NICEP Working Paper: 2023-01

Collective Screening

Vincent Anesi
Peter Buisseret

Nottingham Interdisciplinary Centre for Economic and Political Research

School of Economics, The University of Nottingham, Law & Social Sciences Building, University Park, Nottingham, NG7 2RD
ISSN 2397-9771

Collective Screening*

Vincent Anesi[†] Peter Buisseret[‡]

March 19, 2023

Abstract

We study a dynamic principal-agent model in which the principal is a group whose members hold heterogeneous and evolving values from an agreement with the agent. Learning about the agent's private information reduces the principals' conflicts over their joint offer, mitigating a principal's losses if she is not decisive over future offers. As a consequence, a principal in a group prefers to screen the agent more aggressively than a single principal. We study the dynamics of the principals' collective choice, and obtain conditions under which decisive members of the group successively trade away their decision-making authority, leading inexorably to the concentration of negotiation power in the hands of a single principal.

*We are grateful to Avi Acharya, Daniel Barron, Georgy Egorov, Lars Ehlers, Daniel Garrett, Bard Harstad, Ángel Hernando-Veciana, Nicolas Klein, Antoine Loeper, Ryan Tierney, Rafael Treibich, and many seminar audiences for helpful comments and suggestions. Buisseret gratefully acknowledges the hospitality of Örebro University School of Business.

[†]Department of Economics and Management, University of Luxembourg, *Email*: vincent.anesi@uni.lu

[‡]Department of Government, Harvard University, *Email*: pbuisseret@fas.harvard.edu

1. Introduction

This paper revisits the classical screening problem in which a principal makes repeated offers to an agent with persistent private information. Our key departure is that the principal is a *group*, rather than an individual: the principals jointly propose an offer, and if the offer is accepted the principals enter into a collective agreement with the agent. There are many real-world examples of groups that collectively bargain with an external party. These examples include labor unions negotiating with a firm's management; shareholders contracting with corporate boards; consumer cooperatives negotiating with vendors; suppliers' associations bargaining with firms; and academic departments negotiating with a central administration.

We focus on how heterogeneity amongst the principals shapes their conflicts over the agent's offer. This heterogeneity introduces novel static and dynamic conflicts. At any point during the relationship, different principals have different short-run priorities for negotiations. Principals with relatively high instantaneous benefits from an agreement prefer a generous offer that the agent is sure to accept. Principals with lower instantaneous benefits prefer to gamble on the prospect of securing an agreement with a less generous offer. These heterogeneous short-run priorities generate static conflicts between members of the group over the agent's offer.

Static conflicts between the principals are also augmented by novel dynamic conflicts. In a single-principal context, the key dynamic tension in negotiations arises from the agent's attempt to shape the principal's learning. For example, an agent may be tempted to reject an offer that yields a positive instantaneous net benefit from acceptance, because doing so reveals that her cost of conceding is not too high and thus induces less generous future offers. When the principal is instead a group, the agent and each principal must also account for how the *other* principals' beliefs evolve in response to offers and acceptance decisions.

The principals therefore face the problem of collectively deciding on their offer. This offer must emerge from a collective choice procedure, such as voting. Different voting rules empower different principals to decide offers, and therefore determine the objectives of the group, as a whole.[1] To the extent that these institutions reflect the congealed preferences of their decisive

[1] The difficulty of specifying a union's goals with the 'correct' voting rule represented a "major stumbling

members (Riker, 1980), they may evolve over time along with those preferences. The goals of our paper are two-fold: first, we want to understand how conflicts amongst the principals shape their screening problem. Second, we want to understand how the principals structure their internal decision-making in order to face the agent over the life of the relationship. More broadly, we want to address the question of how within-group bargaining impacts bargaining between groups.

To address these goals, our model studies a long-run (infinite-horizon) relationship between a group of principals and an agent (Acharya and Ortner, 2017). In each period, the principals can collectively make a demand to the agent in exchange for a policy concession. The agent may either concede to the demand, or refuse. The principals derive heterogeneous benefits from securing an agreement with the agent, but they are uncertain about the agent's cost of conceding. We assume that an agreement is always efficient.

We model the principals' collective choice of the agent's offer as an amendment agenda game (Duggan 2006, Austen-Smith and Banks 2005). This game is governed by a *procedure*, which specifies the order in which principals can make proposals, and the voting rule used to select the winning alternative. We allow for deterministic or random recognition rules, and a wide array of voting rules, including quotas, oligarchies, and rules with veto rights. At the start of every period, the principals inherit the previous period's standing procedure. Before facing the agent, however, they may adopt a new procedure. For example, they could amend a unanimous voting rule to a simple majority with veto rights, or change the order in which principals are recognized. This procedural choice is also modeled as an amendment agenda game executed under the (inherited) standing procedure.

