (holding fixed the number of ministries received by each party at seven) that would have been more favorable for the PvdA. Only three of its ministries dealt directly with economic policymaking (Finance, Education and Culture, and Social Affairs), compared to five possessed by the VVD (Prime Minister, Economic Affairs, Infrastructure and Environment, Health and Welfare, and Housing).¹⁹

Is this case emblematic of a more general pattern? Before turning to our analysis, we need to consider a potential confounder, namely the *ideological extremity* of a party in relation to its coalition

¹⁹ Unlike the Prime Minister, the Deputy Prime Minister in the Netherlands has no formal policymaking role, and so this post is not assigned to any policy dimension.

partners. Naturally, a party that finds itself in a government in which it diverges widely from each of its partners on multiple issues should expect to face a higher baseline level of policy risk. If the party’s position vis-à-vis its partners also affects its degree of undercompensation (through its impact on either or both of the covariate’s components, bargaining power or portfolio share), then not taking this factor into account could introduce significant bias into the analysis.

To confront this problem, we include a party-specific distance variable in our model, which we term *Distance from the Coalition Center of Gravity*, and define as:

$$\sum_{d \in D} y_{cd} \left| i_{cd} - \sum_{c \in C} s_c i_{cd} \right|.$$

The final term represents the coalition’s ideological “center of gravity” on an issue dimension d , defined as the weighted-average position of all parties in coalition C on that dimension, with i_{cd} representing the ideal point of each party and the weight s_c corresponding to the proportion of seats each party contributes to the total controlled by the coalition. On each dimension, we measure the absolute distance between party c ’s policy position and the coalition center of gravity, and then weight that distance by the relative importance of the dimension to the party, y_{cd} . Finally, we sum these dimension-specific distances across our six policy dimensions to create the final party-specific measure.

In Table 3, we present results from a multilevel random coefficients model. This type of model has several appealing features. First, it accounts for the fact that our party-level observations are nested (clustered) within bargaining situations. Policy risks may vary across bargaining situations for reasons not accounted for by the party-level covariates, inducing positively correlated errors between government parties. Failing to account for such correlations could yield estimates of the standard errors that are too small.²⁰ Second, the model allows us to deal with possible heterogeneity in the effects of our covariates by providing estimates of both the overall (or average/fixed) intercept and slope coefficients of interest (shown in the top part of the table) as well as the extent to which the intercepts and slope coefficients at the level of the cluster deviate from the average estimates. Using the estimated variance components of the model (shown in the bottom part of the table), we can calculate best linear unbiased predictions of the random deviations and then use those to generate predicted

²⁰ We discuss below the issue of possible negative correlation in risks.

TABLE 3. *Effect of Portfolio Undercompensation on Policy Risk*

Explanatory Variables	Estimates
Undercompensation in Portfolio Share (<i>ups</i>)	-5.449 *** (0.717)
Distance from Coalition Center of Gravity (<i>dccg</i>)	0.698 *** (0.035)
Intercept (α)	4.910 *** (0.255)
$\sigma^2(ups)$	61.638 (16.415)
$\sigma^2(dccg)$	0.149 (0.025)
$\sigma^2(\alpha)$	14.027 (1.781)
$\sigma(ups, dccg)$	0.966 (0.463)
$\sigma(ups, \alpha)$	-13.975 (4.124)
$\sigma(dccg, \alpha)$	-1.253 (0.186)
$\sigma^2(residual)$	2.797 (0.258)

Displayed are maximum likelihood coefficient estimates and standard errors from a multilevel random coefficients model (with no constraints imposed on the covariance matrix). Dependent variable: Portfolio-adjusted policy risk payoffs. N: 910 coalition parties (nested in 308 bargaining situations). A likelihood ratio test against a linear model with fixed coefficients leads us to reject the null hypothesis of no difference between the models ($p < 0.001$). Significance levels for level-one parameters: * : 10% ** : 5% *** : 1%.

random intercepts and random slope coefficients for individual bargaining situations.²¹ Finally, the model allows us to assess how much unexplained variation in policy risk payoffs is due to differences between parties within bargaining situations versus differences between bargaining situations.

