

Portfolio Salience and the Proportionality of Payoffs in Coalition Governments

PAUL V. WARWICK AND JAMES N. DRUCKMAN

A fundamental divide has emerged over how portfolio payoffs are distributed among parties in parliamentary coalitions. On one side lies very strong empirical evidence that the parties in a governing coalition tend to receive portfolios in one-to-one proportion to the amount of legislative support they contribute to the coalition, with perhaps some slight deviations from proportionality coming at the expense of larger parties that lead coalition negotiations. On the other side of the debate lies a stream of formal theories that suggest the opposite – that parties in charge of coalition negotiations ought to be able to take a disproportionately large share of portfolio benefits for themselves. In this article, we address this disjuncture by re-examining the empirical connection between legislative seats and portfolio payoffs with the aid of a new and more extensive dataset, a different method of analysis, and what we see as a more valid operationalization of the dependent variable. This operationalization involves the inclusion, for the first time, of evidence concerning the importance or salience of the portfolios each party receives, as opposed to just their quantity. The article concludes with an assessment of the implications of our findings for the debate over the rewards of coalition membership in parliamentary democracies.

The study of coalition behaviour in parliamentary regimes constitutes one of the most dynamic research projects in political science. Most of the attention in this burgeoning literature has been devoted to two basic issues: which parties will succeed in forming the government and how long will that government survive in office? In recent years, however, a surprising and fundamental divide has emerged over a third aspect of coalition behaviour – how well rewarded can each member-party expect to be in terms of ministerial portfolios?

On one side of this debate lies one of most impressive of empirical findings in all of social science. That finding, first demonstrated by Browne and Franklin,¹ is that the parties in a governing coalition tend to receive portfolios in one-to-one proportion to the amount of legislative support they contribute to the coalition. For example, if a party's share of parliamentary seats constitutes 20 per cent of the total number of seats held by the coalition, then the party can expect to be allocated approximately 20 per cent of the ministerial portfolios. This pattern turns out to hold in parliamentary contexts to a remarkably high degree.

* Department of Political Science, Simon Fraser University, Burnaby, BC; and Department of Political Science, University of Minnesota, respectively. We wish to thank John Huber for supplying us with a computer version of the WKB data, which aided our own coding efforts considerably, and Massimo Morelli for providing useful comments on the nature of the formal models discussed here.

¹ Eric C. Browne and Mark N. Franklin, 'Aspects of Coalition Payoffs in European Parliamentary Democracies', *American Political Science Review*, 67 (1973), 453–69.

The other side of the debate is represented by a number of recent formal theories which suggest that the party in charge of coalition negotiations ought to be able to use that leadership position to extract a disproportionately large share of coalition payoffs.² Since parties are office-seekers in these theories, the payoffs in question derive primarily from holding cabinet portfolios.³ This argument appears to be fundamentally at odds with the high degree of proportionality evident in the allocation of portfolios. Even more unsettling is the empirical evidence that any systematic deviations from proportionality tend to go in precisely the opposite direction to that which these models anticipate. Specifically, Browne and Franklin found that the relatively minor deviations from proportionality in the allocation of portfolios work to the benefit of the smaller coalition parties. As a result, larger parties that tend to be assigned the role of putting together a government are – if anything – slightly under-compensated relative to the proportionality standard.

We thus have a major disjuncture between the dominant theoretical direction on the issue of coalition payoffs and the available evidence. This article seeks to address this disjuncture by re-examining the connection between parliamentary seats and coalition payoffs with the aid of a new and more extensive dataset, a different method of analysis, and what we see as a more valid operationalization of the dependent variable. The last-mentioned feature of the analysis is the most significant because it signals a major weakness in the empirical literature: previous analyses have invariably equated a party's portfolio payoff with its share of portfolios, treating each portfolio as equal in value. This approach ignores the strong likelihood that parties view some portfolios as providing higher payoffs than others; the prime ministership, for instance, is generally seen as more valuable than the public works portfolio. In this article, we address this gap by incorporating, for the first time, systematic evidence about the value or importance of cabinet portfolios into the analysis. Through these and other means, we aim to throw new light on the festering debate over the rewards of coalition membership and the assumptions that should inform theories of coalition government in this regard.

THE COALITION PAYOFFS DEBATE

When two or more parties bargain to form a governing coalition, they face the fundamental task of deciding which and how many portfolios each party will receive. While some approaches, most notably Laver and Shepsle's portfolio allocation theory,⁴ are driven by the 'which' question, others have focused on

² Examples include David P. Baron and John Ferejohn, 'Bargaining and Agenda Formation in Legislatures', *American Economic Review*, 77 (1987), 303–9; David P. Baron and John Ferejohn, 'Bargaining in Legislatures', *American Political Science Review*, 83 (1989), 1181–206; and Joseph E. Harrington Jr, 'The Power of the Proposal Maker in a Model of Endogenous Agenda Formation', *Public Choice*, 64 (1990), 1–20.

³ Baron and Ferejohn, 'Bargaining in Legislatures', p. 1194.

⁴ Michael Laver and Kenneth A. Shepsle, *Making and Breaking Governments* (Cambridge: Cambridge University Press, 1996).

the issue of 'how many'. The initial conjecture in this domain came from Gamson, who asserted that 'Any participant will expect others to demand from a coalition a share of the payoff proportional to the amount of resources which they contribute to a coalition'.⁵ Browne and Franklin subsequently tested Gamson's prediction using data on coalition governments in thirteen parliamentary countries from 1945 to 1969.⁶ Equating 'resources' with the proportion of parliamentary seats that each party contributes to the coalition and 'payoffs' with each party's share of ministries (i.e. ministerial portfolios), they regressed the former variable on the latter and found striking support for Gamson's conjecture: not only were the two variables very closely related ($r = 0.926$), but the estimated intercept and slope were very close to the values of 0 and 1, respectively, which would indicate perfect one-to-one proportionality.

This proportionality relationship has been described as 'one of the strongest relationships to be found anywhere in the realm of the social sciences'⁷ and subsequently labelled, with perhaps some hyperbole, as 'Gamson's Law' of proportionality.⁸ While unquestionably impressive, however, the degree of proportionality that Browne and Franklin found between seats and portfolios was less than total. In fact, they also detected a slight tendency for smaller parties to receive more than their proportional share of portfolios (and for larger parties to receive correspondingly less). Their analysis showed that this 'small-party bias', as we shall term it, became more pronounced as the coalition's size (number of members or parties) decreases.⁹

In view of subsequent developments, it is particularly interesting that Browne and Franklin's analysis implies that the *formateur* – the party charged with putting together a government – will usually be the party that is under-compensated whenever a small-party bias appears.¹⁰ This does not occur simply because the formateur tends to be the largest party in the coalition (although this is indeed the case).¹¹ Rather, it follows from their explanation for the bias. In their view, the bias occurs because the dominant party in a coalition, if one exists, is willing to surrender one or more ministries to which it is entitled under

⁵ William A. Gamson, 'A Theory of Coalition Formation,' *American Sociological Review*, 26 (1961), p. 376.

⁶ Browne and Franklin, 'Aspects of Coalition Payoffs in European Parliamentary Democracies'.

⁷ Michael Laver and Norman Schofield, *Multiparty Government: The Politics of Coalition in Europe* (Oxford: Oxford University Press, 1990), p. 171.

⁸ Massimo Morelli, 'Demand Competition and Policy Compromise in Legislative Bargaining', *American Political Science Review*, 93 (1999), 809–20.