We first study how different collective choice rules affect screening incentives. Relative to a benchmark with a single decisive group member—or equivalently in a group setting, the dictatorship of a single principal—we find that under *any* procedure the collective choice aspect

block" to analysis in models of collective bargaining (Crawford and Blair, 2002). Some authors presumed unions maximize a decisive member's welfare–for example, a median voter—despite the fact that some specifications of the economic environment failed to yield a transitive majority preference relation (Farber, 1978). We show that under a richer array of voting rules, the core of the collective choice problem may alternatively be too permissive to make concrete predictions, motivating our non-cooperative formulation of the collective choice process.

encourages principals to screen the agent more aggressively. The reason is that concentrating the principals' common belief over a smaller set of agent types reduces the principals' scope for disagreement over future offers. Less future disagreement limits a principal's future losses in the event she is not decisive in future periods. Because our result holds across all collective choice procedures, it is relevant to any principal-agent relationship in which the principal is a group.

We then study how the group's organization evolves in the long run. We show that decisive coalitions of relatively impatient principals successively trade away their decision-making authority to other principals, eventually concentrating power entirely in the hands of a single group member. This phenomenon arises when (1) a subset of non-decisive high-benefit principals prioritize a pooling offer the agent is sure to accept, (2) a 'marginal' member prefers to separate some agent types under the status quo rule, but would prefer to pool all types if she were decisive over future offers, and (3) these principals, together, are decisive under the inherited rule. In that context, the non-decisive members sacrifice their own future decision-making authority in order to secure their preferred outcome, today. Complete concentration need not happen immediately: we illustrate the dynamics of how power increasingly and inexorably accumulates in the long run.

Our finding resonates with Robert Michels famous dictum that organizations invariably drift towards the concentration of decision-authority in the hands of an ever-smaller number of individuals (Michels, 1959). There are many accounts of how principals should centralize or delegate authority to leverage an agent's expertise (Dessein, 2002), to adapt decisions to local conditions (Liu and Migrow, 2022), or to improve coordination (Rantakari, 2008). However, organizational theory has had little to say about the oligarchic tendencies Michels identified in his 'iron law'. Williamson (1988) argues that "[t]he incentive literature makes no provision whatsoever for the possibility that oligarchy is a predictable process outcome" (p. 87). We highlight a new channel for the concentration of power: an organization's strategic interaction with an external agent.

We are not the first to integrate the principal-agent framework with collective decision-making (e.g., Laffont, 2000 and Grossman and Helpman, 2001). Yet existing theoretical work on delegation with multiple principals exclusively focuses on common agency environments (Bernheim and Whinston, 1986, Grossman and Helpman, 1994, and Lefebvre and Martimort, 2020) in which the principals non-cooperatively offer distinct and competing contracts to a single agent, or multiple agents (Prat and Rustichini, 2003). While appropriate for some applications, there are many others in which principals collectively sign contracts—for example, a union pursuing a wage settlement that applies to all workers in a given industry. It is often infeasible or even prohibited for a member of the group to individually contract with an agent on his or her own behalf (Tommasi and Weinschelbaum, 2007). In models of dynamic electoral accountability (e.g., Duggan and Martinelli, 2020) multiple principals (voters) contract with an agent (a politician). These papers nonetheless presume a representative voter, thereby suppressing heterogeneity amongst principals.

Our focus on how the evolution of endogenous collective decision-making rules follows Lagunoff (2009), Acemoglu, Egorov and Sonin (2012, 2015, 2021), and Diermeier and Vlaicu (2011) by characterizing self-enforcing institutions when reform is governed by existing rules.[2] At a technical level, the sequences of offers made to the agent in our noncooperative equilibria constitute Markov voting equilibria à la Roberts (2015) or Acemoglu, Egorov and Sonin (2015). While the core is generally too permissive to make concrete predictions for some voting rules—such as large voting quotas—we show that Duggan (2006)'s amendment agenda game serves as a natural and effective approach to refine Markov voting equilibrium to a unique prediction under any procedure.