Based on the results, we can conclude confidently (via a likelihood ratio test) that the multilevel random coefficients specification leads to a significant improvement in fit over a one-level model with

²¹ We allow parameter estimates to vary across bargaining situations because, in most cases, it is preferable to specify a general model and then “test down” (e.g., via *t*-tests or model goodness-of-fit tests) to a more restricted version. Similarly, rather than impose the typical multilevel modeling constraint that random effects are independent, we assume an *unstructured* covariance matrix, letting each variance and covariance component be uniquely estimated.

fixed coefficients and a fixed intercept ($p < 0.001$).²² We can also conclude that most of the unexplained variability in policy risk can be attributed to random (or unmeasured) differences between bargaining situations, rather than differences between parties. Specifically, the level-two variance partitioning coefficient (VPC) is approximately 0.65, meaning that only 35% of the total residual variance is attributable to unmeasured differences between parties.²³




FIGURE 5. *Residuals and Predictive Fit for Policy Risk Model*

But how large is this residual variance, i.e., how well does our model fit the data? In Figure 5, left panel, we present a quantile-normal plot in which the standardized residuals for our 910 coalition parties are compared against the values one would expect for a standard normal distribution (as represented by the solid 45° line). Most residuals stick fairly close to the line, indicating no serious departures from normality, and there are very few outliers (approximately 2% of cases in each tail).

²² Moreover, our model, which *estimates* the covariances between random effects, significantly improves fit over an alternative specification that constrains the covariances to be zero ($p < 0.001$).

²³ The VPC at a given level of the model is simply the variance at that level divided by the total residual variance. In a multilevel model with random slopes, the level-two variance is a quadratic function of the covariates, so in our case: $\sigma^2(\alpha) + (ups)^2 \sigma^2(ups) + (dccg)^2 \sigma^2(dccg) + (ups) 2\sigma(ups, \alpha) + (dccg) 2\sigma(dccg, \alpha) + (ups)(dccg) 2\sigma(dccg, ups)$. We set the two covariates equal to their means for our VPC calculation.

In the right panel, we plot our predicted policy risk payoffs against the actual policy risk payoffs for the 910 parties. In this panel, the dotted line is a 45° line, and the solid line is the linear fit from a regression of predicted policy risk payoffs on actual policy risk payoffs. Any deviation between the two lines suggests a systematic pattern in actual policy risk that is not being adequately explained by the model. Fortunately, the figure shows only a small difference between the lines: The solid line has an intercept of approximately 1.31 and a slope of 0.84, meaning that the model slightly over-predicts policy risk payoffs for parties with low risk and slightly under-predicts payoffs for parties with high risk. Nonetheless, the model appears to fit the data quite well overall, perhaps surprisingly so, since it includes only two covariates and their associated random effects.²⁴

Of course, the most important results from Table 3 concern the estimated effects of these covariates. Beginning with our control variable, we observe that, as expected, the more distant a government party is from the coalition’s ideological center of gravity, the higher its level of policy risk. That is, as a party becomes relatively more extreme, it is more likely to see the ministries not under its control fall into the hands of ideologically hostile partners.

More importantly, the results provide strong support for our central argument. Coalition parties that are undercompensated in their numerical share of portfolios (relative to their bargaining power) tend to receive significantly more favorable policy risk payoffs than parties that are overcompensated. In other words, the set of ministries that are *not* given to undercompensated parties, but rather to their partners, (a) tend to be those with jurisdiction on policy issues that are relatively less important to the undercompensated parties, and/or (b) tend to be assigned to partners with relatively similar policy positions. This supports our central claim that policy risk payoffs appear to be used as “side payments” to parties that are numerically undercompensated.

²⁴ Although multilevel models can effectively solve the problem of residual positive correlation between observations in a cluster, they cannot solve the problem of residual negative correlation. Rather, appropriate level-one covariates must be included in the model in an effort to reduce residual negative correlation to a negligible amount. Regarding our two covariates, their values within clusters also happen to be negatively correlated, and this fact appears to account for much of the negative correlation in policy risk. That is, as we show in Figure 5 and with our VPC calculation, the amount of residual variance is quite low overall and is mostly due to unmodeled differences between bargaining situations rather than between parties. Nonetheless, it would be heroic to assume that there is so little remaining residual negative correlation that it can cause no inferential problems. Fortunately, unmodeled negative correlation presents the opposite problem as unmodeled positive correlation—it results in an *overestimation* of standard errors, and thus an understatement of the parameters’ statistical significance (Norton et al. 1996, 920–24).