⁹ Browne and Franklin labelled it the 'relative weakness effect'; we prefer the term 'small-party bias' because it is more descriptive of the nature of the effect.

¹⁰ Technically, the formateur is the individual who leads negotiations and is expected to be the prime minister. However, we use the term to refer to the formateur's party since it is the party, not the individual, that makes the sacrifice (or captures the excess under the other interpretation) in portfolio benefits.

¹¹ In our dataset, covering approximately forty-five years of coalitions in twelve West European countries, the formateur is the largest party in 84 per cent of the governments.

the proportionality norm because doing so will not jeopardize its dominant (leadership) position and may ensure necessary support from smaller parties.¹² They believe that this accounts for the tendency for the bias to decrease as the number of parties in the coalition increases: more parties constitute a greater threat to the leader's dominance. Since the formateur's party typically leads the coalition formation process and assumes the prime ministership, it follows that the dominant party will usually be the formateur's party.

Browne and Franklin's findings are so strong that few have bothered to conduct further empirical studies of the distribution of coalition payoffs. The rare empirical investigations that have appeared during the last twenty-five years have largely confined themselves to refining the proportionality finding and the small-party bias (albeit in some interesting ways).¹³ However, this consensus by default has recently been challenged by a stream of formal models that generate predictions contradicting both effects.

A highly influential example is Baron and Ferejohn's legislative bargaining model.¹⁴ This model is premised upon what would seem to be rather non-contentious assumptions about the coalition game. In the model, multiple members of a legislature (such as multiple parties) bargain over how to distribute

¹² Browne and Franklin, 'Aspects of Coalition Payoffs in European Parliamentary Democracies', pp. 461–3.

¹³ Eric C. Browne and John P. Frendreis, 'Allocating Coalition Payoffs by Conventional Norm: An Assessment of the Evidence from Cabinet Coalition Situations', *American Journal of Political Science*, 24 (1980), 753–68, show that Browne and Franklin's findings are robust to using a step function instead of linear regression, and that the proportionality norm and the small-party bias apply not only to the distribution of all ministries but also to the distribution of a subset of seemingly desirable ministries (cf., Bruce Bueno de Mesquita, 'Coalition Payoffs and Electoral Performance in European Democracies', *Comparative Political Studies*, 72 (1979), 61–81). Norman Schofield and Michael Laver, 'Bargaining Theory and Portfolio Payoffs in European Coalition Governments 1945–83', *British Journal of Political Science*, 15 (1985), 143–64, replicate Browne and Franklin's analysis and show that a measure of bargaining power (instead of seat share) does not do a better job of predicting allocations, except in multipolar systems. Other recent work on different aspects of portfolios includes David Austen-Smith and Jeffrey Banks, 'Stable Governments and the Allocation of Policy Portfolios', *American Political Science Review*, 84 (1990), 891–906; Laver and Shepsle, *Making and Breaking Governments*; John D. Huber, 'How Does Cabinet Instability Affect Political Performance? Portfolio Volatility and Health Care Cost Containment in Parliamentary Democracies', *American Political Science Review*, 92 (1998), 577–91; and Michael F. Thies, 'Keeping Tabs on One's Partners: The Logic of Delegation in Coalition Formation', *American Journal of Political Science*, forthcoming 2001.

¹⁴ Baron and Ferejohn, 'Bargaining and Agenda Formation in Legislatures' and 'Bargaining in Legislatures'. The basic structure of the model has been widely applied (e.g., David P. Baron, 'Comparative Dynamics of Parliamentary Governments', *American Political Science Review*, 92 (1998), 593–609; Daniel Diermeier and Timothy J. Feddersen, 'Cohesion in Legislatures and the Vote of Confidence Procedure', *American Political Science Review*, 92 (1998), 611–21), and has been identified as a 'common reference point' for non-cooperative office-seeking models of coalition formation (Morelli, 'Demand Competition and Policy Compromise in Legislative Bargaining', p. 809). Although Baron and Ferejohn ('Bargaining in Legislatures', pp. 1193–5) examine legislative bargaining under an open and a closed rule, we focus on the closed rule model with infinite sessions, as this is the model more analogous to coalition bargaining.

benefits. As Baron and Ferejohn make very clear, these benefits consist of the various ministerial portfolios in parliamentary systems.¹⁵ Each party cares only about maximizing its share of these benefits, subject to a discount factor.¹⁶ The game begins with the recognition of one party as proposer or formateur; that party then proposes a government coalition and a way to distribute the portfolios among parties in the coalition. The parties vote on the proposal, using majority rule. If it passes, the game ends and the parties receive the proposed distribution. If the proposal fails, the process repeats.

The model's key result for our purposes is that the formateur always receives a disproportionately *large* share of portfolios; in other words, the formateur is over-compensated relative to the proportionality standard. This result stems from the principle of majority rule combined with the formateur's agenda-setting power. Majority rule entails that some parties will be excluded from the coalition as unnecessary for its survival; as a result, the parties in the coalition have access to extra benefits that would have gone to the excluded parties had the coalition been all-inclusive.¹⁷ This means that the formateur can successfully make a proposal in which it keeps these extra benefits for itself because the other potential coalition members will still receive as much as their expected payoff from rejecting the proposal (and allowing the game to go to another round).¹⁸

The conflict between the theoretical expectation of formateur over-compensation in terms of portfolios and the empirical finding of proportionality (together with a slight small-party bias) indicates the existence of a problem with either the theory or the empirical analysis. There are reasons to indict both suspects. In a recent paper, Morelli challenges the theoretical work by proposing a bargaining model that produces results consistent with both the proportionality principle and the small-party bias.¹⁹ In Morelli's model, multiple parties, none of whom possesses a majority of legislative seats, must decide how to form a majority coalition and distribute benefits. As always, the parties are motivated to maximize their share of private benefits or portfolios, subject to a discount

¹⁵ Baron and Ferejohn, 'Bargaining in Legislatures', pp. 1193–4. Subsequent related models, such as those of Baron, 'Comparative Dynamics of Parliamentary Governments' and Diermeier and Feddersen, 'Cohesion in Legislatures and the Vote of Confidence Procedure', see benefits in terms of government income to be used to reward supporters. However, they assume that the ability to reward supporters depends on holding (cabinet) office (e.g., Baron, 'Comparative Dynamics of Parliamentary Governments', p. 598).

¹⁶ The discount factor simply registers the extent to which a party would prefer to enter a government in the present round of the game, rather than waiting for the next round of the game (other things being equal).

¹⁷ Baron and Ferejohn, 'Bargaining in Legislatures', p. 1193.

¹⁸ Under certain assumptions about how formateurs are selected, such as the assumption that the probability of selection is proportional to seat shares, the expected or ex ante distribution of payoffs can also be proportional to seat shares. However, portfolios are distributed ex post, that is, after the formateur has been selected. At this stage, the formateur is in a position to extract a disproportionately large share of payoffs.

¹⁹ Morelli, 'Demand Competition and Policy Compromise in Legislative Bargaining'.

factor.²⁰ The bargaining begins when the head of state selects a formateur party. The formateur's role is to choose an order in which all parties can make demands, starting with itself. The parties then make sequential demands according to this order of play – that is, each party demands a certain share of the portfolios in turn. A majority coalition forms when the sum of benefits demanded by a sequence of parties that together hold a majority of seats does not exceed the total amount of available benefits. If no coalition forms after all parties have made their demands, the head of state chooses another formateur and the sequential demand game repeats. If no coalition is formed after a finite number of rounds, then a caretaker government forms and no portfolios are distributed.