Our work relates more distantly to the literature on experimentation, e.g., Strulovici (2010), Anesi and Bowen (2021) and Bowen, Hwang and Krasa (2022); Freer, Martinelli and Wang (2020) survey recent contributions. Nevertheless, the strategic interaction with a privately informed agent in our model yields a learning technology that is proper to the dynamic screening problem, and fundamentally different from the experimentation literature. In those papers, a group collectively chooses between a risky reform and a safe status quo in a Poisson bandit framework with exogenous learning costs. Relative to a single-experimenter benchmark, individuals have insuf-

[2] Only institutions, rather than offers, persist across periods. This distinguishes our framework from those in which agreements reached today are the status quo in future negotiations, e.g., Bowen, Chen, Eraslan and Zápal (2017), Buisseret and Bernhardt (2017), Anesi and Duggan (2018), Dziuda and Loeper (2018), and Nunnari (2021), to cite a few of the most recent contributions —Eraslan, Evdokimov and Zápal (2022) provide an extensive survey of that literature.

ficient incentives to learn in a group context.[3] In our setting, the principals collectively determine incentive provision by choosing policy concessions to the agent, which in turn determine both the extent and the (endogenous) costs of learning. Proposition 1 shows that relative to a single-principal benchmark, collective principals have *excessive* incentives to learn in a group context.

2. Two-Period Example

Basic Elements. We consider a two-period interaction between five principals, $N \equiv \{1, \ldots, 5\}$, and an agent. In each period $t = 1, 2$, the principals can collectively make a demand from the agent in return for a policy concession, $x^t \in [0, 1]$. The agent accepts the demand ($a^t = 1$) or rejects it ($a^t = 0$). If the principals do not make an offer ($x^t = \varnothing$), a status-quo policy of zero is implemented.

Principal $i \in N$'s period-t payoff is $a^t \left[b_i^t - x^t\right]$, where b_i^t is a stochastic benefit drawn at the start of every period from a c.d.f. F that is continuous and has full support on $[\underline{b}, \overline{b}]$. The realization $b^t = (b_1^t, \ldots, b_5^t)$ is publicly observed. The agent's period-t payoff is $a^t \left[x^t - c\right]$, where c is her privately observed cost from accepting the principals' demand. The cost is drawn at the outset from $\{c_L, c_H\}$ and persists across both periods, with $\Pr(c = c_L) = p \in (0, 1)$, and $c_H < \underline{b}$. Players share a common discount factor $\delta \in (0, 1)$, and maximize average discounted payoffs.

Collective Choice. In each period $t \in \{1, 2\}$, after the principals' period-t benefits are realized, they collectively vote an offer to the agent, x^t. The voting rule is modeled as a collection $\mathcal{D} \subseteq 2^N \setminus \{\emptyset\}$ of decisive coalitions that we only restrict to be monotonic (e.g., $C \in \mathcal{D}$ and $C \subseteq C'$ imply $C' \in \mathcal{D}$) and proper ($C, C' \in \mathcal{D}$ implies $C \cap C' \neq \emptyset$) — e.g., Austen-Smith and Banks (1999).

We also allow the principals to select the voting rule $\mathcal{D}$ that they use to determine offers in each period. They choose this rule in period 1, after their benefits are realized and before they make their initial offer to the agent. They vote over the rule using the (exogenous) inherited rule $\mathcal{D}^0$, which we presume to be simple majority.[4] The timing is described in Figure 1.

[3] Gieczewski and Kosterina (2020) obtain excessive experimentation in a setting where members can unilaterally take a safe outside option (i.e., exit).

[4] In the sequel the principals can select a new rule in every period.

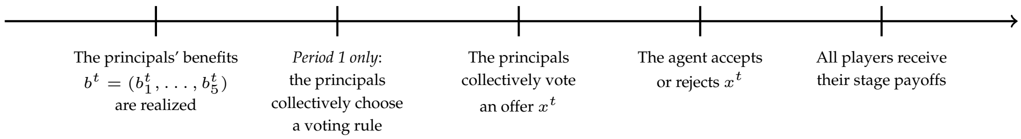

Figure 1 – Timing in each period $t = 1, 2$.

Equilibrium. For this illustrative example, we assume that the principals randomly select their offers (uniformly) from the *core* of the voting rule $\mathcal{D}$ — that is, the set of offers that are undefeated in pairwise voting using $\mathcal{D}$. We impose this selection while still applying standard sequential-rationality and belief-consistency conditions. Our uniform selection is innocuous, and purely for exposition. In the sequel, we characterize the perfect Bayesian equilibria of a fully-fledged noncooperative model of collective decision making among the principals. We further show how the principals endogenously select offers from the dynamic core in every equilibrium of the general model.