Of course, a critical question concerns the substantive significance of these side payments. To illustrate this, we calculate the predicted policy risk payoffs for two types of coalition parties in two different bargaining situations. In both situations, the parties are assumed to control approximately half the portfolios in their respective governments. But for one (the undercompensated party), this numerical payoff is considerably less than what is implied by its bargaining power, while for the other (the overcompensated party), the payoff is considerably more than its bargaining power.²⁵ For each type of party, we predicted its policy risk payoff using the overall intercept and slope estimates from Table 3 (holding distance from the coalition center of gravity at its sample mean), and then generated a distribution of 1,000 simulated predictions (drawn from a normal distribution with a mean equal to the predicted effect and a standard deviation equal to the standard error of that effect).

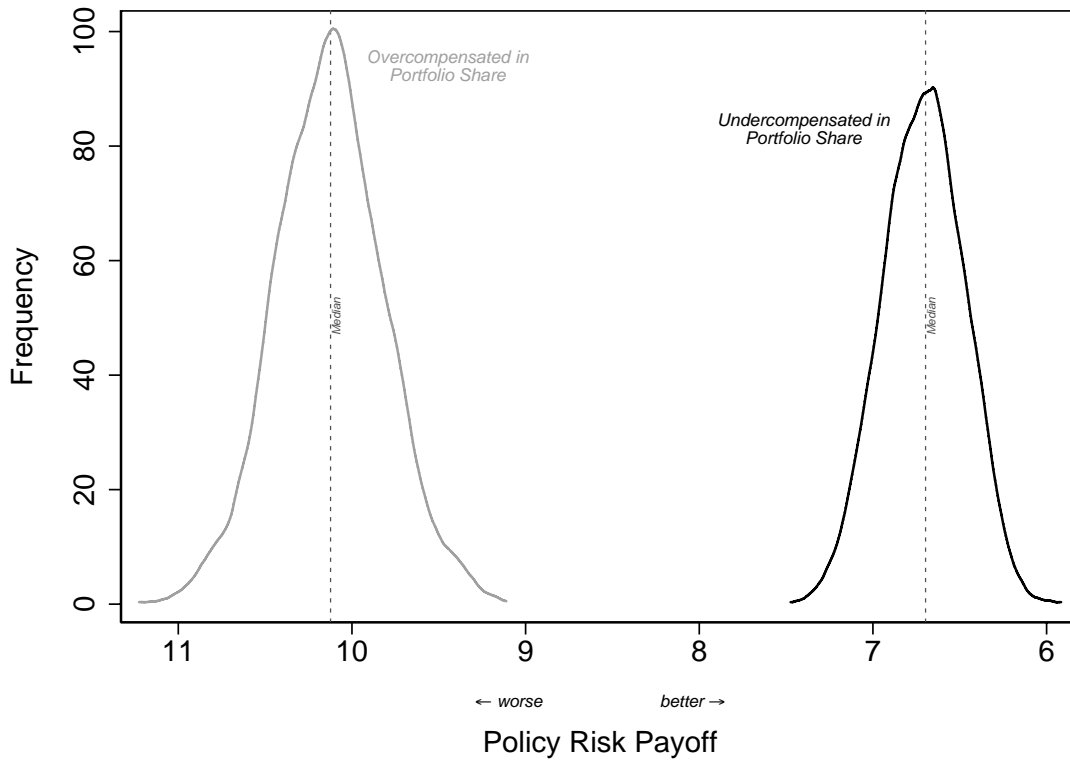


FIGURE 6. *Average Effect of Portfolio Undercompensation on Policy Risk*

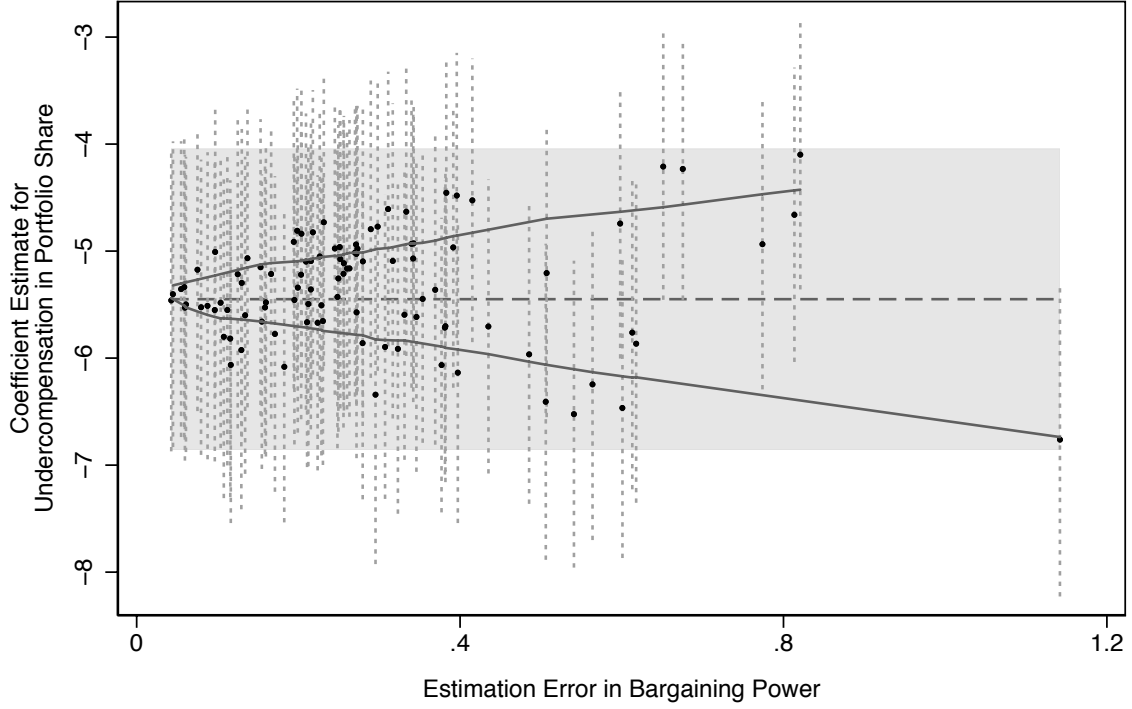
²⁵ To determine the levels of under/overcompensation, we used summary statistics from our sample of bargaining situations. Specifically, we examined the range of undercompensation for parties that received around half the portfolios in the cabinet ($\pm 10\%$) and then assigned the undercompensated party the 90th percentile score from this subsample ($+0.31$), and the overcompensated party the 10th percentile score (-0.31).

We present the two distributions in Figure 6. Clearly, there is a stark difference in expected payoffs between the parties. Specifically, for the party that is less powerful, but nonetheless receives half the portfolios, the figure shows that the median predicted policy risk payoff is over *50% higher* (i.e., less favorable) than for the more powerful party. That is, the party that is undercompensated with respect to the raw number of portfolios can expect to translate its greater bargaining power into a significantly more attractive, albeit less directly observable, set of policy benefits.