In Morelli's model, the formateur is distinguished only by the fact that it chooses the order in which sequential demands are made (and gets to make the first demand). This differs from Baron and Ferejohn's model, where the proposer makes a precise proposal about how portfolios should be distributed. Thus, the sequential demand dynamic strips the formateur of its agenda-setting power. Instead of choosing to accept the formateur's proposal or negotiate in another period, the other parties can now make demands so as to maximize their own share of benefits, given the demands of the parties that precede them. This change in specification profoundly alters the outcome of the game:

The subgame perfect equilibrium payoff distribution of this game is proportional to the ex ante distribution of bargaining power and approximately proportional to the distribution of seats in the winning coalition, consistent with Gamson's Law. Moreover, ... when the number of parties needed in a majority coalition is small, the smaller parties receive more than their relative share of seats in the coalition, and the larger parties receive less ... In the extreme case in which only three parties play the bargaining game, the distribution of ministerial ... payoffs tends to be an equal split ...²¹

With the introduction of Morelli's model, we are faced with opposed theoretical expectations concerning portfolio payoffs: according to the agenda-setting school, formateurs will be over-compensated; according to the sequential-demand framework, proportionality will prevail, albeit with some small-party bias (and consequently formateur under-compensation). The available empirical work appears to support the latter view, as we have seen. Yet that work, too, can be challenged: by operationalizing each party's payoffs in terms of the simple proportion of portfolios it is allocated, previous data analyses have omitted the possibility that a party receives a greater payoff or

²⁰ Morelli also includes a policy component in each party's utility function. This generates some interesting dynamics, but is tangential to the question of portfolio allocation and we do not discuss it further.

²¹ Morelli, 'Demand Competition and Policy Compromise in Legislative Bargaining', p. 810.

benefit when it obtains some portfolios rather than others.²² No account has been taken, for example, of the probability that the prime ministership (which usually goes to the formateur) is worth more than a run-of-the-mill portfolio. Needless to say, this consideration potentially undercuts any empirical conclusions reached about portfolio payoffs.

What is clearly needed is some means of assessing payoffs not just in terms of the number of portfolios each party receives but also in terms of the saliency or importance of these portfolios. In the next section, we describe the characteristics of a new dataset that includes weights that can be assigned to the various portfolios to capture variations in saliency levels. With qualitative distinctions as well as quantitative differences incorporated into the measurement of portfolio payoffs, the stage will be set for a closer assessment of these alternative perspectives on the rewards of coalition membership.

DATA AND MEASUREMENT ISSUES

The dataset that we have built to perform this task covers governments in twelve West European parliamentary democracies from the resumption of normal politics from the Second World War until the end of 1989.²³ All governments in these countries were included, with the exception of single-party governments, non-party governments and caretaker governments. These exclusions reflect our focus on the allocation of portfolio benefits among coalition parties. It should be noted that the identification of governments in these systems is not as straightforward as it might seem, since the criteria for determining when one government ends and another begins often differ from one study to another. Here, we follow the identification of governments provided by Woldendorp, Keman and Budge,²⁴ since this is our source for the allocations of portfolios

²² The existence of differences in portfolio saliences has been widely recognized. Laver and Schofield, *Multiparty Government*, pp. 181–2, for example, suggest that it is conventional to see the prime ministership as the top portfolio, followed by a tier of senior portfolios, and two lower tiers of portfolios. Eric C. Browne and Karen A. Feste, 'Qualitative Dimensions of Coalition Payoffs: Evidence for European Party Governments 1945–70', *American Behavioral Scientist*, 18 (1975), 530–56, find that a plausible ordering of portfolios can be established by examining how often portfolios are held by the largest and therefore most influential party. Ian Budge and Hans Keman, *Parties and Democracy: Coalition Formation and Government Functioning in Twenty States* (Oxford: Oxford University Press, 1990), point out that portfolio preferences may vary across parties and investigate how certain types of parties prefer particular portfolios (e.g., agrarian parties prefer the agriculture portfolio). The important point to note, however, is that no studies incorporate portfolio saliency weights in a cross-national statistical analysis of coalition payoffs.

²³ The countries include Austria, Belgium, Denmark, Finland, Iceland, Ireland, Italy, Luxembourg, the Netherlands, Norway, Sweden and West Germany. Systems that are not fully parliamentary, such as the French Fifth Republic, were excluded so that institutional differences do not influence the results.

²⁴ Jaap Woldendorp, Hans Keman and Ian Budge, 'Political Data 1945–1990: Party Government in 20 Democracies', *European Journal of Political Research*, 24 (1993), 1–120; Jaap Woldendorp, Hans Keman and Ian Budge, 'Party Government in 20 Democracies: An Update (1990–1995)', *European Journal of Political Research*, 33 (1998), 125–64

among parties. ‘WKB’, as we shall term them, define the termination of a government as occurring whenever there is an election, a change in the prime minister or the party composition of the cabinet, or a government resignation.²⁵ Based on these criteria, the dataset comprises a total of 607 parties belonging to some 200 coalition governments.

With respect to independent variables, we are primarily concerned with the resources that each party brings to the coalition. We follow Browne and Franklin in operationalizing resources as the percentage of legislative seats that the party contributes to the coalition, which we will refer to as its ‘seat contribution’.²⁶ In addition, we use a dummy variable, ‘formateur status’, to identify whether (1) or not (0) the party provided the government formateur. This variable will enable us to capture any tendencies towards formateur over-compensation or under-compensation relative to the proportionality principle. Both the seat contributions of parties and their formateur status were taken from *Keesing’s Contemporary Archives*.²⁷ Consistent with previous research, our initial dependent variable, ‘portfolio share’, records each party’s proportion of portfolios relative to all other parties in the coalition. As noted above, information on portfolio allocations was taken from the WKB data.

Browne and Franklin’s finding of a strong element of proportionality in the allocation of portfolio shares is doubly puzzling, as we have seen. The first puzzle is why the allocation should approximate so closely what ‘fairness’ would mandate, given that many models of the coalition game (such as the Baron and Ferejohn model and its heirs) find that the formateur is in a position to extract more than its fair or proportional share. The second puzzle is why proportionality, if that is the guiding principle, should be evident in the *quantitative* allocation of portfolios, where no account is taken of portfolio salience. It is likely that these puzzles are intertwined and, to unravel both of them, some means of assessing the salience or importance of individual portfolios is required.

In view of the need to take the value of portfolios into account, it is striking that this consideration has been so largely ignored in empirical investigations of coalition payoffs. This is not because of information on portfolio saliences is unavailable. In 1989, Laver and Hunt conducted a survey of experts on a variety of democratic systems in which respondents were asked to list the key cabinet positions in ‘their’ country and to rank them in order of importance. The

²⁵ Woldendorp, Keman and Budge, ‘Political Data 1945–1990’, p. 5.

²⁶ We do not employ the measure of bargaining power used by Schofield and Laver (‘Bargaining Theory and Portfolio Payoffs in European Coalition Governments 1945–83’, pp. 154–8), because they find that seat contribution is the best overall predictor. Also, note that measures of seat contribution and bargaining power tend to be highly related (Morelli, ‘Demand Competition and Policy Compromise in Legislative Bargaining’, p. 810).

