

Bargaining and Bicameralism*

Giri Parameswaran[†]

April 26, 2017

Abstract

In bicameral legislatures, the protection of small states often motivates malapportioning in the upper house. Using a legislative bargaining model, I show that malapportionment may produce the opposite effect. Under unicameralism, same-state legislators are shown to not inherently be coordinated to cooperate, diminishing the fear of a big state conspiracy. By contrast, under bicameralism, preference complementarities enable upper house legislators to effectively coordinate their state delegations, and this skews the expected allocation in favor of big states. Hence, unless bicameralism significantly increases their agenda power, small states will fare even worse under bicameralism whenever they are disadvantaged under unicameralism.

*I am grateful to John Londregan, Simon Mongey and Wolfgang Pesendorfer, the editor and 3 anonymous referees, and to seminar participants at Princeton and the University of Western Australia, for their helpful comments, and to Josh Nadel and Caleb Wroblewski for their able research assistance.

[†]Assistant Professor, Department of Economics, Haverford College, 370 Lancaster Ave, Haverford PA 19041; email: gparames@haverford.edu

1 Introduction

"The equality of representation in the Senate.... being evidently the result of compromise between the opposite pretensions of the large and the small States...
" "A government founded on principles more consonant to the wishes of the larger States, is not likely to be obtained from the smaller States." — James Madison, Federalist no. 62

Bicameralism is an institutional feature common to many legislatures around the world. In almost all bicameral systems, legislators in the lower house represent districts of roughly equal size by population (Tsebelis and Money, 1997), reflecting the principle of equal representation. By contrast, representation in the upper house is often intentionally malapportioned to benefit particular constituencies.¹ For example, a malapportioned upper house that gives equal representation to states or provinces has the obvious effect of over-representing regions with smaller populations.

In the United States, the over-representation of small states in the upper house resulted from a compromise to the small states, who feared that their interests would be ignored in the popular chamber where a few large states could command majorities between themselves. A similar compromise, to entice smaller states to join the federation, informed the design of the Australian Senate. The logic seems straightforward: over-representation of small states in the upper house protects them against usurpation of the policy agenda by the larger states. Implicit in the argument's logic is the expectation that state delegations in a unicameral legislature will likely vote as a coordinated unit. This, in turn, requires that the interests of legislators from the same state sufficiently overlap, so that a policy that is held to be desirable by some legislators from a given state, will likely be held desirable by most or all legislators from that state.

But this is a strong assumption. Policies may, after all, benefit voters in some parts of a state but not others. A policy to finance public transit infrastructure, for example, will likely be supported by legislators representing districts in New York City, but is unlikely to be of much value to legislators representing districts upstate. The interests of people in different

electoral districts need not coincide simply because those districts are located within the same state. (Thorpe (2010) finds, for example, that geographic disparities in the allocation of defense contracts are tied to local, rather than state wide, economic conditions.) Over a variety of policy issues, local interests may well be most salient in determining legislator behavior.

This paper investigates the claim that bicameralism provides an advantage to small states. I assume that legislators are concerned with the interests of their immediate constituency, so that parochial and state interests do not necessarily coincide. I show that, contrary to conventional wisdom, bicameralism does not necessarily improve the welfare of small states, and may have the unintended consequence of harming them. Moreover, bicameralism will tend to worsen outcomes for small states precisely when those states would already have been disadvantaged under unicameralism. Hence, the case for bicameralism as a check against policy usurpation by big states, may be mistaken, and warrants further examination.

To study this question, I adapt the bargaining framework in Baron and Ferejohn (1989) to model decision making in a bicameral legislature. The legislature must allocate a fixed surplus amongst various districts. Each district is contained within a state, and small states contain fewer districts than big states. Representation in the lower house is by district, whilst representation in the upper house is by state. Lower house legislators seek to maximize the allocation to their district alone, whilst upper house legislators seek to maximize the total allocation to the districts within their state. If recognized, a legislator will propose an allocation that maximizes funding to her constituency whilst earning the support of majorities in both chambers. In equilibrium, the proposal will target districts whose representatives are ‘cheap’ — i.e. who will demand fewer concessions before supporting a proposal. The price of legislators in turn depends on the size of the allocation they expect to receive. Legislators who rationally expect to receive more, will be more demanding than those who expect to receive less. Since allocating funds to a given district benefits both the legislator representing that district and the legislator representing the associated state, there are complementari-

ties between the preferences of lower and upper house legislators from the same state. This complementarity will be crucial to the analysis that follows.

The main insight of this paper is as follows: Under unicameralism, whilst it is possible for big state legislators to privilege one another when forming coalitions (to the detriment of small states), such behavior would often not be in their best interest. Indeed, a big state legislator may find it more rewarding to collaborate with small state legislators, especially if they are cheaper. Hence, under unicameralism, legislators from the same state are not inherently coordinated to work together; they are not natural coalition partners. Additionally, if the concern of policy usurpation by big states is well founded, then it must be that small state legislators are cheaper coalition partners, since they would be expected to gain less in the bargaining process. Hence, the fear that big states will conspire against small states is least compelling precisely in those cases when small states are expected to fare poorly.

Now, consider a bicameral legislature. The incentives for lower house legislators are analogous to the unicameral case. However, the incentives for upper house legislators differ markedly, since they are concerned with the total allocation across their state. (This is consistent with the empirical findings in Lee (2004) and Lazarus and Steigerwalt (2009), that the state wide interests of upper house legislators differ from the parochial interests of lower house legislators from the same state.) When recognized, an upper house legislator will allocate funds across each of the districts within his state. Given the complementarity in preferences, this naturally earns the support of the lower house legislators from his state, resulting in coordinated support for policies amongst state delegations. Furthermore, this coordinated support from the proposer's delegation reduces the number of out-of-state legislators upon whom the proposer must expend resources to build a winning coalition. Since big state delegations are larger than small state delegations, they are then able to retain more resources within their own state, thus skewing the distribution in their favor. I refer to this as the *reduced requirement effect*. By introducing state-level representatives, bicameralism has the unintended consequence of creating incentives for precisely the sort of coordinated

behavior that generated the concern of policy usurpation by big states under unicameralism.

Of course, by requiring that proposals receive majority support in the upper house, where small state legislators are more numerous, bicameralism may have the additional effect of drawing more small states into coalitions than would otherwise be the case. I refer to this as the *standard effect*. The net effect of bicameralism may thus be ambiguous — the reduced requirement effect tends to work against small states, whilst the standard effect often helps. Which of these effects dominates will depend on the relative ‘prices’ of different types of legislators. If small state legislators are relatively expensive, then the standard effect will dominate. Building an upper house majority requires cooperating with small state legislators, even though they are more expensive, and would have been excluded from unicameral coalitions. By contrast, if small state legislators are relatively cheap, then the reduced requirement effect will dominate. Small state legislators would have already been well represented in the unicameral governing coalition, and the coordination by big states under bicameralism, decreases the number who are needed. Recall that small state legislators will be relatively cheap when they expect to fare worse than big states in the bargaining game. Hence, bicameralism tends to hurt small states in precisely those cases when they would already have been disadvantaged under unicameralism.

Bicameralism may also affect equilibrium policy outcomes by changing the likelihood that the agenda setter is from a big state. Proposition 3 decomposes the overall effect of bicameralism into a coalition composition channel and an agenda setting channel. The proposition shows that, if small states are likely to be disfavored under unicameralism, then they will continue to be disfavored under bicameralism, unless bicameralism sufficiently skews agenda control in their favor.

To be clear, this paper does not claim that bicameralism always disadvantages small states. For example, the model’s predictions are perfectly consistent with the empirical result in Hauk, Wacziarg et al. (2007), which finds no difference in per-capita spending between big and small state districts in appropriations bills originating in the U.S. House of Repre-

sentatives, but a significant advantage to small states in conference committee bills. Rather, this paper makes the conditional claim that if small states were likely to be disadvantaged under unicameralism, then they would likely be even more so disadvantaged under bicameralism. Hence, the case for malapportionment in the upper house, to remedy perceived inequalities between big and small states, may not be well founded.

The characterization of optimal coalitions in a multicameral framework with imperfectly correlated preferences poses significant challenges and is generically intractable. The usual procedure of targeting policy towards the cheapest legislators breaks down when legislators who are relatively cheap in one chamber turn out to be expensive in the other chamber. This paper, by focusing on the simplest bicameral framework that includes imperfectly correlated preferences, provides an explicit characterization of these optimal coalitions. This is the first paper to fully characterize bargaining outcomes under these conditions.

Bicameralism, as a feature of legislative institutions, has generated significant interest in the recent literature. Tsebelis and Money (1997) and Cutrone and McCarty (2006) provide an extensive summary of the existing literature. The model of bicameralism used in this paper draws on the work of McCarty (2000), Ansolabehere, Snyder and Ting (2003), and in particular, Kalandrakis (2004). Kalandrakis presents a model in which there are two types of states — big and small — who send delegations to both houses of the legislature. The government determines an allocation of funds to each state. Legislators represent their state (rather than districts within the state) and so the preferences of all legislators from the same state will be perfectly aligned. Hence, state delegations are guaranteed to vote *en bloc*. Whilst Kalandrakis does not explicitly address the unicameral versus bicameral question, his model setup implicitly coordinates big state legislators in such a way that would cause unicameralism to disadvantage small states. By contrast, in this paper, lower house legislators represent individual districts, rather than the state as a whole. This is important for two reasons. First, as I have argued, recognizing differences in policy interests between legislators from the same state is arguably more reasonable over a range of policy contexts.² Indeed,

much of the empirical literature (see Levitt and Snyder Jr (1995), Heitshusen (2001), Lee (2004), Knight (2005), Lazarus and Steigerwalt (2009), amongst many others) has focused on district level allocations. Second, distinct preferences makes the assumption of coordinated behavior — and the associated concern of big state policy usurpation — far less compelling.

Ansolabehere, Snyder and Ting (2003) present a model in which lower house legislators do represent individual districts. However, they assume that upper house legislators represent the median district in their state, rather than the state as a whole. This approach effectively imposes an arbitrary decision rule on the behavior of upper house legislators, rather than determining their optimal choice. In fact, the chosen decision rule is inconsistent with optimal behavior — it requires the legislator to support proposals that allocate strictly less funds to the state than could be expected in the continuation play. The legislator will support proposals that his constituents, on average, dislike. By contrast, in assuming that upper house legislators are motivated by the gross allocation to their state, my model allows for optimal legislator behavior, rationalized by consistent policy preferences.

Two other models of bicameralism study the effect of preference complementarities in different contexts. Chen (2010) introduces complementarities by considering a polity in which lower and upper house legislators represent geographically overlapping constituencies. In that paper, lower house districts need not be located wholly within upper house districts, and so the application to federal bicameral systems is less compelling. Shepsle et al. (2009) study the effect of the senatorial electoral cycle on the timing of targeted appropriations when credit is shared between upper and lower house legislators. Theirs is a dynamic model that is concerned with cyclicity of appropriations, independent of the big versus small state question.

The model is distinguished from other bargaining models of bicameralism, such as McCarty (2000) and Diermeier and Myerson (1999), that do not account for the complementarities in legislator preferences. Diermeier and Myerson (1999), for example, consider a lobby model in which legislators' preferences are uncorrelated and depend only on the size

of the bribe they individually receive. In contrast to distributional policies, Hammond and Miller (1987) consider a spatial model of bicameralism, although theirs is not a bargaining approach. They demonstrate that, unlike unicameral legislatures, bicameralism may generate stable policies, provided that upper and lower house agents have preferences that are sufficiently distinct. This assumption stands in contrast to this paper, which assumes that preferences are positively correlated. Hammond and Miller’s analysis predicts a range of possible bicameral policies, all of which are intermediate to the ideal policies of each chamber. Heller (1997) uses Nash bargaining to select a unique policy from amongst these. By contrast, this paper shows that bicameralism need not moderate outcomes, and may indeed produce more skewed outcomes that would be generated by either house acting unilaterally.

This paper proceeds as follows. Section 2 outlines the model of bicameralism. Sections 3 and 4 present the intuition for the equilibrium through the exploration of a simple numerical example. Section 5 presents a comparison of unicameralism and bicameralism. Section 6 concludes. The formal analysis of the model, and associated proofs can be found in the Appendix.