These findings provide clear support for our argument. But there are two final issues we need to consider. The first concerns the presence of heterogeneity in the effects of the covariates, as captured by the estimated variance components. As noted earlier, we can use these estimates to predict random effects, which tell us how much the estimated intercepts and slopes in individual bargaining situations deviate from the overall (or average) estimates. In SI Appendix H, we present the estimated effect of the undercompensation variable (i.e., the estimated average effect across bargaining situations *plus* the situation-specific estimated random effect) for each of the 308 governments in the sample.²⁶ We do find variation across bargaining situations in the impact of undercompensation on policy risk. There are a few cases in which the effect is in the opposite direction, and a few in which undercompensation has no significant effect. But in the overwhelming majority (approximately 90% of bargaining situations), the effect is statistically significant ($p < 0.05$) and in the expected direction. Notably, in approximately 25% of cases, the estimated effect is over *35% larger* in magnitude than the average effect shown in Table 3.

The second issue relates to the fact that one component of our undercompensation variable—bargaining power—is itself an *estimate* produced by another model. As such, it is measured with error, which could lead to biased parameter estimates. Given the predictive success of the KOR model, we do not expect the amount of estimation error to be so great as to alter our conclusions. However, this is ultimately an empirical question. Accordingly, we have chosen to conduct a sensitivity analysis using *100 alternative datasets* based on the coefficient estimates and standard errors from the KOR model (see SI Appendix I for a more detailed description of our data generation procedure). Each new dataset contains a new estimate of party bargaining power, and thus a new measure of

²⁶ The estimates from Table 3 also indicate heterogeneity in the effect of the policy distance control variable, but we focus in the appendix only on our main variable of interest.