²⁷ *Keesing’s Contemporary Archives* (London: Keesing’s Publications, 1945–90).

responses were then published in the form of a set of mean portfolio rankings for each of these countries.²⁸

While Laver and Hunt's efforts to measure portfolio importance clearly broke new ground, the exploitation of their data may have been inhibited by doubts about its usability. As is the case with any subjective ranking, a danger exists that the respondents were influenced by the relationship under study (for example, the rankings were affected by knowledge about which parties occupied which portfolios). This danger seems minimal in the present case; indeed, it is difficult to see how any such propensity would affect the rankings systematically, and it also seems more than likely that experts would have focused on the importance of the various portfolios for policy making and/or patronage. A related concern is that the rankings do not pick up differences among parties in the importance attributed to given portfolios, even though some previous research has indicated that these can be significant.²⁹ As before, however, it is not clear how differences of this sort would impinge upon bargaining over payoffs; the theoretical work has always assumed agreement on the value of the various 'spoils' of office.

Neither of the above-mentioned limitations seems prohibitive, but there are certain practical limitations in the Laver and Hunt (LH) data that clearly stand in the way of constructing a measure of portfolio shares that incorporates them. One such limitation is the failure of Laver and Hunt's expert respondents to rank all portfolios.³⁰ In fact, a little more than half the portfolios are ranked, on average, and the unranked positions include both the head of government or prime minister and the deputy prime minister. In addition, the ranked portfolios are taken from the set of portfolios that existed at the time of the survey. This means not only that changes over time in the importance of a given portfolio are not captured but also that changes in the way ministerial responsibilities are grouped can pose problems. For example, if a portfolio in existence in 1989 combined two jurisdictions that had previously been given to separate ministers, Laver and Hunt provide only a ranking for the combined ministry, not the separate ministries that once existed. Finally, while the rankings give us some idea of how a typical expert in each country would order certain portfolios in terms of their importance, they do not tell us how much more important a given portfolio is in comparison with the one that follows or precedes it in the ranking. In other words, they provide an ordinal ranking of portfolios rather than a cardinal scale of portfolio weights.

²⁸ Michael Laver and W. Ben Hunt, *Policy and Party Competition* (New York: Routledge, 1992). It is worth noting that they collected these data with an eye towards the type of exercise that we are undertaking here (p. 105).

²⁹ This was a major finding in Budge and Keman, *Parties and Democracy*.

³⁰ An extension of the LH rankings is provided in Wolfgang C. Müller and Kaare Strøm (eds), *Regierungskoalitionen in Westeuropa* (Vienna: Signum Verlag, 1997). We chose not to use these extended rankings because the additional information consists of the ranking of minor portfolios (and thus the distinctions are not likely to be clear cut) and are based on the judgement of only one or two experts.

These limitations are certainly serious, but they did not strike us as insuperable barriers to the exploitation of this unique source for portfolio ratings. We therefore decided to develop means of overcoming them as much as possible. Our procedure involved the following steps. To fill out the rankings, we gave the prime ministership the top ranking in each country and the deputy prime ministership the position immediately below it. The portfolios ranked in the LH data then followed in the order of their mean rankings, with all unranked portfolios placed at the bottom of the scale, that is, one rank below the lowest ranked portfolio. In justification for these decisions, we note that the placement of the prime minister and deputy prime minister at the top of the rankings is consistent with the practice adopted by the country experts who contributed to the Müller and Strøm volume (see fn. 30). As for the other unranked portfolios, their placement at the bottom reflects Laver and Hunt's interpretation of these portfolios; as they observe, 'Most experts ignored minor portfolios; thus while the rankings for all key portfolios are reported, some less salient portfolios are not listed'.³¹

In dealing with inconsistencies between portfolio jurisdictions at the time of the survey and at other times, the key fact to note is that WKB treat portfolios that combine jurisdictions as separate portfolios. Thus, if a minister holds the 'Culture and Education' portfolio, he or she will be listed as holding two portfolios. In cases where the LH data provide only a ranking for the combined portfolio, we ranked the separate jurisdictions by dividing the ranking for the combined portfolio in half. For example, if Culture and Education is ranked eighth in the country in question, we gave Education and Culture each a rank of 16 and re-ordered the portfolios accordingly. This assumes that when the ministries were separate, they were each worth half the combined ministry, whereas in reality the decision to combine two separate ministries may indicate that they have become less valuable than they once were. Some assumption had to be made in these cases, however, and this one seemed ineluctable. This situation highlights an important limitation in the reliability of these portfolio rankings noted earlier, which is that we must assume that they are constant over time.

The final step was to create weights from these rankings and apply them to the portfolio shares. The procedure we adopted to derive weights was to invert the rankings by subtracting them from the total number of portfolios ever occupied in any one government in the given country. This ensured that the most important portfolio (the prime ministership) received the largest weight and the least important portfolios received the smallest weight, albeit a weight well above 1.³² These values were then used as weights to calculate a 'weighted portfolio share' variable for each party in each government.

³¹ Laver and Hunt, *Policy and Party Competition*, p. 133.

³² It is important not to allow the lowest category to receive this value because to do so would be to create enormous disparities in portfolio values. If portfolios are ranked from 1 to 15, for example, it would mean that the top portfolio (the prime ministership) is fifteen times as valuable as the bottom portfolios (which include all unranked portfolios) – which seems rather unlikely.

This procedure can be made clearer with an example. In the case of Luxembourg, Laver and Hunt provide rankings for nine portfolios. The largest cabinet in Luxembourg had eighteen portfolios. Therefore, under our scheme, the prime ministership in Luxembourg receives a value of 17 (18 minus its ranking of 1), the deputy prime ministership a value of 16, the nine ranked portfolios values from ranging 15 to 7, and all unranked portfolios the lowest value of 6. A weighted portfolio share was then calculated for each party in a governing coalition by summing up the weights of the portfolios it holds and dividing by the sum of the weights of all portfolios held by all coalition parties.

Clearly, a number of the assumptions that went into the construction of weighted portfolio shares could prove to be unwarranted and misleading. Apart from the fact that a substantial portion of portfolios have been placed in a residual ranking, the most threatening of these probably is the use of ordinal ranks as weights. This step assumes that the ordinal rankings approximate a cardinal scale; in other words, that the deputy prime ministership in Luxembourg is worth 16/17 or 94.1 per cent that of the prime ministership, that the prime ministership is worth 2.8 (17/6) times the value of the lowest or baseline portfolio, and so forth. It is possible, however, to recalibrate the scales in various ways to correct for certain kinds of flaws that become apparent in analyses that involve weighted portfolio shares. For instance, the weights can be exponentiated to increase the distances between portfolios near the top of the scale more than portfolios near the bottom; alternatively, logarithms can be taken to expand the bottom end of the scale. It is also possible to increase the value of the prime ministership so that it is more clearly distinguished from the other portfolios. We shall be fully prepared to take such steps as appropriate to produce a weighting system that better captures the underlying phenomenon; as will become apparent, the evidence provides ample grounds for believing that this objective can be satisfied reasonably well.

A RE-ANALYSIS OF PORTFOLIO PAYOFFS

Although the main concern in this article is with the introduction of portfolio saliences into the explanation of portfolio payoffs, we shall begin the data analysis with the numerical or quantitative payoffs to cabinet parties. Our purpose is to re-evaluate the small-party bias and, to the extent that it exists, to determine whether it derives from a tendency towards under-compensation of the formateur party. We shall then turn to the investigation of weighted portfolio shares in order to determine whether the formateur is compensated, or more than compensated, for any numerical shortfall by receiving portfolios of greater importance.