2 The Model

The formal model is adapted from Kalandrakis (2004). There is a polity that is divided into geographical regions, or states. There are $b \geq 1$ big states and $s \geq 1$ small states. The polity is also divided into electoral districts. For simplicity, I assume that each small state contains just one electoral district (i.e. the state and district coincide), whilst each big state contains $k > 1$ districts. Hence, there are $s + kb$ districts located within $s + b$ states.

The government must determine the allocation of resources (e.g. funding for highways or other local public goods) amongst the various districts in the polity. Congressional districts are assumed to be the finest level at which legislators may direct resources. The total supply of government resources is normalized to 1. A policy is a vector $x = (x_1, \dots, x_{s+bk})$, where x_i

is the share of the pie that is received by the i^{th} district.

Government policy is determined by a legislature comprised of two chambers: an upper and lower house. In the upper house, each state is represented by one legislator, whilst in the lower house, each district is represented by one legislator. The procedure by which the legislature adopts policy is based on the framework of Baron and Ferejohn (1989). In a given bargaining period, a member of the legislature is chosen at random to propose a division of the ‘pie’. Once the proposal is made, the legislators in both chamber simultaneously vote to either accept or reject the proposal. The proposal is accepted if at least $M_L > \frac{s+kb}{2}$ legislators in the lower house and at least $M_U > \frac{s+b}{2}$ legislators in the upper house vote to accept the proposal. (The setup, therefore, allows for super-majority rules in either chamber.) If the proposal is accepted, then the allocation is implemented and the game ends. If the proposal is rejected, then the legislature adjourns and reconvenes in the following period, when the above procedure is repeated. This process continues until a proposal is accepted. Whenever a comparison is made to unicameralism, the unicameral legislature is assumed to be identical to the lower house of the bicameral legislature.

Since the policy space is purely distributional, each legislator’s goal is to direct as large a slice of the pie as possible to his or her constituency. Formally, $u_i^L(x) = x_i$ is the utility of the lower house legislator representing the i^{th} district. Similarly $u_j^U(x) = \sum_i x_i$ is the utility of the upper house legislator from the j^{th} state, where the summation is over the set of districts within that state.³ All agents share a common discount factor $\delta \in [0, 1]$.

This specification of preferences is a significant point of departure from Kalandrakis (2004). That paper assumes that allocations are made to states rather than districts, and so all legislators from the same state have identical preferences over allocations, independent of the chamber in which they reside. The result is that state delegations (in both chambers) vote *en bloc*. By contrast, since this paper assumes that resources are targeted at the district level, state delegations need no longer be unanimous in their assessment of different proposals. Indeed, lower house legislators from the same state are motivated purely by

attracting resources to their district, independent of the allocation to other districts within their state. However, there is a complementarity between the preferences of lower and upper house legislators from the same state. Since upper house legislators seek to maximize the total resources accruing to districts within their state, an allocation that improves outcomes for a given district, also improves outcomes for that district's state, *ceteris paribus*. In this model, since in small states, the state is the district, the preferences of lower and upper house agents perfectly coincide. By contrast, in big states, where there are many districts, the preferences of lower and upper house legislator only partially correlate. The complementarity in legislator preferences is crucial to the analysis that follows.

Denote the set of legislator types by $T = \{S^L, S^U, B^L, B^U\}$, where S^L refers to a lower house legislator from a small state, and the other types are similarly defined. Let $p_t \geq 0$ be the probability that a type- t legislator is recognized as the proposer.⁴ These probabilities must satisfy: $b(kp_{B^L} + p_{B^U}) + s(p_{S^L} + p_{S^U}) = 1$. (When comparing to unicameralism, there are two types $T \in \{S, B\}$ and the probabilities satisfy $bkp_B + sp_S = 1$.) The setup places no restriction on the allocation of proposal power between chambers or between large and small states. For example, recognition probabilities are not assumed to be independent across chambers, and a feature of bicameralism may be to skew proposal power towards small states.

The goal of this paper is to analyze the effect of bicameralism on big versus small states. Accordingly, I abstract from features that may generate within-type differences (i.e. differences in outcomes between big state districts or between small state districts), in order to focus on between-type differences. Such abstractions are analogous to controlling for factors that may be salient, but which are unrelated to the effect of size. For example, implicit in the type-dependent recognition probabilities above, is the assumption that all legislators of a given type (e.g. all lower house legislators from big states) are identical in their likelihood of being recognized as the proposer. Whilst within-type differences in proposal power may exist, such differences cannot explain aggregate differences in outcomes *between* big and

small states. For similar reasons, I restrict attention to strategies and equilibria that are stationary and symmetric. This requires that, whilst legislators may distinguish between districts or legislators from differently sized states, they cannot arbitrarily distinguish districts or legislators from states of the same size. Symmetry ensures that the equilibrium outcomes generated by this model are a consequence of differences in size, rather than some other arbitrary or unmodeled factor.

Pure strategy equilibria do not always exist in this game, and so agents may be required to play mixed strategies. Let $\mu \in \Delta(X)$ be an assignment of probabilities over the set of feasible allocations. A stationary, symmetric strategy for a type- t legislator is an assignment of probabilities μ_t over feasible allocations whenever they are the proposer, and a decision rule $a_t : X \rightarrow \{0, 1\}$ which indicates whether they will accept or reject a given proposal when they are not the proposer. The stationarity assumption requires that agents choose the same action in every structurally equivalent sub-game, which amounts to asserting that strategies are history independent. A stationary, symmetric, sub-game perfect equilibrium is a set of strategies $\{(\mu_t, a_t)_{t \in T}\}$, such that, for each type $t \in T$, there is no other strategy (μ'_t, a'_t) which gives some legislator strictly higher utility, given the strategies of all other players.

I make two additional refinements which are standard in the literature. First, I restrict attention to equilibria in which the agents' decision rules satisfy the weak dominance property. This requires that an agent accept a proposal x , only if she weakly prefers the outcome under that policy to her expected payoff if the proposal were rejected. Weak dominance requires that agents behave as though they were pivotal. This rules out perverse equilibria in which agents accept (or reject) all proposals, independent of their preference, because they do not expect their choice to affect the outcome. Second, where necessary, I restrict attention to no-delay equilibria. Banks and Duggan (2000),(2006) show that no-delay equilibria exist in a general bargaining environment that embeds this model.

Before moving to the analysis, I briefly discuss some of the modeling choices underlying this model. First, the model focuses on purely distributive policies. This in itself is not

unreasonable — a significant function of government is to allocate funding and resources — although it admittedly narrows the scope of the analysis. Several factors motivate this choice. Chief amongst these is that the measurement of welfare is simplest in a distributive policy space — it suffices to simply compare per-capita spending across districts. Given the ease of measurement, the distributional policy space is also the natural context to empirically study the consequences of bicameralism. Indeed, Lee (2000), Lee (2004), Hauk, Wacziarg et al. (2007), Knight (2008), Lazarus and Steigerwalt (2009), amongst many others, all examine the distributional consequences of bicameralism by measuring funding for various local public goods in appropriations bills. By contrast, suppose we considered a spatial policy. To determine the policy success of each legislator, we must now compare their ideal policy (which, empirically, is not observed) to the equilibrium policy. The overall implication for big and small states will thus depend on the arrangement of the ideal policies of legislators from big and small states.⁵ Unless there are strong reasons, *a priori*, to believe that the ideal policies of legislators from big states systematically differ from those of small state legislators in a known way, it becomes difficult to draw general conclusions about the impact of institutions on the outcomes for big and small states.

The distributive policy space also enables the simplest specification of legislator preferences and of complementarities across chambers (since money outcomes are additive). By contrast, the specification of preferences and the connection between the preferences of lower house and upper house legislators is made much more complicated in a spatial model.⁶ As we will see below, complementarities play a crucial role in predicting how bicameral policies will differ from unicameral ones. Absent these complementarities, the sense in which upper house legislators broadly represent the aggregate interest of voters across their state becomes lost.

Second, the model assumes that legislators are unconcerned with the allocation to constituencies other than their own, thus ignoring the possibility of spillovers. Even if the interests of disparate districts from the same state are unconnected, we might expect the

interests of neighboring districts to correlate. That this is true is not disputed. As before, the simplifying assumption is made for tractability. Furthermore, externalities are typically not contained within state boundaries — spillovers will also flow between neighboring districts separated by state lines. Hence, the inclusion of externalities may be just as likely to spur between-state cooperation as within-state cooperation. Moreover, the results will depend crucially upon the specific geographic arrangement of districts and states, and this will again obfuscate the identification of general insights into the effect of bicameralism on big versus small states.

Third, although the model of the legislature is consistent with many models of bicameralism (including Ansolabehere, Snyder and Ting (2003), Banks and Duggan (2000), Kalandrakis (2004), and McCarty (2000)), it nevertheless simplifies the bargaining process in a variety of ways. For example, the model simplifies the bargaining dynamic by assuming that offers continue to be made *ad nauseum*, until a resolution is achieved; navette continues until a proposal is accepted. Tsebelis and Money (1997) note that navette is the most common method of resolving inter-cameral disputes, and as such, I take this as a reasonable modeling choice. Nevertheless, I acknowledge that alternative dispute resolution mechanisms (such as a conference committee, or dissolution of the legislature) are used in some legislatures. In doing so, I note the following: First, the Baron and Ferejohn framework can be easily modified to reflect these alternative procedures, and doing so will not change the inherent logic of the model. The continuing navette assumption is made, because the model requires that *some* procedure be specified, and amongst the various alternatives, it is simplest. Second, Kalandrakis (2004) discusses conditions under which the *ad nauseum* navette assumption can be mapped onto a model with a conference committee.

Notwithstanding its starkness, there is considerable evidence (see Wilson (1986) and Knight (2005)) that the Baron and Ferejohn model captures important features of the legislative bargaining dynamic. Furthermore, as will become clear, the channel that this paper highlights focuses purely on the composition of optimal coalitions. The crucial assumption

is that proposers will build coalitions comprised of the cheapest of legislators. (Lee (2000) finds empirical support for this assumption.) The remaining assumptions (e.g. regarding the recognition rule, and the nature of the continuation game) are important in determining the ‘price’ of legislators, and these prices may differ under different specifications. Nevertheless, the insights regarding the composition of optimal coalitions will remain true, no matter how the legislator prices are determined. Hence, the process highlighted in this paper will be robust to different specifications of the bargaining protocol, making any given modeling choice relatively benign.

3 Optimal Coalitions & Reduced Requirement Effect

In this section, I provide an informal account of equilibrium in the bicameral legislature by solving a particular numerical example. The interested reader can find a formal presentation of the general results in the Appendix.

Suppose there are four small states ($s = 4$) each containing only one district, and one big state ($b = 1$) containing five districts ($k = 5$). Each district has one representative in the lower house and each state has one representative in the upper house. Accordingly, there are nine lower house, and five upper house legislators. A bill is accepted if it receives the support of a simple majority of legislators in each chamber — five legislators in the lower house ($M_L = 5$) and three legislators in the upper house ($M_U = 3$). This example is instructive because it captures the scenario in which big state legislators could control the agenda if there were only a lower house, whilst small states constitute a majority in the upper house. For simplicity, I assume there is no discounting of the future ($\delta = 1$).

Equilibrium characterization in the Baron and Ferejohn (1989) framework proceeds in two steps. One step determines the composition of the equilibrium coalition, *taking as given* the cost of ‘buying’ each legislator’s support (i.e. how demanding they are). The other step determines this ‘price’ as a consequence of the legislator’s expected payoff under the future

coalitions that will likely arise if the current proposal fails. Naturally these steps interact with one another — the optimal coalitions depend on legislators’ equilibrium ‘prices’, and these ‘prices’ themselves depend on (future) optimal coalitions. An equilibrium is a specification of ‘prices’ and coalition building rules for which these two procedures are consistent. Since this paper’s novelty stems from the composition of the optimal coalition, I begin with a consideration of that step.

Let v_S be the share of the pie that a small state district expects to receive, *ex ante* (i.e. before the proposer is recognized).⁷ Similarly, let v_B be the *ex ante* share of the pie that big state districts expect to receive. If $v_S = v_B$, then expected per capita spending is the same in both large and small state districts. By contrast, if $v_S > v_B$, then per capita spending is larger (on average) in small state districts, and the converse is true if $v_S < v_B$.