Note: The dashed horizontal line is the estimated coefficient for *Undercompensation in Portfolio Share* from the model shown in Table 3, encompassed by a 95% confidence region. Circles represent the coefficient estimates from 100 simulated data sets (with 95% confidence bounds), arranged in ascending order of estimation error in bargaining power. The solid lines are lowess smoothing curves for simulated estimates greater than and less than the original coefficient estimate.

FIGURE 7. *Impact of Measurement Error*

undercompensation. For each dataset, we reestimate our model and compare the results to those from Table 3.

In Figure 7, the dashed horizontal line represents the original coefficient estimate from Table 3, with the shaded region corresponding to its 95% confidence interval. The 100 dots in the figure are the coefficient estimates for the undercompensation covariate (along with 95% confidence bounds) from the alternative datasets. For easier visualization, we have ordered the estimates according to the amount of measurement error in the datasets (calculated as the *sum of squared deviations* between the original and new estimates of bargaining power). Also displayed are two lowess smoothing curves, one for estimates above the original coefficient, and one for estimates below. As one might expect, the curves move away from the original coefficient as measurement error increases. However, even for values of measurement

error in the tails, the new coefficient estimates fall within the original estimate’s confidence region, and in all cases we can reject the null of no effect with a very high level of confidence.²⁷

CONCLUSION

The negotiation of a coalition bargain constitutes one of the most consequential political events in multiparty parliamentary democracies. We have highlighted a central feature of such bargaining that has only received implicit treatment by scholars, but that has profound consequences for coalition negotiations: the fact that party elites must consider how the bargains they strike will be perceived by their supporters. The argument rests on two foundational assumptions. First, supporters vary in their political sophistication. Less knowledgeable supporters rely more heavily on simple heuristics to assess their party’s performance. Second, the goods that parties negotiate over vary in the degree to which they can be assessed by less knowledgeable supporters. Some aspects of the bargain are relatively easy to observe, while others are more complex and difficult for less sophisticated supporters to evaluate. The combination of these features generates incentives for party leaders to strike bargains that distribute goods that less knowledgeable supporters are more likely to be aware of—and can assess with respect to some heuristic—in proportion to the expectations generated by their heuristic, and to use those goods that these supporters are less likely to be aware of as side payments to reconcile the overall bargain with the underlying bargaining power of the parties. Our analysis of the allocation of ministries in the process of government formation and the “policy risk” that emanates from this allocation provides clear initial support for the argument.

Our argument and findings have broad implications. Most immediately, it offers a way to reconcile the long-standing tension between the empirical regularity of Gamson’s Law and the theoretical expectation that coalition bargains should reflect party bargaining power. More importantly, the argument has profound implications for political accountability, for the incentives of parties to be responsive to the preferences of citizens, and thus for the quality of representation. The theory and findings suggest that what coalition governments *appear* to do (as reflected in the easily observable features of their bargains) will typically not be an accurate reflection of the *actual* influence of the coalition partners.

²⁷ We also conduct an additional analysis, presented in SI Appendix J, in which we substitute a leading alternative index of bargaining power (the normalized Banzhaf index) for our model-derived measure. The results are very similar to what we report in Table 3.

In return for allowing their weaker partners to “look good” with respect to those pieces of the bargain that supporters can easily evaluate, parties that are in a strong bargaining position are able to extract significant concessions in areas that are less amenable to supporter scrutiny. Put differently, the reality of who shapes coalition government outcomes may diverge from supporter perceptions on the basis of what they can easily observe. This raises obvious hurdles to the effectiveness with which supporters can hold parties accountable. It also implies—with respect to the quality of representation—that those interests represented by parties with strong bargaining positions exercise more influence than may be apparent on the basis of the “observable” aspects of coalition governance.

We close with a brief discussion of avenues for future work. First, there is obvious value in expanding empirical applications beyond portfolio allocations to other, less tangible objects of coalition bargaining (such as budget allocations or policy compromises). More importantly, future work should move beyond the evidence we have presented here—which focuses on the behavioral implications of our argument for bargaining outcomes—to include an analysis of supporter (especially voter) behavior. Specifically, building on such work as that of Lin et al. (2017), Duch and Stevenson (2013), and Bowler, Gschwend and Indridason (N.d.), it is critical to investigate the following questions: To what extent do supporters employ heuristics to evaluate party performance under coalition governance? And to what extent does future support for a party depend on meeting supporter expectations induced by these heuristics?