Quantitative Proportionality and the Small-Party Bias

As we have seen, previous empirical research has revealed a very strong pattern of proportional allocations together with a much weaker small-party bias. In Browne and Franklin's seminal analysis, the strength of the proportionality effect was indicated by the very high correlation ($r = 0.926$) between the portfolio shares of cabinet parties and the seat contributions.³³ The small-party bias was suggested by a slight deviation in their regression line from perfect one-to-one proportionality in a direction that indicated greater rewards for smaller parties.

Both of these effects are also present in our data. Although the sample is substantially different from Browne and Franklin's, the correlation between our portfolio share and seat contribution variables is virtually identical at $r = 0.925$. Moreover, when portfolio share is regressed on seat contribution, the estimated intercept and slope coefficients of 0.081 (s.e. = 0.005) and 0.755 (s.e. = 0.013), respectively, point to the presence of a small-party bias.³⁴ Specifically, it would appear that parties contributing less than about one-third of the cabinet's parliamentary basis are over-compensated relative to the proportionality standard and that parties larger than that are under-compensated.

Before we can accept this interpretation, however, we need to examine the possibility that the deviation from one-to-one proportionality is merely an artefact of the nature of the data. One source of concern in this regard is the 'lumpiness' of the dependent variable, portfolio share.³⁵ Although it is possible to imagine otherwise, in reality portfolios are always allocated to single parties. This means that if there are twenty portfolios in a given cabinet, coalition parties can only receive one of nineteen possible shares (0.05, 0.10, ..., 0.95). If a party contributes, say, 8 per cent of this cabinet's total seat share, it cannot possibly receive the same proportion of portfolios; it must be either under-rewarded or over-rewarded.

A second source of concern is the possibility that there is also a degree of randomness in the translation of seats into portfolio shares. In other words, it may be the case that occasionally a party gets a portfolio too many, or too few, for its seat contribution, but the pattern is globally random: it does not systematically favour small parties, formateurs, or any other distinct category. Random error of this sort would not deflect the relationship from one-to-one proportionality by itself, but when combined with the need to translate seat shares to discrete portfolio shares, it may have this consequence. The reason is

³³ Browne and Franklin, 'Aspects of Coalition Payoffs', p. 460.

³⁴ Browne and Franklin's regression coefficients are not directly comparable to ours since they regressed seat contribution, which they took to be the predicted portfolio payoff under the proportionality hypothesis, on portfolio shares (the actual payoff). Based on their reported slope of 1.07 and R^2 of 0.855, the slope relating seat contribution to portfolio shares in their data would be 0.799, very similar to the one reported here.

³⁵ Browne and Frendreis, 'Allocating Coalition Payoffs by Conventional Norm', p. 758.

that the dependent variable, portfolio share, tends to be ‘flattened’: in our hypothetical cabinet of twenty portfolios, any portfolio share above 0.95 or below 0.05 must be collapsed to those values (otherwise the party would either be excluded from the cabinet or would form a single-party government). The effect of this flattening out of the relationship might be to raise the intercept above 0 and lower the slope below 1.

In order to assess the potential impact of these circumstances, we conducted a series of simulation experiments. Each simulation involved the generation of a hypothetical portfolio share variable that embodies both types of error and its regression on (actual) seat contributions. The construction of hypothetical portfolio shares for each simulation involved two steps. First, to capture the idea that portfolio shares are allocated proportionally but with some degree of random error, the hypothetical share for each party was set equal to its actual seat contribution plus a random error drawn from a normal distribution. A mean of zero was specified for the normal distribution to ensure that the errors are unbiased and the standard deviation was set at a value (0.1065) that results in a relationship between the final variable and seat contribution that closely matches the strength of the actual relationship. The second step involved the incorporation of a typical level of lumpiness. Since cabinets in our dataset average about twenty portfolios, this was achieved by collapsing the hypothetical party shares to the nearest 5 per cent. In other words, portfolio shares up to 0.07499 were set equal to 0.05 (one portfolio), shares between 0.075 and 0.12499 were set to 0.10 (two portfolios), and so forth.³⁶

The purpose in generating a series of hypothetical portfolio share variables in this fashion is to simulate situations in which the relationship between seat contributions and portfolio shares mimics the actual relationship observed in our data, but without the inclusion of any small-party bias. The outcomes of fifty such simulation experiments are summarized in Table 1. The results, as intended, show a mean level of explained variance (84.8 per cent) in the regressions that is very close to the actual value of 85.5 per cent. This suggests that the amount of error introduced in the simulation procedure is reasonable. The outcome, however, is striking: the regressions produce intercepts that are significantly greater than 0 and slopes that are significantly less than 1. In other words, the assumption that the actual relationship between the two variables is the result of some random deviations from one-to-one proportionality plus an inevitable need to allocate whole portfolios to single parties generates the same pattern that has been taken as signalling the existence of a small-party bias. Granted, the average deviation from proportionality in the simulations is smaller than we found with the real data, but some or all of this shortfall may be due

³⁶ A plausible alternative specification of the allocation process would be to collapse seat contributions to the nearest whole portfolio first, then introduce random error, then collapse again. This would simulate a process in which each party’s proportional share of portfolios is calculated first, then as a result of bargaining and other considerations, a final allocation emerges that may deviate in random ways from proportionality. Further simulations show that the conclusions of the simulation experiment do not change under this alternative specification.

TABLE 1 *Simulating the Effects of Random Error and 'Lumpiness' on the Proportionality Relationship*

Number of simulations	Mean level of explained variance (standard error)	Mean intercept (standard error)	Mean slope (standard error)	Mean slope when $a = 0$ (standard error)
50	0.848 (0.011)	0.026 (0.006)	0.951 (0.016)	1.003 (0.010)

Note: Entries summarize the results of regressing simulated versions of portfolio share on seat contribution. Portfolio share was simulated by setting it equal to seat contribution plus some random error drawn from a normal distribution with mean zero and standard deviation of 0.1065, then collapsing it into twenty discrete categories.

to limitations in the simulation procedure (for example, the assumption that all cabinets have twenty portfolios). Clearly, an intercept greater than 0 and a slope less than 1 are not, in and of themselves, proof of the existence of a small-party bias.

Since the slope and intercept coefficients themselves do not point unequivocally to the existence of non-proportionality in the allocation of portfolios, a better approach might be to see if the actual data pattern themselves in a manner that corresponds to the explanations that have been offered for non-proportionality. Browne and Franklin, for example, argued that the small-party bias is the result of a willingness on the part of coalition leaders or formateurs to take less than their proportional share in order to form coalitions they expect to dominate, which tend to be coalitions containing relatively few other parties.³⁷ They demonstrated the existence of this effect by running the regression separately on coalitions of different sizes; the results showed a clear tendency for the regression line to move closer to one-to-one proportionality as the size of the coalition increases.

The present data also show a tendency for the intercepts to move towards 0 and the slopes to move towards unity as coalition sizes increase from two parties to five or more parties.³⁸ The problem with the test is that it is very indirect: it does not show that the formateur is the one making the portfolio concessions. In addition, it is not clear that the premises of the explanation are correct. For instance, is it true that formateurs are more likely to dominate smaller coalitions than larger ones? After all, a formateur faced with multiple coalition partners may be in a better position to dominate than a formateur with just one coalition partner because it can play partners off against each other. One can also ask why

³⁷ Browne and Franklin, 'Aspects of Coalition Payoffs', p. 463.