Legislators from a small state reason that if the current proposal is rejected, their constituency will on average receive v_S in the following round of bargaining. Hence, a small state legislator (from either chamber) will accept any current proposal which allocates at least v_S to their district.⁸ This is the ‘price’ of a small state legislator. It is the most that they can credibly demand in the bargaining process. Similarly, v_B is the ‘price’ of a lower house legislator from the big state, and $5v_B$ is the price of the upper house legislator from the big state (since the expected total allocation across the big state is $5v_B$).

To see the effect of preference complementarities across the chambers, note that allocating at least v_S to a small state district earns the support of both the lower and upper house representatives from that state. Similarly, allocating at least $2.5v_B$ to two of the big state districts will earn the support of the two lower house legislators representing those districts as well as the upper house legislator from the big state. Alternatively, allocating v_B to each of the five big state districts would suffice to earn the support of all five lower house legislators as well as the upper house legislator from the big state. Amongst all allocations that earn the support of the upper house legislator from the big state, the equal division allocation is most ‘efficient’ in the sense that it maximizes support amongst lower house legislators from

the big state. A proposer who seeks to build the cheapest winning coalition will naturally prefer efficient allocations to inefficient ones.

A corollary to this result is that, if a proposal is efficient, the upper house legislator from the big state will only accept the proposal if it is supported by *all* of the lower house legislators from that district. This result may seem counter-intuitive — we might, for example, have expected the upper house legislator to support any proposal that receives the support of a majority of lower house legislators from her state.⁹ But note that allocating v_B to a particular big state district causes that district’s legislator to essentially be indifferent between supporting or rejecting the proposal.¹⁰ If the proposer allocated v_B to three big state districts and 0 to the remaining two, then the constituents in three of the five big state districts would be indifferent to the proposal, whilst constituents in the remaining two districts would strictly prefer rejection. The ‘average’ constituent in the big state would clearly prefer that the proposal be rejected. (Intuitively, accepting the proposal only brings $3v_B$ into the state, whereas rejecting the proposal will bring $5v_B$, in expectation, in the following period.) By contrast, allocating v_B to all five big state districts causes all constituents in the big state to be indifferent to the proposal, and so the upper house legislator will be indifferent as well. Indeed, further increasing the allocation to any big state district will then cause the upper house legislator to strictly prefer accepting the proposal, since the proposal is weakly preferred in all districts and strictly in some districts.

3.1 Optimal Coalition in a Unicameral Legislature

Begin by considering coalition building in a unicameral legislature with the same composition as the lower house. The proposer must propose an allocation that receives the support of at least five legislators. In equilibrium, the proposer will always support her own proposal. (Intuitively, if the price of each legislator is equal to her expected share of the pie, and not all legislators are included in the coalition, then the proposer will keep a larger share of the pie than she would expect to receive if she were not the proposer.) Hence, the proposer need

only ‘purchase’ the support of four other legislators. Clearly she will choose the cheapest coalition partners, in order to keep the largest possible share of the pie for herself.

Recall — at this step, we take legislator ‘prices’ (v_S and v_B) as given. There are three cases to consider. First, suppose $v_B < v_S$, so that big state legislators are cheaper coalition partners than small state legislators. If the proposer is a big state legislator, she will invite the four remaining big state legislators to join the coalition. If the proposer is a small state legislator, she will randomly select four of the five big state legislators to join the coalition. In both cases, the ‘purchased’ coalition members are offered v_B each, whilst the proposer keeps the remaining $1 - 4v_B$ for her own district. Small state legislators are never brought into the coalition.

Second, suppose $v_B > v_S$, so that big state legislators are more expensive coalition partners than small state legislators. If the proposer is a big state legislator, she will offer v_S to the four small state legislators to join the coalition, keeping the remaining $1 - 4v_S$ for herself. If the proposer is a small state legislator, she will offer v_S to the three remaining small state legislators and v_B to one randomly selected big state legislators (from the five), keeping $1 - 3v_S - v_B$ for herself. Small state legislators are always included in every coalition.

Finally, suppose $v_B = v_S$, so that big and small state legislators are equally expensive as coalition partners. Then, the proposer is indifferent between coalition partners and will randomly choose between them.¹¹

To summarize the above discussion, let σ_t (resp. β_t) denote the number of small (resp. big) state legislators that a type- t proposer invites into her coalition, where $t \in \{S, B\}$.¹² To be clear, this count does not include the proposer; it only includes those legislators whose

support was directly ‘purchased’. We have¹³:

$$\begin{aligned} (\sigma_S^u, \beta_S^u) &= \begin{cases} (3, 1) & v_S < v_B \\ (0, 4) & v_S > v_B \end{cases} \\ (\sigma_B^u, \beta_B^u) &= \begin{cases} (4, 0) & v_S < v_B \\ (0, 4) & v_S > v_B \end{cases} \end{aligned}$$

The composition of the optimal coalition is determined by the relative prices of big and small state legislators. If the system favors big state districts (so that $v_B > v_S$), then small states will always be included in the coalition, and big state districts will be routinely excluded, regardless of the identity of the proposer. The incentive for legislators to get the best deal for their constituents precludes strategies in which big state legislators conspire against small states and vice versa. Despite commanding a majority of the legislature between them, it may be in interest of big state legislators to form coalitions with small state legislators rather than with each other. Legislators from the big state are not innately coordinated to keep resources within their state.

3.2 Optimal Coalition in a Bicameral Legislature

Now, consider coalition building in a bicameral legislature. The proposer’s task of constructing the optimal coalition is analogous to the previous sub-section, except that, now, the proposer must simultaneously build majorities in both chambers. Analogous to the unicameral legislature, the proposer will always support her own proposal in equilibrium. Additionally, as I show in Lemma 1 in the Appendix, in a symmetric equilibrium, the proposal is also guaranteed to receive the support of *all* of the legislators from the proposer’s state residing in the opposite chamber. (As above, the intuition is that, even after enticing other out-of-state legislators into the coalition, the proposer will keep sufficiently many resources in her constituency to ensure the support of the legislator(s) in the other chamber

who also value that constituency.¹⁴) Hence, if the proposer is from a small state, she can take for granted the support of one legislator in each chamber, and so must propose an allocation that ‘purchases’ the support of an additional four lower house, and two upper house, legislators. The same is true of a lower house legislator from the big state. By contrast, the upper house legislator from the big state can take for granted her own support in the upper house, and the support of five legislators (all of the lower house legislators from the big state) in the lower house. The proposer does not need to purchase the support of as many legislators to achieve a majority in the lower house. I refer to this as the *reduced requirement effect*; it is the main insight that distinguishes this paper from Kalandrakis (2004). In fact, in this case, a lower house majority automatically exists. The proposer need only ‘purchase’ the support of two additional upper house legislators.

The implications of the reduced requirement effect are made apparent by constructing the optimal bicameral coalitions. Again, there are several cases to consider. First, suppose $v_S < v_B$ so that small state legislators are cheaper in both houses. Regardless of their type, proposers will want to build coalitions containing small state legislators whenever possible. If the proposer is from a small state, she will allocate v_S to each of the remaining three small state districts, v_B to one randomly selected big state district, and keep the remaining $1 - 3v_S - v_B$ for her own district/state. Similarly, if the proposer is a lower-house legislator from the big state, she will allocate v_S to each of the four small state districts, and keep the remaining $1 - 4v_S$ for her own district. In both cases, all small districts are included in the coalition. A minimum winning coalition will exist in the lower house, and a super majority will exist in the upper house. Notice that the optimal coalitions are identical to the unicameral case. The complementarity in preferences ensured that, in constructing a lower house majority, the upper house constraint would be automatically satisfied. Malapportionment notwithstanding, bicameralism has conferred no particular advantage to small states.

Now suppose that the proposer is the upper house legislator from the big state. Given the reduced requirement effect, the proposer can presume the existence of a majority in the lower

house. To achieve a majority in the upper house, she need only allocate v_S randomly amongst two of the four small state districts. The remaining $1 - 2v_S$ will be equally allocated amongst the five districts in the big state.¹⁵ The properties of the equilibrium have reversed — a bare majority exists in the upper house, whereas a super-majority exists in the lower house. Although big state districts are more expensive than small state districts, the equilibrium coalition excludes more small state districts, and directs more resources towards the big state districts. Here we see the *reduced requirement effect* at work. To the extent that the upper house legislator from the big state is recognized as the proposer, bicameralism can skew the allocation away from small states, even when small state legislators are cheaper coalition partners than big state legislators.

The optimal coalitions can similarly be determined for the remaining cases. We have:

$$\begin{aligned}
 (\sigma_S^b, \beta_S^b) &= \begin{cases} (3, 1) & v_S < v_B \\ (2, 2) & v_B < v_S < 3v_B \\ (1, 5) & v_S > 3v_B \end{cases} \\
 (\sigma_{BL}^b, \beta_{BL}^b) &= \begin{cases} (4, 0) & v_S < v_B \\ (2, 2) & v_S > v_B \end{cases} \\
 (\sigma_{BU}^b, \beta_{BU}^b) &= (2, 0)
 \end{aligned}$$

As before, as small states become more expensive, the proposer switches towards coalitions that privilege big states.¹⁶ However, unlike the unicameral case, small state legislators are not completely excluded from the coalition as they become relatively expensive. Instead, resources are always directed towards at least one small state, and often two. This is the ‘standard channel’ by which bicameralism is perceived to benefit small states. The need to achieve a majority in the upper house makes it more likely that small states will be included in coalitions, since they are more numerous in that chamber.

Comparing the optimal coalitions under the two regimes, we see that bicameralism has two potential effects. On the one hand, the standard channel tends to increase representation of small states. This is especially apparent in the case when $v_S > v_B$, where small state legislators would have been entirely excluded from coalitions under unicameralism, but at least one and often two small states are included under bicameralism.¹⁷ On the other hand, the *reduced requirement effect* tends to decrease representation of small states, by coordinating big state proposers to keep resources within their state. This is made apparent in the case when $v_S < v_B$, where all small states are included in the coalition under unicameralism, but some are excluded under bicameralism when the proposer is an upper house legislator from the big state. Notice that the standard channel tends to benefit small states in cases where they would otherwise have been excluded from coalitions under unicameralism (i.e. when $v_S > v_B$). Similarly, the reduced requirement channel particularly hurts small states in cases where they would otherwise have been included in coalitions under unicameralism (i.e. when $v_S < v_B$). Hence, the reduced requirement channel hurts small states under bicameralism precisely in those cases when unicameralism would already disadvantage small states. This is the main insight of this paper — whilst bicameralism *can* benefit small states, it is not guaranteed to do so, and may in fact be detrimental to their interests. Indeed, bicameralism can have the effect of amplifying inequities that arise under unicameralism. If there were a genuine concern that small states would fair poorly under unicameralism, then bicameralism, with a malapportioned upper house, may merely make things worse.

I stress that the process driving these results stems purely from the nature of optimal coalition building. The effects outlined in this section will be present whenever proposers build coalitions comprising of the cheapest legislators. As I show in the following section, equilibrium legislator prices will depend on various assumptions about the bargaining framework, including the recognition rule, the process following rejection of a proposal and so on. Varying these details may cause equilibrium legislator prices to change. However these details do not effect the coalition building process itself. The insights regarding the optimal

coalition will remain true, no matter how the legislator prices are determined. Hence, the process highlighted in this paper will be robust to different specifications of the bargaining protocol.

I end this subsection with a technical note. As should be clear from the above example, coalition building is simple when either $v_S < v_B$ or $v_S > kv_B$. In the former case, small state legislators are cheaper than big state legislators in both chambers, and so the optimal allocation will tend to privilege small state districts. In the latter case, the opposite is true, and so the optimal allocation will tend to privilege big state districts. Coalition building in the intermediate case ($v_B < v_S < kv_B$) is more interesting. Big state legislators are cheaper in the lower house, but more expensive in the upper house. The proposer would ideally target big state districts to build lower house majorities, whilst targeting small state districts to build upper house majorities. How should the proposer allocate resources, given that the same allocation must achieve both goals simultaneously, and there is no clear sense about which approach is most attractive? In Proposition 1 in the Appendix, I provide a simple algorithm that characterizes the optimal coalition in this intermediate case for the general model (where b , s and k are arbitrary, and super-majority requirements are allowed). This paper is the first to provide a succinct characterization of the composition of the optimal coalition in this intermediate case. I briefly summarize it here.