There are also important directions for further theoretical development. One path relates to the “observability” of goods. We have treated this as an *exogenously* given quality of the objects of coalition bargaining, and we think this represents a useful approximation. However, observability is also, within limits, subject to manipulation by parties—most obviously through prominent and clear public commitments to particular results that can be reasonably easily “checked” by supporters. The fact that party elites can attempt to shape observability opens up rich strategic possibilities that represent a fruitful area for further research. A second path is suggested by our empirical results. The model fits the data extremely well, but there remains unexplained variation in policy risk. As we have shown, this variation is mostly attributable to differences across bargaining situations, rather than differences between parties. This opens up interesting theoretical possibilities. One is that differences between bargaining situations are related to heterogeneity in voter knowledge about coalition outcomes:

If supporters are generally more aware of the policy implications of different portfolio distributions in some contexts, it may be more difficult for parties to use more qualitative aspects of portfolio allocation to achieve numerical proportionality. Institutional features of the bargaining environment may also play an important role. For example, the presence of written coalition agreements may be a significant factor that allows parties both to secure other side payments and to constrain policymaking in ways that interact with the particular allocation of portfolios (which is our focus here). Significantly, scholars are already making important theoretical and empirical progress in understanding how such agreements shape coalition governance (Müller and Strøm 2008; Klüver and Bäck 2019). While the data that would be required to incorporate this dimension into the analysis are not yet available, moving in this direction is an obvious avenue for future work.

REFERENCES

- Bäck, Hanna, Henk Erik Meier and Thomas Persson. 2009. "Party Size and Portfolio Payoffs: The Proportional Allocation of Ministerial Posts in Coalition Governments." *Journal of Legislative Studies* 15(1):10–34.
- Bäck, Hanna, Marc Debus and Patrick Dumont. 2011. "Who Gets What in Coalition Governments? Predictors of Portfolio Allocation in Parliamentary Democracies." *European Journal of Political Research* 50(4):441–78.
- Bassi, Anna. 2013. "A Model of Endogenous Government Formation." *American Journal of Political Science* 57(4):777–93.
- Benoit, Kenneth and Michael Laver. 2006. *Party Policy in Modern Democracies*. London: Routledge.
- Bowler, Shaun, Thomas Gschwend and Indridi H. Indridason. N.d. "Coalition Policy Perceptions." *Journal of Politics*. In Press. <https://doi.org/10.1086/708629>.
- Browne, Eric C. and Mark N. Franklin. 1973. "Aspects of Coalition Payoffs in European Parliamentary Democracies." *American Political Science Review* 67(2):453–69.
- Carroll, Royce and Gary W. Cox. 2007. "The Logic of Gamson's Law: Pre-election Coalitions and Portfolio Allocations." *American Journal of Political Science* 51(2):300–313.
- Downs, Anthony. 1957. *An Economic Theory of Democracy*. New York: HarperCollins Publishers.
- Duch, Raymond M. and Randolph T. Stevenson. 2013. "Voter Perceptions of Agenda Power and Attribution of Responsibility for Economic Performance." *Electoral Studies* 32(3):512–516.
- Falcó-Gimeno, Albert and Indridi H. Indridason. 2013. "Uncertainty, Complexity, and Gamson's Law: Comparing Coalition Formation in Western Europe." *West European Politics* 36(1):221–47.
- Fortunato, David. 2019. "Legislative Review and Party Differentiation in Coalition Governments." *American Political Science Review* 113(1):242–47.
- Fortunato, David and Randolph T. Stevenson. 2019. "Heuristics in Context." *Political Science Research and Methods* 7(2):311–30.
- Fortunato, David, Randolph T. Stevenson and Greg Vonnahme. 2016. "Context and Political Knowledge: Explaining Cross-National Variation in Partisan Left-Right Knowledge." *Journal of Politics* 78(4):1211–1228.