³⁸ The estimated regression coefficients are $a = 0.196$ and $b = 0.608$ for two-party coalitions; $a = 0.109$ and $b = 0.671$ for three-party coalitions; $a = 0.055$ and $b = 0.780$ for four-party coalitions; and $a = 0.045$ and $b = 0.773$ for coalitions of five or more parties. All coefficients are significant at $p < 0.001$.

a dominant formateur should be obliged to over-compensate smaller parties – are they not better off in government than outside of it, regardless of the size of the portfolio payoff they receive?

Fortunately, Morelli's model, which also anticipates a small-party bias, suggests a much more focused test.³⁹ As noted earlier, in the extreme case of a three-party legislature where any two parties (but no single party) can form a majority government, the model prescribes equal payoffs to both government members, which means that the larger party (usually the formateur) will be under-rewarded. In our dataset, three systems have frequently exhibited party configurations in which a minimal winning government can only be formed by a coalition of any two from a set of three parties: Austria, Luxembourg, and (West) Germany. Of the forty-five governments in these countries that are included in the sample, some twenty-eight exhibit this basic configuration.⁴⁰ For the parties in these governments, however, the pattern of portfolio allocations is a very far cry from equal shares: the correlation between portfolio share and seat contribution is $r = 0.850$ ($n = 52$, $p < 0.001$).⁴¹

Thus, we have no clear evidence that the proportionality principle is systematically compromised by the existence of a modest small-party bias. As Table 1 illustrates, however, there is a way to get at the underlying proportionality if the deviations from it are the result of random error and lumpiness as modelled here. That way is simply to run the regression through the origin. In the simulations, this tactic produced a mean slope of 1.003 (s.e. = 0.010), which is not statistically different from unity.⁴² Eliminating the intercept makes sense from a theoretical perspective as well. The positive intercept in the original specification creates the expectation that the minimum allocation of portfolios, in this case an average of 8.1 per cent ($a = 0.081$), would go to parties with no legislative seats. This expectation is, of course, totally unrealistic: any such party would get no portfolios at all since it would not be offered cabinet membership. In other words, the estimated regression generates an impossible predicted value for parties that are not represented in the legislature. As Snedecor and Cochran note, in situations where 'the nature of

³⁹ Morelli, 'Demand Competition and Policy Compromise in Legislative Bargaining'.

⁴⁰ This does not necessarily mean that all legislative seats are held by three parties; only that the other parties (if any) are so small that their presence does not alter this basic bargaining logic.

⁴¹ This coefficient, moreover, may be too low. There is one highly deviant government here: the 1979 Werner government in Luxembourg for which WKB, but not Keesing's, give a very low portfolio allocation to the dominant Christian Socials. If this government is excluded, the correlation rises to $r = 0.953$ ($n = 50$, $p < 0.001$).

⁴² The reason is that forcing the line through the origin prevents the regression from responding to the 'flattening' of the relationship noted earlier. Because of this, these slopes are not especially sensitive to the size of the random error introduced in the simulations. In a further series of simulations, the standard deviation of the distribution from which the errors were drawn was increased nearly threefold to 0.3. This drastically lowered the strength of the relationship between hypothetical portfolio share and seat contribution (the mean R^2 fell to 0.401), but the mean slope increased just slightly to 1.007 (s.e. = 0.024), which is still statistically indistinguishable from unity.

TABLE 2 *Re-assessing the Connection between Seats and Portfolio Shares*

	Dependent variable			
	Model 1 Portfolio share	Model 2 Portfolio share	Model 3 Weighted portfolio share	Model 4 Weighted portfolio share
Independent variables				
Seat contribution	0.915 (0.009)	1.049 (0.016)	1.032 (0.017)	0.859 (0.018)
Formateur status \times Seat contribution	—	-0.182 (0.019)	-0.152 (0.019)	0.135 (0.021)

Note: Entries are unstandardized regression slopes, with intercepts set equal to zero and standard errors given in parentheses. All coefficients are significant at $p < 0.001$.

the variables Y and X makes it clear that when $X = 0$, Y must be zero', the appropriate procedure is normally to run the regression through the origin.⁴³

The results of the simulations indicate that the use of a regression model with this restriction allows us to capture the underlying proportionality of the relationship between seat contributions and portfolio shares, if it exists. There is no guarantee that it does, however. Remember that we have not shown that there is no small-party bias, only that the evidence provided by Browne and Franklin is equivocal and the limited test of Morelli's model is unresponsive. Now that we can differentiate systematic non-proportionality from the effects of random error and lumpiness, let us look again at the issue of proportional allocations.

The first model of Table 2 presents the results when portfolio share is regressed on seat contribution without an intercept: the slope coefficient rises

⁴³ George W. Snedecor and William G. Cochran, *Statistical Methods*, 8th edn (Ames: Iowa State University Press, 1989), p. 174. See also John O. Rawlings, Sastry G. Pantula, and David A. Dickey, *Applied Regression Analysis: A Research Tool*, 2nd edn (New York: Springer, 1998), p. 21. One could also argue that the model ought to be altered to preclude predicted values above unity, which are also theoretically impossible. The standard method for transforming variables that are bounded at 0 and 1 is to take the log-odds. Indeed, with both seat contribution and portfolio share transformed into log-odds, a regression slope of 1.0 would be produced in the case of perfect proportionality. Further simulations show, however, that the introduction of random error and lumpiness prior to transforming the dependent variable destroys this result, making it impossible to capture the underlying proportionality if deviations from it are entirely or partly the result of these factors. For this reason, we chose to stick with linear regression on the untransformed variables. It should be noted, however, that running the regression through the origin produces levels of explained variance that are often very misleading. The reason for this is that the original variance in the dependent variable is calculated from the origin, rather than from the dependent variable's mean. Unless that variable is centred around the origin, this tends to make the amount of explained variance excessively high. For this reason, we do not report R^2 values for any of these regressions.

to 0.915 (s.e. = 0.009). This coefficient is much closer to one-to-one proportionality, to be sure, but it still falls significantly short of that standard. Clearly, something else besides lumpiness and random error must be involved. A common factor in the hypotheses and models considered here is that deviations from one-to-one proportionality involve the formateur party in some way. Those who anticipate that small parties are favoured see the advantage of these parties as coming at the formateur's expense (for example, Browne and Franklin), while the opposite approach views the formateur as positioned to capture a disproportionately large portfolio payoff (for example, Baron and Ferejohn). What we must do, then, is to separate the formateur's portfolio share from those of the other coalition parties.

This can be achieved in the present context by creating and introducing into the regression model an interactive term between seat contribution and formateur status, the dichotomous indicator of whether a party is the government formateur. A significant negative coefficient for this interactive term would indicate that the slope for formateurs is significantly lower than that for other cabinet parties, suggesting that formateurs are relatively under-compensated in terms of the number of portfolios they receive. By the same token, a significant positive slope would indicate that formateurs are relatively over-rewarded in this regard.