First, I say that the *lower house dominance property* holds if, in allocating v_B to as many big state districts as are needed to achieve the required majority in the lower house, the upper house majority is satisfied as well.¹⁸ The algorithm is as follows: First, determine if the lower house dominance property holds. If so, build the coalition by allocating resources to big state districts until the lower house majority is achieved. If not, allocate v_S to a randomly selected small state district. This reduces the residual majority requirements in each chamber by one. Repeat the above procedure, first checking whether the lower house dominance property holds on the residual requirement, and adding a big or small district legislator to the coalition as appropriate. Following this procedure, the proposer will allocate

v_S to small state districts until the residual majority requirements satisfy the lower house dominance property, whereupon she switches to allocating v_B to sufficiently many big state districts as to achieve a majority in both chambers.

4 Equilibrium Distribution and Legislator ‘Prices’

The previous section characterized the optimal coalitions, taking as given the ‘prices’ of different legislators. Since the equilibrium prices had not yet been determined, I characterized the optimal coalitions for any possible arrangement of prices. In this section, I pin down the prices that will prevail in equilibrium.

In the previous section, I conceived of each legislator’s price as their expected future payoff if the current proposal were rejected. These payoffs will in turn depend on the anticipated composition of future coalitions — which will themselves depend upon the legislators’ beliefs about future prices. Thus, an equilibrium is a set of beliefs about legislator prices and likely future coalitions which are consistent with one another. To make this clearer, let (v_S, v_B) denote the legislators’ conjecture about future prices, which determine the optimal coalitions. Similarly, let (\hat{v}_S, \hat{v}_B) denote the expected payoffs implied by these coalitions. In equilibrium, the conjectured and implied prices must coincide.

4.1 Distribution under Unicameralism

As in the previous section, I first study equilibrium in the unicameral legislature. Suppose legislators conjecture that $v_S < v_B$. Given the discussion above, small states will always be included in coalitions. Now, consider the expected payoffs implied by these coalitions. A small state will be allocated v_S whenever the proposer is not from that state, and will retain $1 - 3v_S - v_B$ when the proposer is. Since that state’s legislator is recognized with probability

p_S , the state's expected payoff is:

$$\hat{v}_S = p_S(1 - 3v_S - v_B) + (1 - p_S)v_S$$

Similarly, a big state district will have a one-in-five chance of being allocated v_B if the proposer is from a small state, will be allocated 0 if the proposer is from a different big state district, and will retain $1 - 4v_S$ when the proposer is from the district. Hence, the expected payoff to a given big state district is:

$$\hat{v}_B = p_B(1 - 4v_S) + 4p_S\frac{v_B}{5}$$

If the conjectured beliefs are indeed equilibrium consistent, then $\hat{v}_B = v_B$ and $\hat{v}_S = v_S$. Solving this system of equations gives $v_S = \frac{1}{4}$ and $v_B = 0$. Since small states are drawn into every coalition, and big state districts are excluded from most, small states are expected to receive a larger share of the pie than big states. But this is inconsistent with the conjecture that $v_S < v_B$. The belief that small states will do worse than big states is not equilibrium consistent. We can reject such beliefs.

We can similarly show that any conjecture with $v_S > v_B$ cannot be equilibrium consistent, and so in equilibrium $v_S = v_B = \frac{1}{9}$.¹⁹ On average, we should not expect the bargaining game to favor big states over small states, and vice versa. Notwithstanding the big state's numerical advantage, allocations within the unicameral legislature will be equalized across districts, on average. Here again, we see that the concern about policy usurpation by the big state may well be misplaced.

This result follows intuitively from the nature of coalition building. If it were the case that $v_S < v_B$, then small state legislators would make cheaper coalition partners, and so small states would always receive funds whilst big state districts would routinely be excluded. But this has the effect of directing resources away from big states and towards small states, which counteracts the effect of small states appearing cheaper. The bargaining protocol

exhibits a strong force towards equality. Any change that tends to give one legislator a relative advantage immediately makes that legislator more expensive, and thus excludes them from coalitions more often. This force counteracts the effect of the original change. In this particular example, the two forces exactly counter-balance one another, so that no district can achieve a lasting advantage over the others.

In a more general setting (e.g. with positive discounting $\delta < 1$), the two forces may not always exactly counter-balance, although the equal allocation of resources is still predicted over a broad range of recognition probabilities. We can sustain equilibria with unequal allocations, if recognition probabilities are sufficiently skewed (see Kalandrakis (2006)). As I show in Proposition 2 in the Appendix, generically, per capita funding will favor big states if big states have significant proposal power. Similarly, per capita funding will favor small states if they have significant proposal power. If recognition probabilities are relatively moderate, then per-capita funding will tend to be equalized across big and small states.

4.2 Distribution under Bicameralism

Now, consider the expected payoffs under bicameralism. Suppose the legislators conjecture that $v_S < v_B$. Then a given small state will receive v_S if the proposer is from a different small state or is a lower house legislator from the big state; it has a one-in-two chance of receiving v_S if the proposer is the upper house legislator from the big state; and it will retain $1 - 3v_S - v_B$ if the proposer is from that state. Hence, its expected allocation will be:

$$\hat{v}_S = (p_{S^L} + p_{S^U})(1 - 3v_S - v_B) + (3p_{S^L} + 3p_{S^U} + 5p_{B^L})v_S + p_{B^U}\frac{1}{2}v_S$$

Similarly, a big state district will have a one-in-five chance of receiving v_B if the proposer is from a small state; it will receive zero if the proposer is a lower house legislator from a different big state district; it will retain $1 - 4v_S$ if the proposer is the lower house legislator from that district, and it will retain its equal share of $1 - 2v_S$ if the proposer is the upper

house legislator from the big state. Hence, its expected allocation is:

$$\hat{v}_B = p_{BL} (1 - 4v_S) + p_{BU} \frac{1 - 2v_S}{5} + \frac{4}{5} (p_{SL} + p_{SU}) v_B$$

Again, if the conjectured beliefs are indeed equilibrium consistent, then $\hat{v}_B = v_B$ and $\hat{v}_S = v_S$.

Solving this system of equations gives:

$$v_S = \frac{4(p_{SL} + p_{SU})}{16(p_{SL} + p_{SU}) + \frac{5}{2}p_{BU}}$$

$$v_B = \frac{p_{BU}}{32(p_{SL} + p_{SU}) + 5p_{BU}}$$

Consistency requires that $v_S < v_B$, and this will be the case provided that $p_{SL} + p_{SU} < \frac{1}{8}p_{BU}$. Hence, if the agenda setting power of the upper house legislator from the big state is large relative to small state legislators, then bicameralism will skew allocations towards the big state. And this is true, in spite of unicameralism producing equal outcomes, *ex ante*. Small states are adversely affected through the reduced requirement effect when the upper house legislator from the big state has sufficient proposal power.

We can similarly find conditions under which bicameralism is likely to produce equal *ex ante* outcomes, or to advantage small states. Analogous to the previous discussion, bicameralism will tend to favor small states when their proposal power is high, and will tend to produce equal allocations when differences in proposer power are moderate. Proposition 2, in the Appendix, fully characterizes the relationship between proposal power and equilibrium payoffs.

I conclude this section with the following insight. Suppose, the unicameral legislature were modeled on the upper house rather than the lower house, so that the chamber represented states equally. Again, in this framework, we can find a range of recognition probabilities over which small states do better under unicameralism than they would under bicameralism. Hence, the *reduced requirement effect* has caused the equilibrium allocation to be

more generous to big states than would be the case either if the lower house or upper house could enact policy unilaterally. As such, and in contrast to Hammond and Miller (1987) and Heller (1997), bicameralism does not necessarily produce outcomes that are a compromise between those that would likely be chosen if each chamber could act unilaterally.

5 Comparison of Institutions

In the preceding two sections, I examined optimal coalition and equilibrium outcomes under both unicameralism and bicameralism, through the lens of a simple example. Equilibrium in the general case (e.g. for differently composed legislatures) is characterized in the Appendix. The lessons from the simple example extend to the general setting. In this section, I compare distributional outcomes under unicameralism and bicameralism in the general setting. In particular, I am interested in the following question: If small states are expected to fare poorly (in the sense of lower per-capita allocations) under unicameralism, under what conditions will they fare better under bicameralism?

To facilitate comparison, it is helpful to conceive of the unicameral legislature as simply the lower house of the bicameral legislature acting unilaterally. Hence, the composition and the majority requirements are the same in the unicameral legislature and in the lower house of a bicameral legislature. Similarly, the recognition probabilities in the unicameral legislature are identical to recognition probabilities in the bicameral legislature, conditional upon proposers being drawn from the lower house.

The effect of bicameralism will arise through two channels: First is the *coalition composition channel* — bicameralism may skew the nature of the optimal coalitions. Second is the *agenda power channel* — bicameralism may change the odds that the proposer represents a small rather than big state. The direction of the coalition composition channel will, in turn, depend on the size of the (super)-majority requirements in each chamber. Each point in Figure 1, below, represents a possible configuration of majority requirements in the bicameral

legislature. The further along the horizontal (resp. vertical) axis, the more demanding is the majority requirement in the lower (resp. upper) house.

Figure 1: The Coalition Composition Channel and Effect of (Super)-Majority Requirements

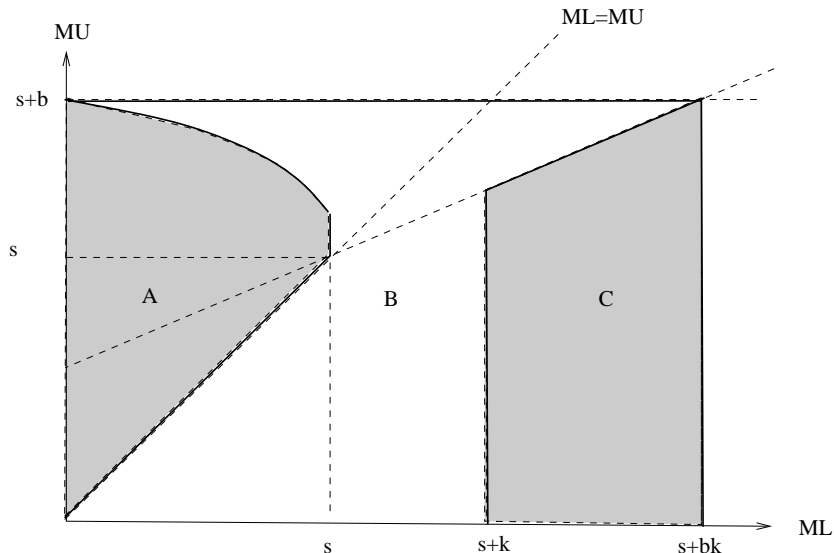


Figure 1 partitions the space of feasible majority requirements into 3 regions. These regions demonstrate the qualitative effect of bicameralism on the allocation to small states, solely through the coalition composition channel. Suppose that small states are disadvantaged under unicameralism, and assume that the agenda control channel is inoperative (i.e. bicameralism does not alter the proposal power of big states relative to small states). Then, legislatures whose majority requirement configurations lie in region *A* will see an improvement in the expected allocation to small states, whilst legislatures contained in region *B* make small states even worse-off. Bicameralism does not confer any advantage or disadvantage to legislatures contained in region *C*.

In the previous section, we noted that, fixing the composition of coalitions, the expected allocation to a legislator's constituency is weakly increasing in their proposal power. Hence, introducing the agenda power channel extends the result in the following way: For legislatures in region *A*, bicameralism will improve the expected allocation to small states unless the upper house sufficiently dilutes their proposal power. By contrast, for legislatures in

region B , bicameralism will worsen the expected allocation to small states unless the upper house sufficiently improves their proposal power. Finally, legislatures in region C will make small states better (resp. worse) off whenever bicameralism increases (resp. decreases) their relative control of the agenda. Proposition 3 (in the Appendix) states these results formally, and quantifies how much proposal power must improve (or may worsen) before small states are made better off.