- Gamson, William A. 1961. "A Theory of Coalition Formation." *American Sociological Review* 26(3):373–82.
- Goldstein, Daniel G. and Gerd Gigerenzer. 2012. "Models of Ecological Rationality: The Recognition Heuristic." *Psychological Review* 109(1):75–90.
- Healy, Andrew and Neil Malhotra. 2013. "Retrospective Voting Reconsidered." *Annual Review of Political Science* 16(1):285–306.
- Kahneman, Daniel and Amos Tversky. 1979. "Prospect Theory: An Analysis of Decision Under Risk." *Econometrica* 47(2):263–291.
- Kayser, Mark A., Matthias Orlowski and Jochen Rehmert. N.d. "Coalition Inclusion Probabilities: A Dynamic Measure of Party Competitiveness and Cabinet Leverage." Unpublished Manuscript, Hertie School of Governance. http://mark-kayser.com/papers/KOR_CoalProbs_190125.pdf.
- Klüver, Heike and Hanna Bäck. 2019. "Coalition Agreements, Issue Attention, and Cabinet Governance." *Comparative Political Studies* 52(13–14):1995–2031.
- Laver, Michael and Kenneth A. Shepsle. 1996. *Making and Breaking Governments: Cabinets and Legislatures in Parliamentary Democracies*. Cambridge, UK: Cambridge University Press.
- Laver, Michael and Norman Schofield. 1998. *Multiparty Government: The Politics of Coalition in Europe*. Ann Arbor: University of Michigan Press.
- Laver, Michael, Scott de Marchi and Hande Mutlu. 2011. "Negotiation in Legislatures over Government Formation." *Public Choice* 147(3–4):285–304.
- Lin, Nick, Randolph T. Stevenson, Mathias Wessel Tromborg and David Fortunato. 2017. "Gamson's Law and Voters' Perceptions of Portfolio Allocation." *European Journal of Political Research* 56(4):912–40.
- Lodge, Milton, Kathleen M. McGraw and Patrick Stroh. 1989. "An Impression-Driven Model of Candidate Evaluation." *American Journal of Political Science* 83(2):399–419.
- Lodge, Milton, Marco Steenbergen and Shawn Brau. 1995. "The Responsive Voter: Campaign Information and the Dynamics of Candidate Evaluation." *American Political Science Review* 89(2):309–326.
- Lowe, Will, Kenneth Benoit, Slava Mikhaylov and Michael Laver. 2011. "Scaling Policy Preferences from Coded Political Texts." *Legislative Studies Quarterly* 36(1):123–55.

- Martin, Lanny W. and Georg Vanberg. 2011. *Parliaments and Coalitions: The Role of Legislative Institutions in Multiparty Governance*. Oxford: Oxford University Press.
- Martin, Lanny W. and Georg Vanberg. 2015. Coalition Formation and Policymaking in Parliamentary Democracies. In *Handbook of Comparative Political Institutions*, ed. Jennifer Gandhi and Rubén Ruiz-Ruffino. New York: Routledge. pp. 181–194.
- 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(3):503–18.
- Montero, Maria and Juan Vidal-Puga. 2011. “Demand Bargaining and Proportional Payoffs in Majority Games.” *Games and Economic Behavior* 71(2):395–408.
- Morelli, Massimo. 1999. “Demand Competition and Policy Compromise in Legislative Bargaining.” *American Political Science Review* 93(4):809–20.
- Müller, Wolfgang C. and Kaare Strøm. 2008. Coalition Agreements and Cabinet Governanc. In *Cabinets and Coalition Bargaining: The Democratic Life Cycle in Western Europe*, ed. Kaare Strøm, Wolfgang C. Müller and Torbjörn Bergman. Oxford: Oxford University Press. pp. 159–200.
- Norton, E.C., G.S. Bieler, S.T. Ennett and G.A. Zarkin. 1996. “Analysis of Prevention Program Effectiveness with Clustered Data Using Generalized Estimating Equations.” *Journal of Consulting and Clinical Psychology* 64(5):919–26.
- Schelling, Thomas. 1960. *The Strategy of Conflict*. Cambridge, MA: Harvard University Press.
- Strøm, Kaare. 1990. “A Behavioral Theory of Competitive Political Parties.” *American Journal of Political Science* 34(2):565–98.
- Tullock, Gordon. 1967. *Towards a Mathematics of Politics*. Ann Arbor: University of Michigan Press.
- Verzichelli, Luca. 2008. Portfolio Allocation. In *Cabinets and Coalition Bargaining: The Democratic Life Cycle in Western Europe*, ed. Kaare Strøm, Wolfgang C. Müller and Torbjörn Bergman. Oxford: Oxford University Press. pp. 237–68.
- Warwick, Paul V. and James N. Druckman. 2006. “The Portfolio Allocation Paradox: An Investigation into the Nature of a Very Strong but Puzzling Relationship.” *European Journal of Political Research*

45(4):635–65.