Model 2 in Table 2 presents the results when the interactive term is added to the regression. Its estimated coefficient is both significant and negative, indicating a tendency towards formateur under-compensation. In fact, the size of this effect suggests a degree of under-compensation for formateurs that is substantially greater than Browne and Franklin found. They calculated that under-compensated parties receive, on average, a portfolio share that falls 8.3 per cent below proportionality;⁴⁴ Model 2 implies that the payoff formateurs receive falls short of their seat contribution by an average of 13.3 per cent (based on a mean seat contribution of 0.5734). With the interactive term capturing this effect, the seat contribution coefficient ($b = 1.049$) now shows the non-formateur parties to be slightly *over*-compensated relative to the proportionality norm. Indeed, these two effects off-set each other almost perfectly, producing an average slope across all parties of 0.989 – very close to one-to-one proportionality.⁴⁵

Introducing the Quality of Portfolios

The results thus far indicate that, apart from the effects of random errors and lumpiness, deviations from one-to-one proportionality between seat contributions and portfolio payoffs can be attributed to a tendency for formateurs to be

⁴⁴ Browne and Franklin, 'Aspects of Coalition Payoffs', p. 464.

⁴⁵ This figure is calculated by averaging the 407 non-formateur slopes of 1.049 and the 200 formateur slopes of 0.867 ($1.049 - 0.182$). Incidentally, further testing shows that a slope that falls slightly below 1.0 could simply be a consequence of the 'lumpiness' of portfolio allocations.

under-compensated and non-formateurs over-compensated. This would seem to support the hypothesis that formateurs do pay a price after all, which in turn would cast doubt on the models of Baron and Ferejohn and others that anticipate formateur over-compensation. In reality, however, the picture is not so clear. While formateurs evidently are under-compensated in terms of the numbers of portfolios they receive, there is no reason why portfolio payoffs should be evaluated only in those terms. Indeed, as noted earlier, it is highly implausible that parties would rate all portfolios equally and simply count the number that each party receives. In particular, since the formateur almost always gets the prime ministership (95 per cent of them do in our sample), it is quite possible that any numerical shortfall is made up – perhaps more than made up – by the inherent value of this portfolio along with the others it receives. In assessing the nature of portfolio payoffs, we must take account of quality as well as quantity.

The introduction of considerations of portfolio quality or value involves the use of the portfolio assessments produced from Laver and Hunt's expert survey. We have seen that these assessments are less than ideal for this purpose, primarily because they are incomplete and ordinal rather than cardinal in nature. The potential for error in the method we have developed for translating them into weights is correspondingly large. It is striking, therefore, to find that this potential does not appear to have been realized to any significant degree. In fact, the correlation between the original portfolio share variable and its weighted version is virtually perfect ($r = 0.986$). Because of this close correspondence between weighted and unweighted portfolio shares, substituting the weighted measure for its unweighted version scarcely affects the regression results, as the third model of Table 2 shows.

Why should the inclusion of portfolio weights change so little? One possible answer is that the weights have little effect because they are all approximately equal. This possibility can be easily dismissed. The weighting scheme gives the top portfolio, the prime ministership, a weight that is, on average, more than twice (2.41 times) that of the lowest or baseline portfolio. This is not a function of any special weighting given the prime ministership itself; excluding that post leaves the range of portfolio weights virtually the same (the average ratio of highest weight to lowest becomes 2.26). Clearly, the weights are far from being only trivially different from one another. A more likely explanation is that parties are receiving distributions of portfolios that are reasonably balanced with respect to salience. In other words, proportionality prevails even when the weights are applied because parties are getting their proportional shares of portfolios of different values; apparently, certain parties are not being fobbed off with portfolios of lesser value so that other parties can scoop up the more valuable ones.

This is not the end of the story, however. The overall one-to-one proportionality, which also characterizes Model 3, is made up of a tendency for formateurs to receive less than a proportional share, while the other parties receive slightly more. But consider these two groups of parties separately. While overall proportionality requires that one group must be over-rewarded if the

other group is under-rewarded, one would still expect proportional payoffs to be in evidence when we consider only one group. In other words, non-formateurs may be over-compensated relative to formateurs, but they should still receive proportional shares vis-à-vis one another (on average) if proportionality is the operative principle.

This turns out to be the case for the non-formateur parties. If we consider just these parties and create new versions of the two variables so that the formateur's part in each is removed, we find that one-to-one proportionality prevails: the slope coefficient becomes 0.987 (s.e. = 0.008), which is statistically indistinguishable from unity.⁴⁶ Among non-formateur parties, then, the weighting scheme developed here generates a near-perfect one-to-one correspondence between seat contributions and weighted portfolio shares. Whether the value or simply the quantity of portfolios is considered, relative payoffs among these parties seem to follow very closely the proportionality principle.

The same is not true of formateurs, however: the slope for formateurs in Model 3 ($b = 1.032 - 0.152 = 0.880$) is clearly well below unity. Is this an accurate portrayal of the payoffs for formateurs or is it simply the result of some deficiency in the weighting scheme? Note that if there is a deficiency in the weighting scheme, it is unlikely to apply across the board; the fact that it produces near-perfect proportionality for non-formateur parties suggests that it is reasonably accurate. But there is one portfolio that non-formateur parties almost never hold – the prime ministership – and the analysis of non-formateur proportionality can tell us little about it. It may be, then, that the formateur under-compensation evident in Model 3 is the result of nothing more than an under-valuation of the post of prime minister relative to other portfolios.

The possibility that the prime ministership is under-weighted in our scheme is by no means far-fetched; after all, apart from making it the top post, we gave it no special consideration. The consequence is that, on average, the weighting scheme allocates it only 6.4 per cent more weight than the next ranked portfolio (usually the deputy prime ministership) and makes it just 59.9 per cent more valuable than the average non-prime ministerial portfolio. It is highly probable that the prime ministership is worth a good deal more than this, but how much more? The available data cannot answer this question in any definitive sense, but they do allow us to explore a particularly interesting scenario, that in which the slope for formateurs who hold the prime ministership is exactly equal to unity.

As it happens, the slope for the formateur regression increases steadily as the weight for the prime ministership is increased. If the weight is increased by a factor of 3.65, a slope of 1.000 is reached. With this adjustment, we now have

⁴⁶ 'Weighted portfolio share' becomes the party's weighted proportion of cabinet portfolios, excluding those held by the formateur, and 'seat contribution' becomes the party's proportion of cabinet-held legislative seats, excluding those contributed by the formateur. There is a risk that this relationship is inflated by parties in two-party coalitions, since eliminating the formateur leaves one party with 100 per cent of both seats and portfolios. Excluding these cases, however, still leaves the slope very close to unity ($b = 0.968$, s.e. = 0.015)

a weighting scheme that produces slopes implying perfect or near-perfect one-to-one proportionality among non-formateurs and among formateurs. This does not mean, however, that it produces this result across the board. In fact, as Model 4 in Table 2 shows, the revised weights actually produce a very different outcome: formateurs are now receiving a significantly *larger* share of portfolio payoffs than other cabinet parties. Increasing the value of the prime ministership (which non-formateurs almost never hold) has altered the balance of payoffs between non-formateurs and formateurs decisively in favour of the latter.

Model 4, by showing formateurs capturing more than their proportional share of portfolio benefits, presents a very different picture from that gleaned from the examination of portfolio numbers. We must interpret this outcome with caution, however, since it was achieved entirely on the basis of a large increase in the assumed value of the prime ministership. As noted earlier, this appears to be the *only* way in which formateurs can be shown to be advantaged in terms of portfolio payoffs. In fact, further investigation indicates that the non-prime ministerial portfolios held by formateurs are distributed across the rankings in very much the same fashion as the portfolios held by non-formateurs, making it very unlikely that any transformation of the weights of these other portfolios (such as by increasing the weights at the top end of the scale at the expense of those at the bottom end or vice versa) would have affected the relative payoffs of formateurs and non-formateurs.⁴⁷ The focus on the value of the prime ministership is thus necessary, but it means that this post has been allocated levels of salience that are, on average, 5.83 times those of the typical non-prime ministerial portfolio. Is this a realistic assessment or have we now erred in the opposite direction by assigning too much value to the top post in the cabinet?