To parse this result, begin by considering region B , where bicameralism tends to worsen outcomes for small states. As Figure 1 makes clear, B is itself comprised of two distinct regions. The lower section is of particular interest because it showcases the reduced requirement effect. Focus on this region. Given the motivating concern, suppose that small states would be disadvantaged under unicameralism. (This assumption is maintained throughout the remainder of the discussion.) If so, then the optimal unicameral legislature will be largely (or solely) comprised of small states. Now consider a bicameral legislature. Since $M_L \geq M_U$ in this region (i.e. the upper house majority requirement is not too demanding), this optimal unicameral coalition will automatically satisfy the upper house majority requirement. Hence, bicameralism does not intrinsically draw more small states into the coalition. Furthermore, since $M_L < s + k$, an upper house legislator from a big state can satisfy the lower house majority requirement without utilizing all of the small states. Since such a proposer benefits from the reduced requirement effect, she will draw fewer small state legislators into her coalition than would other types of proposers. This entails that small states will be excluded from bicameral coalitions more frequently than under unicameralism. Furthermore, since small states were assumed to fare poorly under unicameralism, they will be made even worse off under bicameralism. (By contrast, if $M_L > s + k$, so that the legislature is in region C , then even after accounting for her reduced requirement, an upper house proposer from the big state would still invite every small state into the coalition — and so small states suffer no loss.) Hence, generically, the reduced requirement effect will have its strongest bite when the majority requirements in neither chamber are too demanding (and when recognition rights

are not too favorable to small states).

Given the above discussion, the following comparative statics are immediate: First, increasing the majority requirements in either chamber decreases the likelihood that bicameralism will further disadvantage small states through the reduced requirement channel; doing so pushes the legislature closer to regions A or C . Second, legislatures with a relatively large number of big states will be less likely to disadvantage small states through the reduced requirement effect. Third, and relatedly, the reduced requirement effect will be less likely to be operative in legislatures where there is a large imbalance between the size of states (i.e. high k).²⁰ To build intuition for these latter results, suppose the lower house decides by simple majority. Then the requirement that $M_L < s + k$ simplifies to $k(b - 2) < s - 2$.²¹ When there are many big states or when big states are large relative to small states, proposers will quickly exhaust the supply of small states when building coalitions. Hence, notwithstanding the reduced requirement effect, small states will continue to be included in coalitions with high probability.

Whilst the lower section of region B is of most interest, I briefly discuss the other regions. Each of these regions distinguish themselves from the lower section of B by having majority requirements in at least one chamber that are sufficiently demanding. In region A , the upper house majority requirement is more demanding than the lower house requirement in absolute terms. This has the effect of pulling more small states into the coalition, thus conferring an advantage to them. In the upper section of B , the upper house majority requirement is more demanding still. Having exhausted all the small states, the optimal coalition must additionally draw more big state legislators into the coalition. (This effect is particularly strong, because big state districts must be added k at a time, to earn the support of that state's upper house legislator.) This has the opposite effect of increasing the payoff to big state districts, at the expense of small states. In region C , the lower house majority requirement is sufficiently demanding that all small states will be included in every coalition, so the coalition composition channel confers neither advantage nor disadvantage.

I stress the conditional nature of these results. The above discussion was premised on the assumption that small states would be disadvantaged under unicameralism. As was demonstrated in Section 4, distributional outcomes are jointly determined by both the composition of the chamber and the assignment of proposal power, and there is no guarantee that unicameralism will favor big states simply because of their numerical advantage. Nevertheless, assuming that unicameralism would favor big states — as the motivating concern does — Proposition 3 determines whether small states will likely fare better or worse under bicameralism. I further stress that the main channel that generates these results depends purely on the composition of optimal coalitions, which, in turn follows solely from the assumption that proposers seek out the cheapest coalition partners. Hence, notwithstanding the many other assumptions invoked in this paper, the underlying logic of these results should be fairly robust to alternative formulations of the legislature.

The analysis demonstrated that the coalition composition effect is ambiguous and depends on where a legislature is located in Figure 1. I conclude this section by arguing that most legislatures are likely to be found in the lower section of region B , and so the region in which the reduced requirement effect is salient is also the region which is most empirically relevant. To see this, I calibrate the model's parameters to reflect the current legislative arrangements in four polities (United States, Australia, Switzerland and the European Union). Additionally, I calibrate the model against the historical arrangements in the United States and Australia immediately after the adoption of the relevant constitutions, to capture the legislature as it would have been anticipated by the constitution's framers. The calibration is similar to Kalandrakis (2004), and is constructed by estimating a single level regression tree (see Kuhn and Johnson (2013)). The results are displayed in Table 1.

As Table 1 shows, with one exception, legislatures are all found in the lower section of region B — precisely the region in which the reduced requirement effect tends to adversely affect small states via the coalition composition channel.

That this is the case should not be unexpected. To see this, take as a baseline those

Table 1: Calibrated Parameters for Current and Historical Legislatures

Country	s	b	k	N_L	N_U	M_L	M_U	Region
United States								
1st Congress (1789-80)	8	5	2.25	19.26	13	9.62	6.5	B^-
112th Congress (2015-16)	46	4	5.63	68.53	50	34.26	30	B^-
Australia								
At Federation (1901)	4	2	3.77	11.54	6	5.77	3	B^-
Current (2016)	3	3	3.71	14.03	6	7.07	3	C
European Union	22	6	4.68	50.07	28	25.03	14	B^-
Switzerland	20	3	4.26	32.49	23	16.39	11.5	B^-

Note: N_L and N_U represent the total number of legislators in the lower and upper house, respectively. B^+ and B^- denote the upper and lower sections of region B , respectively. The upper house is the Senate in the U.S. and Australia, the Council of Ministers for the E.U, and the Council of States in Switzerland. The lower house is the House of Representatives in the U.S. and Australia, the European Parliament for the E.U, and the National Council in Switzerland. The calibration procedure is detailed in the Appendix. The size of legislatures is rescaled so that each small state has one representative in the lower house.

For the United States, the big states in 1789 were VA, MA, PA, NY, and MA; the big states in 2013-14 were CA, TX, NY, and FL. For Australia, the big states in 1901 are NSW and VIC; the big states in 2013 are NSW, VIC and QLD. For the European Union, the big states are France, Germany, Italy, Spain, U.K and Poland. For Switzerland, the big states (cantons) are Zurich, Bern and Vaud. (The three pairs of semi-cantons are each combined to form three cantons.) Simple majorities are assumed in all cases, except the U.S. 112th Congress, for which a filibuster-proof majority (60%) is required in the Senate.

legislatures with simple majority requirements in each chamber; five of six of the calibrated legislatures satisfy this property. Then $M_L > M_U$, which implies that such legislatures cannot be in region A . (Indeed, we should expect few legislatures in region A , given that the upper house majority requirement in this region is more demanding than the lower house requirement *in absolute terms*.²² Not only are such arrangements rarely observed, they would be impossible whenever the lower house is at least twice as large as the upper house). Hence, a typical legislature is likely to be found in regions B or C . Moreover, as we have seen, with the simple majority requirement, legislatures will be in region C only if they have relatively many big states and the imbalance in the size of big and small states is large. As Table 1 demonstrates, most bicameral legislatures have many more small states than big states, which explains why few legislatures are contained in that region. Together with these intuitions, Table 1 provides anecdotal evidence that the novel effect identified in this paper may be relevant in many real world bicameral legislatures.

6 Conclusion

A significant motivation for the malapportioned design of bicameral legislatures is to protect small states from policy usurpation by big states, who would ordinarily dominate the lower chamber. This paper investigated whether bicameralism is effective in remedying this concern. I developed a model of bicameralism, in which lower house legislators are primarily motivated by outcomes in their district, whilst upper house legislators are concerned with state-wide outcomes. Given this assumption, under unicameralism, big state legislators are not inherently coordinated to form coalitions that privilege one another — implying that the fear that big states will conspire against small states is misplaced.

By contrast, by introducing legislators with state-wide interests to the decision making process, bicameralism may have the unintended consequence of generating precisely this sort of coordinated behavior. Crucial to this result is the role of preference complementarities

between legislators across chambers. This implicit coordination generated the *reduced requirement effect*, which reduced the number of out-of-state legislators that big state actors need to entice into the coalition, thereby retaining more resources within their own state. This effect followed solely from the assumption that proposers would build coalitions comprising of the cheapest legislators, and is robust to alternative specifications of the remaining assumptions in the model.

Assuming that unicameralism favored big states, I examined the distributional effect of bicameralism. Over a range of parameters, the reduced requirement effect is salient, and unless bicameralism significantly increases their agenda setting power, small states are made even worse off. Through a calibration, the model’s insights are shown to be relevant to many existing legislatures, suggesting important policy considerations for future constitutional design.

Notes

¹For example, in Australia, Mexico, Russia, Switzerland and the United States, amongst other nations, states are represented equally in the upper chamber (Lijphart, 1999). Similarly, in the European Union, member countries are equally represented on the European Council, in contrast to the European Parliament, where countries are represented in proportion to population size. In other nations, Burundi, Canada, Germany, South Africa and Spain amongst them, states (or administrative units) are not necessarily equally represented, however, the essential feature that smaller regions are over-represented, tends to remain true. By contrast, the upper house is not malapportioned in nations including Italy and Japan, nor in any of the U.S. states with bicameral state legislatures (Tsebelis and Money, 1997).

²Kalandrakis’s framework is perhaps most salient when considering federal block grants to states.

³Formally, there is a function $\rho : \{1, \dots, s + kb\} \rightarrow \{1, \dots, s + b\}$ which maps every district into a unique state. Then $u_j^U(x) = \sum_{\{i|\rho(i)=j\}} x_i$.

⁴Although legislators have equal voting power, some (e.g. committee chairs or party leaders) may exert greater control of the agenda than others. Type-dependent recognition probabilities allows for the possibility that agenda setters are on average more (or less) likely to be drawn from large states. For example, 21 of Australia’s 29 Prime Ministers (as of 2016) represented districts located in the country’s two largest states

(New South Wales and Victoria).

⁵For example, suppose there are two big states and one small state. If the small state legislators are centrists, whilst the legislators from the big states are either right- or left-wing, then small state legislators will likely find themselves in every governing coalition. By contrast, if the small state legislators are on one extreme, then they may find themselves frequently excluded from coalitions.

⁶Even in the simplest case of a uni-dimensional policy space, where the median voter theorem applies, we would need to specify the distribution of voter preferences in every district, in order to determine the state-wide medians. In multidimensional spaces, median policies do not generically exist, complicating this task even further.

⁷Stationarity ensures that this value does not vary across bargaining rounds.

⁸Legislators are assumed to be risk-neutral, and so will accept any offer in which their current payoff is at least as their expected future payoff after discounting. The model is readily generalized to include risk-averse agents. Since the model's qualitative predictions are unaffected by risk preferences, the focus on risk neutral agents is benign.

⁹Ansolabehere, Snyder and Ting (2003) invoke this decision procedure. As I argued in the Introduction, and again below, such a procedure cannot be rationalized as a consequence of the maximizing behavior of agents.

¹⁰Technical considerations motivate the assumption that legislators support proposals to which they are indifferent. If the policy space were discrete — i.e. if there is a minimal increment in which resources can be allocated — then it suffices to supplement the allocation by one unit, so that the allocation is strictly preferred.

¹¹It will turn out that to sustain an equilibrium, the probabilities with which the proposer includes big and small legislators in the coalition may differ. Since the proposer is indifferent between all coalitions, she is free to randomize amongst these at any rate.

¹²Note: in Kalandrakis (2004) β_t denotes the number of big states that receive funds, whereas in this paper β_t denotes the number of big state districts receiving funds. This is the main difference in notation between the papers, and stems from the differing assumptions on whether policies are distributed at the district or state level.

¹³The superscript u indicates a unicameral legislature. Below, a superscript b will denote a bicameral legislature.

¹⁴The most interesting case is that of a big state proposer in the upper house. Since she is willing to support her own proposal, she must allocate at least $5v_B$ amongst the districts that comprise the big state. If she allocates these resources equally across the five districts in her state, then each district will receive at

least v_B , which is sufficient to earn the support of each of the five lower house legislators from the big state.

¹⁵The equal division is not essential; any allocation that assigns at least v_B to each big state district is consistent with optimal coalition building. The equal division requirement follows from the symmetry assumption discussed in the previous section. An unequal allocation would result in arbitrary differences in outcomes between big state districts.

¹⁶In the case of small state proposers, the switch towards big state coalition partners occurs at two points: when v_S exceeds v_B , and when it exceeds $3v_B$. The second jump occurs because for $v_S \in (v_B, 5v_B)$, big state legislators are cheaper coalition partners in the lower house, but more expensive in the upper house. Taken in isolation, the proposer would ideally build a coalition of big state lower house legislators and small state upper house legislators. But given the complementarity in preferences, targeting small states to meet the upper house constraint also contributes to the lower house requirement, and vice versa. The proposer faces a choice about which type of districts to target. For $v_S < 3v_B$, small state legislators are relatively less expensive in the lower house than big state legislators are in the upper house. Hence, at the margin, the proposer prefers targeting small states. The opposite is true when $v_S > 3v_B$.

¹⁷The standard channel can also benefit small states if the super-majority requirement in the upper house is so demanding that more legislators are needed to pass a bill in the upper house (in absolute terms) than in the lower house. If so, small states may benefit even when $v_S < v_B$, if some were excluded from the optimal unicameral coalition, but are now included in the optimal bicameral coalition.

¹⁸I implicitly assume that the allocation to big state districts is clustered within states to ensure that the allocation most efficiently target lower and upper house legislators from big states.

¹⁹Such an equilibrium can be sustained by mixed strategies which result in each small district receiving an allocation of $\frac{1}{9}$ with probability $\sigma_S = \frac{1-5p_S}{1-p_S}$ and each big state district receiving an allocation of $\frac{1}{9}$ with probability $\sigma_B = \frac{5p_S}{1+p_S}$. Notice that σ_S is decreasing in p_S and σ_B is increasing in p_S . As small states become more likely to be recognized as the proposer, they are more likely to keep the residual surplus. To ensure an equal allocation across all districts, this effect must be counter-balanced by small states being omitted from coalitions more frequently, when they are not the proposer.

²⁰I thank an anonymous referee for suggesting this comparison to me.

²¹This assumes there are an odd number of lower house legislators. If even, we need $k(b-2) < s-1$.

²²I.e. the total number of legislators (not just the fraction) required to achieve a majority in the upper house is larger than the number required in the lower house.

7 Appendix

7.1 Formalities

In this section, I present the formal analysis of the model. Since the intuition has been discussed above, I proceed without discussion. I begin by characterizing the equilibrium of the bicameral game. This embeds the equilibrium of the unicameral game as the special case, in which $M_U = 0$ and $p_{SV} = p_{BV} = 0$.

7.1.1 Optimal Coalitions

A proposal by a type- t proposer is a pair (σ_t, β_t) , where σ_t and β_t are the number of small and big districts to which are allocated resources. Since mixed strategies are allowed, let $\mu_t(\sigma, \beta)$ be the probability that a type- t proposer builds a (σ, β) -coalition. Define $\bar{\sigma}_t = \sum_{(\sigma, \beta)} \mu_t(\sigma, \beta) \sigma$ and $\bar{\beta}_t = \sum_{(\sigma, \beta)} \mu_t(\sigma, \beta) \beta$.

Lemma 1. *In equilibrium, the proposer will always accept his own proposal. Moreover, a proposal from a lower house agent will be accepted by that state's upper house legislator. Similarly, in a symmetric equilibrium, a proposal from an upper house agent will be accepted by every lower house agent from that state.*

Let (v_S, v_B) be a pair of equilibrium shares. Given Lemma 1, an optimal coalition can be formulated as the solution to the following cost minimization problem:

$$\begin{aligned}
 (\sigma_t, \beta_t) &\in \arg \min_{\sigma, \beta} v_S \sigma + v_B \beta \\
 \text{s.t } \sigma + \beta &\geq M_L - 1 - (k - 1) \mathbf{1}_{BV} [t] \\
 \sigma + \left\lfloor \frac{\beta}{k} \right\rfloor &\geq M_U - 1 \\
 \beta &\leq bk - \mathbf{1}_{BL} [t] - k \mathbf{1}_{BV} [t] \\
 \sigma &\leq s - \mathbf{1}_S [t]
 \end{aligned}$$

where $\lfloor x \rfloor$ denotes the largest integer weakly less than x . The first two constraints are the

lower and upper house majority constraints (respectively), whilst the last two constraints are the legislator supply constraints. Naturally, the optimal proposal rule μ_t may only put positive weight on minimizers of the above problem.

Let $l_t = 1 + (k - 1) 1_{B^U} [t]$, be the number of lower house legislators (including possibly the proposer herself) whose support the proposer gets for free. The following proposition characterizes the composition of the optimal coalition:

Proposition 1. *Let $\sigma'_t = \left\lfloor \frac{k(M_U - 1) - (M_L - l_t)}{k - 1} \right\rfloor$ and $\theta = \text{mod}_{k-1}(kM_U - M_L) + 1$. Given the equilibrium expected share ratio $v = \frac{v_S}{v_B}$, the optimal coalition for a type- t proposer is given by:*

$$\sigma_t(v) = \begin{cases} \min \{ \max \{ M_U - 1, M_L - l_t \}, s - 1_S[t] \} & v < 1 \\ \min \{ \max \{ 0, M_L - 1_S[t] - bk, M_U - 1_S[t] - b, \sigma'_t + 1 \}, s \} & 1 < v < \theta \\ \min \{ \max \{ 0, M_L - 1_S[t] - bk, M_U - 1_S[t] - b, \sigma'_t \}, s \} & \theta < v < k \\ \max \{ 0, M_L - 1_S[t] - bk, M_U - 1_S[t] - b \} & v > k \end{cases}$$

$$\beta_t(v) = \begin{cases} \max \{ 0, M_L - s - l_t 1_B[t], k(M_U - s - 1_B[t]) \} & v < 1 \\ \max \{ M_L - l_t - \sigma_t(v), k(M_U - 1 - \sigma_t(v)) \} & v > 1 \end{cases}$$

Furthermore, unless $t = B^L$ and $\sigma_t(v) = M_U - b$, $\beta_t(v) = \min \{ \max \{ M_L - l_t, k(M_U - 1), bk - l_t 1_B[t] \} \}$ whenever $v > k$. If $v = \psi$, where $\psi \in \{1, \theta, k\}$, then any mixture of the optimal coalitions for $v = \psi - \frac{1}{2}$ and $v = \psi + \frac{1}{2}$ is optimal.

7.1.2 Equilibrium Shares

Let $P = b(kp_{B^L} + p_{B^U})$ denote the aggregate probability that the proposer is from a big state, and let $\alpha_B = \frac{bkp_{B^L}}{bkp_{B^L} + bp_{B^U}}$ and $\alpha_S = \frac{p_{S^L}}{p_{S^L} + p_{S^U}}$ be conditional probabilities that a proposer is from the lower house. Given optimal coalition formation, the expected share of the pie

that is received by big state districts is:

$$v_B = \frac{P}{bk} - \frac{P}{bk} (\alpha_B \bar{\sigma}_{BL} + (1 - \alpha_B) \bar{\sigma}_{BU}) \delta v_S + (1 - P) \frac{\bar{\beta}_S}{bk} \delta v_B \quad (1)$$

Similarly, the expected share accruing to small state is:

$$v_S = \frac{1 - P}{s} - \frac{1 - P}{s} \bar{\beta}_S \delta v_B + P \frac{(\alpha_B \bar{\sigma}_{BL} + (1 - \alpha_B) \bar{\sigma}_{BU})}{s} \delta v_S \quad (2)$$

Solving (1) and (2) gives:

$$\frac{v_S}{v_B} = \frac{1 - P}{P} \frac{bk - \delta \bar{\beta}_S}{s - \delta (\alpha_B \bar{\sigma}_{BL} + (1 - \alpha_B) \bar{\sigma}_{BU})} \quad (3)$$

which is analogous to (A.2) in Appendix A of Kalandrakis (2004) and Proposition 2 in McCarty (2000).

Let $v = \frac{v_S}{v_B}$ be the share ratio, and let $\phi(v)$ be defined by:

$$\phi(v) = \frac{bk - \delta \bar{\beta}_S(v)}{s - \delta (\alpha_B \bar{\sigma}_{BL}(v) + (1 - \alpha_B) \bar{\sigma}_{BU}(v))} \quad (4)$$

An equilibrium is simply a fixed point of the mapping $v = \frac{1-P}{P} \phi(v)$. Since the optimal coalitions $\{(\sigma_t, \beta_t)\}_{t \in T}$ are piece-wise constant over the intervals $[0, 1), (1, \theta), (\theta, k)$ and (k, ∞) , so is $\phi(v)$. (For convenience, denote $\phi(v) = \phi_1$ when $v < 1$. Similarly let $\phi(v)$ be denoted by ϕ_2, ϕ_3 and ϕ_4 and over the remaining intervals, respectively.)

Proposition 2. *The equilibrium shares can be characterized (uniquely) as follows:*

$$v = \begin{cases} \frac{1-P}{P} \phi_1 & P \in \left(\frac{\phi_1}{1+\phi_1}, 1 \right] \\ 1 & P \in \left[\frac{\phi_2}{1+\phi_2}, \frac{\phi_1}{1+\phi_1} \right] \\ \frac{1-P}{P} \phi_2 & P \in \left(\frac{\phi_2}{\theta+\phi_2}, \frac{\phi_2}{1+\phi_2} \right) \\ \theta & P \in \left[\frac{\phi_3}{\theta+\phi_3}, \frac{\phi_2}{\theta+\phi_2} \right] \\ \frac{1-P}{P} \phi_3 & P \in \left(\frac{\phi_3}{k+\phi_3}, \frac{\phi_3}{\theta+\phi_3} \right) \\ k & P \in \left[\frac{\phi_4}{k+\phi_4}, \frac{\phi_3}{k+\phi_3} \right] \\ \frac{1-P}{P} \phi_4 & P \in \left[0, \frac{\phi_4}{k+\phi_4} \right) \end{cases}$$

7.1.3 Comparing Unicameralism and Bicameralism

Proposition 3. *Suppose the environment in the unicameral legislature favors big state districts in equilibrium. (i.e. the parameters are such that $v^u < 1$). Then the distribution under a bicameral legislature favors the big district even more (i.e. $v < v^u$) if and only if $\alpha_S > \underline{\alpha}_S = \alpha_B \frac{\phi_1(\alpha_B)}{\phi_1^u}$. Furthermore $\underline{\alpha}_S < \alpha_B$ whenever $(M_L, M_U) \in A_2$ and $\underline{\alpha}_S > \alpha_B$ whenever $(M_L, M_U) \in A_1$.*

7.2 Proofs

Proof of Lemma 1. Since there is no delay, we have $bkv_B + sv_S = 1$. Suppose the proposer is from a small state. Then the most he will allocate to districts outside his state is $bk\delta v_B + (s-1)\delta v_S = \delta(1 - v_S)$. Hence, the small state/district retains at least $1 - \delta(1 - v_S) \geq \delta v_S$. This implies that a small state proposer will always support his own proposal. Similarly, if the proposer is from a big state, then the most he will allocate to districts outside of his state is $(b-1)k\delta v_B + s\delta v_S = \delta(1 - kv_B)$, and so at least $1 - \delta(1 - kv_B) \geq \delta kv_B$ will be retained within the proposer's district. This is sufficient to earn the support of the upper house legislator and any single lower house legislator from that state. Moreover, if the proposer is

from the upper house, an equal division of the pie within the proposer's state is sufficient to earn the support of all lower house legislators. \square

Proof of Proposition 1 . The legislator supply constraints imply that:

$$\begin{aligned} \max \{0, M_L - bk - \mathbf{1}_S[t], M_U - b - \mathbf{1}_S[t]\} &\leq \sigma_t \leq \min \{s - \mathbf{1}_S[t]\} \\ \max \{0, M_L - s - l_t \mathbf{1}_B[t], k(M_U - s - \mathbf{1}_B[t])\} &\leq \beta_t \leq \min \{bk - l_t \mathbf{1}_B[t]\} \end{aligned}$$

Consider the first expression. The lower bound is the minimum number of small states (in addition to the proposer, if she is from a small state) who must be enticed into the coalition to satisfy the majority constraints in both houses, given the supply of big state legislators. The upper bound is the maximum number of small states who can be enticed into the coalition (in addition to the proposer, if she is from a small state), given the supply of small state legislators. The second expression is analogous for big state districts. I refer to these inequalities as the ‘feasibility constraints’.