Clearly, this cannot be determined with the data at hand, but we can gain an indication of the plausibility of the basic result portrayed in Model 4 by determining the minimum valuation that has to be put on the prime ministership (given the weights attributed to the other portfolios) in order to yield formateur over-compensation. It turns out that increasing the weight for the prime ministership by a factor of 2.62 is sufficient to produce a statistically significant advantage in terms of portfolio payoffs for the formateur. With this adjustment, the prime ministership becomes 4.19 times more valuable than a typical portfolio, on average. In other words, with the present data and assuming the weights for the other portfolios are correct, any valuation for the prime ministership that makes it worth more than about four typical non-prime ministerial portfolios will generate over-compensation for the formateur.⁴⁸

⁴⁷ For example, the mean non-prime ministerial portfolio held by formateurs is valued at 5.8 per cent of the total portfolio payoff; for non-formateurs, the corresponding value is 6.0 per cent. This difference is not statistically significant ($F = 0.667$, $p = 0.414$).

⁴⁸ Conversely, a weighting for the prime ministership that makes it no more than 3.05 times as valuable as a typical non-prime ministerial portfolio would produce formateur under-compensation. Values in-between produce formateur proportionality.

Thus, the prime ministership must be very valuable indeed before the present data can show formateur over-compensation.

Another indication of the plausibility of formateur advantage hypothesis can be gleaned from the amount of over-compensation that accrues to formateurs. This turns out to be remarkably slight. For example, if the prime ministership were worth 4.19 typical portfolios (the minimum value needed to produce over-compensation), the average formateur party can expect to receive 58.2 per cent of the total portfolio payoff on the basis of a seat contribution of 57.3 per cent – an over-compensation of about 1.5 per cent. Even at the much higher valuation of 5.83 non-prime ministerial portfolios, the over-compensation would average just 5 per cent. This seems very little for so high a valuation put on the prime ministership. Although neither the finding that the prime ministership requires a very large weight to show any formateur advantage nor the finding that even large re-weightings produce only minimal levels of advantage should be seen as conclusive, given the possibilities for inaccuracy in our estimates of portfolio salience, both tend to discredit the notion that formateurs occupy and exploit a privileged position in the coalition formation game.

DISCUSSION

In parliamentary systems, the allocation of portfolios within coalitions critically affects the ability of member-parties to reward supporters with patronage and ‘pork’, and it has at least some impact on the direction that public policy takes. For many years, it appeared as if Browne and Franklin had satisfactorily disposed of this very consequential element of the coalition game. Indeed, one would be hard pressed to find another social science relationship as strong as their finding of one-to-one proportionality between seat contributions and portfolio shares, and the little that was left unaccounted for seemed attributable to nothing other than a slight small-party bias. So impressive were these results that they have gone virtually unquestioned for nearly thirty years, despite the fact that Baron and Ferejohn’s widely applied bargaining model predicts what amounts to a large-party bias in the form of formateur over-compensation. To the best of our knowledge, only Morelli has attempted to reconcile theory with data and his reconciliation takes the form of a bargaining model that produces results consistent with Browne and Franklin’s findings.

In this article, we used a new and much larger dataset to re-investigate the evidentiary side of this debate. We found, for one thing, that Browne and Franklin’s results can be reproduced very accurately with these data. Whether their analysis revealed a small-party bias was unclear, however: simulations indicated that a regression intercept above 0 and a slope below 1 do not necessarily indicate its presence. Moreover, Morelli’s model produces incorrect predictions concerning that bias in situations where the model expects it to be at its strongest. Notwithstanding these concerns, the use of what we see as a more appropriate method (regression without an intercept) has enabled us to produce strong evidence not only for overall proportionality in the quantitative allocation

of portfolios but also for a substantial formateur under-compensation (and thus a small-party bias).

The key assumption in this analysis is, of course, that the proportion of portfolios each party receives can be equated with its payoff – in other words, that the obvious differences in the salience or importance of portfolios are not relevant. Many observers have noted the vulnerability of this assumption and our main goal in this investigation has been to find a way to dispense with it by introducing portfolio salience directly into the analysis. We found that the rank orderings of key portfolios produced by Laver and Hunt's expert survey, despite their weaknesses, can be adapted for this purpose. We also found that the weighted portfolio shares produced with the aid of this information are extremely similar to their original or unweighted values, indicating that portfolio allocations are balanced with respect to quality as well as quantity. If the portfolio weights as initially formulated are reasonably correct, the conclusion that overall proportionality prevails in the allocation of portfolio benefits (albeit with some formateur under-compensation) continues to hold.

The 'fly in the ointment' here is the weight put on the post of prime minister: it clearly may stand farther above the other portfolios than our weighting scheme has placed it. Indeed, simply adjusting the value of the prime ministership to produce a slope of unity for formateurs who hold the prime ministership generates significant formateur over-compensation relative to other parties, consistent with the models that see formateurs as agenda-setters. Under this scenario, non-formateurs still receive proportional shares relative to other non-formateurs, but the increase in value of the prime ministership means that their payoffs are not as lucrative as those of formateurs.

What prevents us from endorsing this scenario is that it entails a very large increase in the relative value of the prime ministership: it becomes nearly six times as valuable as the typical other portfolio. In fact, it would have to be more than four times as valuable just to produce any significant over-compensation at all, and any weight that leaves it no more than three times as valuable would yield significant formateur under-compensation. Moreover, it is only by making such large increases in the valuation of the prime ministership that formateurs can be shown to be anything other than under-compensated; the other portfolios that formateurs receive are, on average, no more valuable (and they are less numerous) than the portfolios allocated to non-formateurs. Finally, even very large increases in the value of the prime ministership produce surprisingly little over-compensation. Setting its value at about six average portfolios, for example, yields only a 5 per cent over-compensation for formateurs.

These points must call into question the application of proposer models such as Baron and Ferejohn's to coalition bargaining situations.⁴⁹ By assuming an

⁴⁹ Experimental evidence has also called it into question. Daniel Diermeier and Rebecca Morton, 'Proportionality versus Perfectness: Experiments in Majoritarian Bargaining' (unpublished paper, 2000), for example, ran experiments to test various implications of the Baron and Ferejohn model and found little support for the model. Most notably, they found that proposers failed to take full advantage of their proposal power.

exogenous order of play where the formateur offers an all or nothing proposal to the other parties, the models appear to predict too much power for the formateur. Morelli's model is based upon a very different framework in which the order of play is endogenous with the formateur choosing the sequence in which parties make their own payoff demands; the result is that the formateur is not privileged and payoffs are expected to be approximately proportional. While one aspect of this model did not find empirical support here, the test was certainly too limited for us to conclude that the sequential-demand framework is inappropriate.

The choice of an appropriate set of assumptions to model the coalition formation process in West European parliamentary systems therefore remains unresolved; all we can say at this point is that the available evidence on portfolio quality, sketchy though it is, points to the same basic conclusion that the evidence on portfolio quantity indicated. That conclusion is that portfolio payoffs are allocated in a roughly proportional manner, with perhaps some degree of formateur under-compensation. Whatever structures the coalition game, it does not appear to be the agenda-setting power of formateurs.