Define $\sigma'_t = \max \left\{ \sigma \in \mathbb{Z} \mid \sigma < \frac{k(M_U - 1) - (M_L - l_t)}{k-1} \right\}$ and for a given $\sigma \in \mathbb{Z}$, let $\beta_t(\sigma) = \max \{k(M_U - 1 - \sigma), M_L - l_t - \sigma\}$. For a given σ , β_t is the minimum number of big state districts that must be included in the coalition to ensure that the majority constraints are satisfied in both chambers (ignoring legislator supply constraints). (NB — at this stage I do not insist that σ'_t be feasible. Indeed, σ'_t may be negative.) Clearly:

$$\beta_t(\sigma) = \begin{cases} k(M_U - 1 - \sigma) & \sigma \leq \sigma'_t \\ M_L - l_t - \sigma & \sigma > \sigma'_t \end{cases}$$

Let $C_t(\sigma) = v_S \sigma + v_B \beta_t(\sigma)$ be the cost for a type- t proposer to build a (σ, β_t) -coalition. Let $\Delta C_t(\sigma) = C_t(\sigma + 1) - C_t(\sigma)$. If $\Delta C_t(\sigma) < 0$, then there is strict incentive for the proposer to increase the number of small states in the coalition, and vice versa. From the

above expression for $\beta_t(\sigma)$, it is easily verified that $\Delta C_t(\sigma)$ satisfies:

$$\Delta C_t(\sigma) = \begin{cases} v_S - kv_B & \sigma < \sigma'_t \\ v_S - [k(M_U - 1 - \sigma'_t) - (M_L - l_t - \sigma'_t)]v_B & \sigma = \sigma'_t \\ v_S - v_B & \sigma > \sigma'_t \end{cases}$$

Let $\psi_t = k(M_U - 1 - \sigma'_t) - (M_L - l_t - \sigma'_t)$. It can easily be shown that $\psi_t = \text{mod}_{k-1}[kM_U - M_L] + 1$. (To see this, note that $(k-1)\sigma'_t < k(M_U - 1) - (M_L - l_t) \leq (k-1)(\sigma'_t + 1)$, which implies that $0 < \psi_t \leq k-1$. Additionally, $k - l_t \in \{0, k-1\}$ by construction.) Note that $0 < \psi < k$ implies that $\Delta^2 C_t(\sigma) \geq 0$, and so the marginal cost of adding one more small state to the coalition is weakly increasing.

Suppose $v < 1$ (i.e. $v_S < v_B$). Then $\Delta C_t(\sigma) < 0 \forall \sigma$, and so the proposer can decrease the cost of a coalition by adding another small state, whenever it is feasible to do so. Adding the feasibility constraints, the optimal coalition will contain $\sigma_t = \min\{\max\{M_L - l_t, M_U - 1\}, s - \mathbf{1}_S[t]\}$ and $\beta_t = \max\{0, M_L - s - l_t \mathbf{1}_B[t], k(M_U - s - \mathbf{1}_B[t])\}$.

Suppose $v > k$ (i.e. $v_S > kv_B$). Then $\Delta C_t(\sigma) > 0 \forall \sigma$, and so the proposer can decreasing the cost of a coalition by reducing the number of small states, whenever it is feasible to do so. Again adding the feasibility constraints, the optimal coalition will contain $\sigma_t = \max\{0, M_L - bk - \mathbf{1}_S[t], M_U - b - \mathbf{1}_S[t]\}$ and $\beta_t = \max\{M_L - l_t - \sigma_t, k(M_U - 1 - \sigma_t)\}$. Unless $0 \not\geq M_U - b > M_L - bk$ and $t = B^L$, this can be expressed more simply as $\beta_t^* = \min\{\max\{M_L - l_t, k(M_U - 1)\}, bk - l_t \mathbf{1}_B[t]\}$. (This is easily verified. If $\sigma_t = 0$, then there are sufficiently many big state districts to satisfy both majority constraints. Hence $\beta_t = \max\{M_L - l_t, k(M_U - 1)\} = \beta_t^*$. If $\sigma_t > 0$, then there are insufficiently many big state districts to satisfy both constraints. Hence $\beta_t^* = bk - l_t \mathbf{1}_B[t]$. Now, if the lower house majority constraint is binding, then the optimal coalition will use every available big state districts, and so $\beta_t = bk - l_t \mathbf{1}_B[t] = \beta_t^*$. However, if the upper house majority constraint

(only) is binding, then the optimal coalition will purchase every available big state senator, and so $\beta_t = (b - \mathbf{1}_B[t])k$. If $t \in \{S, B^U\}$, then $\beta_t = \beta_t^*$. By contrast, if $t = B^L$, then $\beta_t = (b - 1)k < bk - 1 = \beta_t^*$. When a lower house legislator from a big state is the proposer, he automatically gets the support of the upper house agent from his state - so the coalition needn't include the $k - 1$ other districts in his state. This problem does not arise when $t = S$ (obviously) or when $t = B^U$, since in the latter case, all k districts from the proposer's state are automatically in the coalition.)

Suppose $1 < v < \psi$ (i.e. $v_B < v_S < \psi v_B$). Then $\Delta C(\sigma) > 0$ if $\sigma > \sigma'_t$ and $\Delta C(\sigma) < 0$ if $\sigma \leq \sigma'_t$. Hence, in the unconstrained problem, the optimal coalition contains $\sigma'_t + 1$ small states (since, when $\sigma = \sigma_t$, it is still profitable to add one more small state to the coalition). Adding the feasibility conditions gives $\sigma_t = \min\{\max\{0, M_L - bk - \mathbf{1}_S[t], M_u - b - \mathbf{1}_S[t], \sigma'_t + 1\}, s - \mathbf{1}_S[t]\}$ and $\beta_t = \beta_t(\sigma'_t + 1)$. Suppose $\psi < v < k$ (i.e. $\psi v_B < v_S < kv_B$). Then $\Delta C(\sigma) > 0$ if $\sigma \geq \sigma_t$ and $\Delta C(\sigma) < 0$ if $\sigma < \sigma_t$. Hence, in the unconstrained problem, the optimal coalition contains σ'_t small states. Adding the feasibility conditions gives $\sigma_t = \min\{\max\{0, M_L - bk - \mathbf{1}_S[t], M_u - b - \mathbf{1}_S[t], \sigma'_t\}, s - \mathbf{1}_S[t]\}$ and $\beta_t = \beta_t(\sigma'_t)$.

Finally, if $v \in \{1, \psi, k\}$ then the optimal coalition is found by taking arbitrary mixtures of the adjacent coalitions. This follows since the correspondence $(\bar{\sigma}_t(v), \bar{\beta}_t(v))$ is upper hemicontinuous. (To see this, let $\{v_n\} \rightarrow v$ and let $\{(\sigma_n, \beta_n)\} \rightarrow (\sigma, \beta)$ be a sequence s.t. (σ_n, β_n) is optimal for v_n . Suppose (σ, β) is not optimal for v . Then, $\exists(\sigma', \beta')$ feasible s.t. $v_S\sigma' + v_B\beta' < v_B\sigma + v_S\beta - 3\varepsilon$. But for n large enough, $v_S^n\sigma_n + v_B^n\beta_n > v_S^n\sigma + v_B^n\beta - \varepsilon > v_S\sigma + v_B\beta - 2\varepsilon$. Moreover, for n large, $v_S\sigma' + v_B\beta' > v_S^n\sigma' + v_B^n\beta' - \varepsilon$. Then $v_S^n\sigma' + v_B^n\beta' < v_S^n\sigma_n + v_B^n\beta_n$, which contradicts the assumption that (σ_n, β_n) is optimal for v_n .) Take $v = 1$. Let (σ_1, β_1) be the optimal coalition whenever $v < 1$ and (σ_2, β_2) be the optimal coalition whenever $1 < v < \psi$. Then, by upper-hemicontinuity $\lim_{v \rightarrow 1^-} (\sigma_t(v), \beta_t(v)) = (\sigma_1, \beta_1) \in (\sigma_t(1), \beta_t(1))$. Similarly, $\lim_{v \rightarrow 1^+} (\sigma_t(v), \beta_t(v)) = (\sigma_2, \beta_2) \in (\sigma_t(1), \beta_t(1))$. Hence (σ_1, β_1) and (σ_2, β_2) are both optimal coalitions when $v = 1$, and so any mixture of these is also optimal. A similar

argument holds for $v = \psi$ and $v = k$. □

Lemma 2. *Let $(\sigma, \beta) \neq (\sigma', \beta')$ both be optimal coalitions for $v = v_0$ and let μ be the probability that a (σ, β) -coalition is chosen. Let $\phi(v_0, \mu) = \frac{bk - \delta\bar{\beta}}{s - \delta\bar{\sigma}}$, where $\bar{\sigma} = \mu\sigma + (1 - \mu)\sigma'$ and $\bar{\beta} = \mu\beta + (1 - \mu)\beta'$, and let $\phi = \phi(v_0, 1)$ and $\phi' = \phi(v_0, 0)$. Then, for every $\lambda \in [0, 1]$, there is a unique $\mu \in [0, 1]$ s.t. $\phi(v_0, \mu) = \lambda\phi + (1 - \lambda)\phi'$.*

Proof. Clearly, for $\lambda = 1$ (respectively $\lambda = 0$), $\mu = 1$ (respectively $\mu = 0$) satisfies the claim. Suppose $\lambda \in (0, 1)$. $\phi(v_0, \mu)$ is continuous in μ , since it is the ratio of two non-zero continuous functions. Since (σ, β) and (σ', β') are both optimal, it cannot be that $(\sigma, \beta) \geq (\sigma', \beta')$ or $(\sigma', \beta') \geq (\sigma, \beta)$. Accordingly, since $(\sigma', \beta') \neq (\sigma, \beta)$, then $\sigma > \sigma'$ iff $\beta < \beta'$. Hence $\phi \neq \phi'$. WLOG suppose $\phi > \phi'$. Let $\phi_\lambda = \lambda\phi + (1 - \lambda)\phi'$ and note that $\phi > \phi_\lambda > \phi'$ and ϕ_λ is strictly increasing in λ . Then, by the intermediate value theorem, there is some $\mu(\lambda) \in (0, 1)$ s.t. $\phi(v_0, \mu) = \phi_\lambda$, for each $\lambda \in (0, 1)$. Moreover, since ϕ_λ is strictly increasing, $\mu(\lambda)$ is unique. □

Proof of Proposition 2. Let $\Phi(v; P) = \frac{1-P}{P}\phi(v)$. It suffices to show that the conjectured v is a fixed point of $\Phi(v; P)$. The proof proceeds piecewise. First consider $v \in \mathbb{R} \setminus \{1, \theta, k\}$, so that the optimal coalitions are unique and so $\phi(v)$ is a singleton. As above, let $\phi(v) = \phi_1$ for $v < 1$, and define ϕ_2, ϕ_3 and ϕ_4 similarly. Since $\bar{\sigma}(v)$ is weakly decreasing in v and $\bar{\beta}(v)$ is weakly increasing, then $\phi(v)$ is weakly decreasing in v . Hence $\phi_1 \geq \phi_2 \geq \phi_3 \geq \phi_4$. To fix ideas, consider $v < 1$. Then $\Phi(v; Q) = \frac{1-P}{P}\phi_1$. Hence $v = \frac{1-P}{P}\phi_1$ is a fixed point of Φ as long as $\frac{1-P}{P}\phi_1 < 1$. But this implies that $P < \frac{\phi_1}{1+\phi_1}$. A similar argument is used for the remaining cases: $1 < v < \theta$, $\theta < v < k$ and $v > k$.

Now, suppose $v \in \{1, \theta, k\}$, so that $\phi(v)$ is set-valued. Again to fix ideas, consider $v = 1$. Define $\phi(v, \mu)$ as in Lemma 2. It suffices to find some $\mu \in [0, 1]$ s.t. $\frac{1-P}{P}\phi(1, \mu) = 1$. This requires that $P = \frac{\phi(1, \mu)}{1+\phi(1, \mu)}$. By Lemma 2 there is unique $\mu \in [0, 1]$ for each $\phi_\lambda \in [\phi_2, \phi_1]$ s.t. $\phi(1, \mu) = \phi_\lambda$. Since $\phi_2 \leq \phi_\lambda \leq \phi_1$, then $\frac{\phi_2}{1+\phi_2} \leq \frac{\phi_\lambda}{1+\phi_\lambda} \leq \frac{\phi_1}{1+\phi_1}$. Hence, a fixed point of Φ exists whenever $P \in \left[\frac{\phi_2}{1+\phi_2}, \frac{\phi_1}{1+\phi_1} \right]$. Moreover, since ϕ_λ is strictly increasing in μ , this fixed point is

unique. The proof for $v = \theta$ and $v = k$ is analogous.

Finally, I show that the expected equilibrium shares are unique. Suppose not. Then for some probability triple (P, α_B, α_S) , there exist v, v' with $v \neq v'$ such that both are fixed points of Φ . WLOG suppose $v < v'$. Since $\sigma(v)$ is decreasing in v and $\beta(v)$ is increasing in v , $\min \phi(v) \geq \max \phi(v')$. Then $\phi(v, \mu) \geq \phi(v', \mu')$ and so:

$$v = \frac{1-P}{P} \phi(v, \mu) \geq \frac{1-P}{P} \phi(v', \mu') = v'$$

which contradicts $v' > v$. □

Proof of Proposition 3. First, let P^u be the aggregate recognition probability under unicameralism, and note by construction $P^u = \frac{\alpha_B P}{\alpha_B P + \alpha_S (1-P)}$ and so $\frac{1-P^u}{P^u} = \frac{\alpha_S}{\alpha_B} \frac{1-P}{P}$. Let (P, α_B, α_S) be such that the equilibrium ratio satisfies $v^u < 1$. By Proposition 2, $v^u = \frac{\alpha_S}{\alpha_B} \frac{1-P}{P} \cdot \phi_1^u$. By similar argument, $v = \frac{1-P}{P} \phi(\alpha_B) \leq \frac{1-P}{P} \phi_1(\alpha_B)$, since ϕ is decreasing in v . Suppose $\alpha_S > \alpha_B \frac{\phi_1(\alpha_B)}{\phi_1^u}$. Then $v \leq \frac{1-P}{P} \phi_1(\alpha_B) < \frac{\alpha_S}{\alpha_B} \frac{1-P}{P} \phi_1^u = v^u < 1$. Hence $v < v^u$. Suppose instead that $v < v^u < 1$. Then $v = \frac{1-P}{P} \phi_1(\alpha_B) < \frac{\alpha_S}{\alpha_B} \frac{1-P}{P} \phi_1^u$, which implies that $\alpha_S > \alpha_B \frac{\phi_1(\alpha_B)}{\phi_1^u}$. This proves the first part of proposition.

To prove the second part, define the following sets which partition the space of feasible majority requirements:

$$\begin{aligned} A_1 &= \left\{ (M_L, M_U) \mid M_L < M_U < s + b \frac{s - (M_L - 1)}{s - \delta(M_L - 1)} \text{ and } M_L \leq s \right\} \\ A_2 &= \left\{ (M_L, M_U) \mid \frac{kM_U - M_L}{k-1} > s \text{ and } M_L > s, \text{ or } M_U \leq s \text{ and } M_U < M_L < s + k, \right. \\ &\quad \left. \text{or } M_L \leq s \ \& \ M_U > s + b \frac{s - M_L}{s - \delta M_L} \right\} \\ A_3 &= \left\{ (M_L, M_U) \mid M_L \geq s + k \text{ and } \frac{kM_U - M_L}{k-1} \leq s \right\} \end{aligned}$$

Let $\Delta = bk(\sigma_B - \hat{\sigma}_B) - s(\beta_S - \hat{\beta}_S) + \delta(\beta_S \hat{\sigma}_B - \sigma_B \hat{\beta}_S)$, where $\sigma_B = \alpha_B \sigma_{B^L} + (1 - \alpha_B) \sigma_{B^U}$. Note that this expression simplifies to $\Delta = (bk - \delta \beta_S)(\sigma_B - \hat{\sigma}_B)$ if $\beta_S = \hat{\beta}_S$ and $\Delta = -(s - \delta \sigma_B)(\beta_S - \hat{\beta}_S)$ if $\sigma_B = \hat{\sigma}_B$. It is easily verified that $\Delta > 0$ iff $\phi_1(\alpha_B) > \hat{\phi}_1$ (which

implies that $\frac{\phi_1(\alpha_B)}{\phi_1} > 1$). To show that $\alpha_S \leq \alpha_B$ it suffices to show that $\Delta \leq 0$.

Suppose $(M_L, M_U) \in A_1$. Then $M_L < M_U$ and $M_L \leq s$. If $M_U \leq s$, then $\sigma_B = M_U - 1 > M_L - 1 = \hat{\sigma}_B$ and $\beta_S = \hat{\beta}_S = 0$, and so $\Delta = bk(M_U - M_L) > 0$. If instead, $s < M_U < s + b\frac{s-(M_L-1)}{s-\delta(M_L-1)}$, then $\sigma_B = s > M_L - 1 = \hat{\sigma}_B$, and $\beta_S = k(M_U - s) > 0 = \hat{\beta}_S$. Then

$$\begin{aligned}\Delta &= bk(s - (M_L - 1)) - sk(M_U - s) + \delta k(M_U - s)(M_L - 1) \\ &= k[b(s - (M_L - 1)) - (M_U - s)(s - \delta(M_L - 1))] \\ &> 0\end{aligned}$$

where the last inequality is implied by $M_U < s + b\frac{s-(M_L-1)}{s-\delta(M_L-1)}$.

Suppose $(M_L, M_U) \in A_2$. There are three possibilities. (1) If $\frac{kM_U - M_L}{k-1} > s$ and $M_L > s$ (which implies that $M_U > s$), then $\sigma_B = \hat{\sigma}_B = s$ and $\beta_S = k(M_U - s) > M_L - s = \hat{\beta}_S$. Hence $\Delta = -(s - \delta\sigma_B)(\beta_S - \hat{\beta}_S) < 0$. (2) Suppose $M_U \leq s$ and $M_U < M_L \leq s + k - 1$. Since $M_U \leq s$, then $\beta_S = \hat{\beta}_S$, and since $M_U < M_L$, then $\sigma_{B^L} = \hat{\sigma}_B$. Hence $\Delta = (bk - \delta\beta_S)(1 - \alpha)(\sigma_{B^U} - \sigma_B)$ and $\Delta < 0$ if $\sigma_{B^U} < \hat{\sigma}_B$. Since $M_L \leq s + k - 1$ (i.e. $M_L - k < s$), then $\sigma_{B^U} \in \{M_L - k, M_U - 1\}$. Furthermore, by the assumptions on M_L and M_U , $M_L - k < \min\{M_L - 1, s\}$ and $M_U - 1 < \min\{M_L - 1, s\}$. Hence $\sigma_{B^U} < \min\{M_L - 1, s\} = \hat{\sigma}_B$, and so $\Delta < 0$. (3) Suppose $M_L \leq s$ and $M_U > s + b\frac{s-(M_L-1)}{s-\delta(M_L-1)}$. Then $\beta_S = k(M_U - s) > 0 = \hat{\beta}_S$ and $\sigma_B = s > M_L - 1 = \hat{\sigma}_B$. Hence

$$\begin{aligned}\Delta &= bk(s - (M_L - 1)) - sk(M_U - s) + \delta k(M_U - s)(M_L - 1) \\ &= k[b(s - (M_L - 1)) - (M_U - s)(s - \delta(M_L - 1))] \\ &< 0\end{aligned}$$

where the last inequality is implied by $M_U > s + b\frac{s-(M_L-1)}{s-\delta(M_L-1)}$. Hence $\Delta < 0$ whenever $(M_L, M_U) \in A_2$. \square

References

- Ansolabehere, Stephen, Jr. Snyder, James M. and Michael M. Ting. 2003. "Bargaining in Bicameral Legislatures: When and Why Does Malapportionment Matter?" *The American Political Science Review* 97(3):pp. 471–481.
- Banks, J. and J. Duggan. 2006. "A general bargaining model of legislative policy-making." *Quarterly Journal of Political Science* 1(1):49–85.
- Banks, Jeffrey S. and John Duggan. 2000. "A Bargaining Model of Collective Choice." *The American Political Science Review* 94(1):pp. 73–88.
- Baron, David P. and John A. Ferejohn. 1989. "Bargaining in Legislatures." *The American Political Science Review* 83(4):pp. 1181–1206.
- Central Intelligence Agency. 2011. "The CIA World Factbook."
- Chen, Jowei. 2010. "The effect of electoral geography on pork barreling in bicameral legislatures." *American Journal of Political Science* 54(2):301–322.
- Cutrone, M. and N. McCarty. 2006. "Does bicameralism matter?" *Oxford Handbook of Political Economy*. Oxford: OUP pp. 180–195.
- Diermeier, Daniel and Roger B. Myerson. 1999. "Bicameralism and Its Consequences for the Internal Organization of Legislatures." *The American Economic Review* 89(5):pp. 1182–1196.
- Hammond, Thomas H. and Gary J. Miller. 1987. "The Core of the Constitution." *The American Political Science Review* 81(4):pp. 1155–1174.
- Hauk, William R, Romain Wacziarg et al. 2007. "Small states, big pork." *Quarterly Journal of Political Science* 2(1):95–106.

- Heitshusen, Valerie. 2001. "The Allocation of Federal Money to House Committee Members Distributive Theory and Policy Jurisdictions." *American Politics Research* 29(1):79–97.
- Heller, William B. 1997. "Bicameralism and budget deficits: the effect of parliamentary structure on government spending." *Legislative Studies Quarterly* pp. 485–516.
- Kalandrakis, T. 2006. "Proposal rights and political power." *American Journal of Political Science* 50(2):441–448.
- Kalandrakis, Tasos. 2004. "Bicameral Winning Coalitions and Equilibrium Federal Legislatures." *Legislative Studies Quarterly* 29(1):pp. 49–79.
- Knight, Brian. 2005. "Estimating the value of proposal power." *The American Economic Review* 95(5):1639–1652.
- Knight, Brian. 2008. "Legislative representation, bargaining power and the distribution of federal funds: Evidence from the US congress." *The Economic Journal* 118(532):1785–1803.
- Kuhn, Max and Kjell Johnson. 2013. *Applied predictive modeling*. Springer.
- Lazarus, Jeffrey and Amy Steigerwalt. 2009. "Different Houses: The distribution of earmarks in the US House and Senate." *Legislative Studies Quarterly* 34(3):347–373.
- Lee, Frances E. 2000. "Senate representation and coalition building in distributive politics." *American Political Science Review* 94(01):59–72.
- Lee, Frances E. 2004. "Bicameralism and geographic politics: Allocating funds in the House and Senate." *Legislative Studies Quarterly* 29(2):185–213.
- Levitt, Steven D and James M Snyder Jr. 1995. "Political parties and the distribution of federal outlays." *American Journal of Political Science* pp. 958–980.
- Lijphart, A. 1999. *Patterns of democracy*. Yale University Press.

- McCarty, N. 2000. "Proposal rights, veto rights, and political bargaining." *American Journal of Political Science* 44(3):506–522.
- Shepsle, Kenneth A, Robert P Van Houweling, Samuel J Abrams and Peter C Hanson. 2009. "The senate electoral cycle and bicameral appropriations politics." *American journal of political science* 53(2):343–359.
- Thorpe, Rebecca U. 2010. "The role of economic reliance in defense procurement contracting." *American Politics Research* 38(4):636–675.
- Tsebelis, G. and J. Money. 1997. *Bicameralism*. Cambridge Univ Pr.
- Wilson, Rick K. 1986. "An empirical test of preferences for the political pork barrel: District level appropriations for river and harbor legislation, 1889-1913." *American Journal of Political Science* pp. 729–754.