Analysis. We start with the second period. The type-c agent accepts the principals' period-2 offer x^2 if and only if $x^2 \geq c$. On the path, the principals' common belief is either the prior, or degenerate. Define $b^* \equiv \frac{c_H - pc_L}{1-p}$. If the principals learned that the agent's cost is $c \in \{c_L, c_H\}$, they unanimously prefer to offer $x^2 = c$, which the agent accepts. Under the prior belief, instead, principal i with benefit b_i's preferred offer is

$$x^2(b_i) \equiv \begin{cases} c_H & \text{if } b_i \geq b^*, \\ c_L & \text{otherwise.} \end{cases}$$

Henceforth, we assume $\underline{b} < b^* < \bar{b}$. Thus, for any (equilibrium) beliefs of the principal, they select from at most two possible offers to the agent.

Given period-2 belief p^2 and voting rule $\mathcal{D}$, the *core* of the principals' collective choice problem in period-2 is the set of offers that are undefeated in pairwise voting. We denote this set $\mathcal{K}(p^2, \mathcal{D})$: it is non-empty, and may contain either one or two offers. If the core contains two elements, our selection presumes that each is equally likely to be chosen. This selection generates a unique equilibrium outcome in period 2 for any belief and voting rule.

We therefore turn to period $t = 1$. Let $\overline{x}^H$ denote a period-1 *pooling* offer that both types accept, and $\overline{x}^L$ denote a *separating* offer that only the low-type agent accepts. Then, $\overline{x}^L$ solves:

$$\overline{x}^L - c_L + \delta \times 0 = 0 + \delta(c_H - c_L).$$

If the agent accepts the offer, she reveals her type is c_L, and receives a payoff of zero at period 2. If she rejects the offer, the principals infer that her type is c_H, and she receives the pooling offer at period 2. Similarly, $\overline{x}^H$ solves:

$$\overline{x}^H - c_H + \delta \times 0 = 0.$$

To see why, recognize that if the period-2 offer is c_L, the high-cost agent rejects and receives zero; if the period-2 offer is c_H, the agent's payoff is zero. Notice that the low-cost agent also accepts this offer: if she rejects, her second-period payoff is bounded above for all possible beliefs of the principals by $\delta(c_H - c_L)$, which is strictly less than her payoff from accepting c_H, today.

We conclude that, in an equilibrium, the principals face a period-1 collective choice between two possible offers—one for each agent-type in the support of the principals' beliefs. As is standard, the highest type's dynamic incentive constraint coincides with her static constraint, and the low type extracts an information rent. The information rent does not depend on the principals' voting rule because the agent has only two possible types, implying that the principals' belief is degenerate after the agent either accepts or rejects the separating offer. As a consequence, the principals unanimously agree on their preferred period-2 offer after a period-1 separating offer, and the voting rule does not impact the agent's period-2 offer.

Collective versus Individual Screening. We now study how different voting rules shape a principal's induced preferences over offers in period 1. Principal i's continuation value from a period-1 separating offer is:

$$W^{\text{sep}} \equiv \mathbb{E}[b_i] - pc_L - (1-p)c_H.$$

When the principals' beliefs about the agent are degenerate, they unanimously agree on their preferred period-2 offer. As a consequence, the continuation value from separation does not depend on the voting rule. Matters are different when the principals are uncertain about the agent's type. Let $\tau(b, \mathcal{D})$ denote the probability that the period-2 offer is c_H when the principals hold prior belief p, the benefits realization is b, and the voting rule is $\mathcal{D}$. Principal i's continuation value from a period-1 pooling offer is therefore:

$$W^{\text{pool}}(\mathcal{D}) \equiv \int_b \Big[\tau(b,\mathcal{D})(b_i - c_H) + \big[1 - \tau(b,\mathcal{D})\big]p(b_i - c_L)\Big]\, dF(b),$$

where $F(\cdot)$ is the joint distribution of the benefits profile $b = (b_1, \ldots, b_5)$. Principal i with period-1 benefit b_i therefore prefers the separating offer if and only if

$$b_i \leq \frac{1}{1-p}(c_H - p\overline{x}^L) + \frac{\delta}{1-p}\big[W^{\text{sep}} - W^{\text{pool}}(\mathcal{D})\big] \equiv \beta(\mathcal{D}).$$

Our goal is to compare this threshold under two classes of voting rules. A voting rule $\mathcal{D}$ is a *dictatorship of principal* i if:

$$\mathcal{D} = \{S \subseteq N : S \ni i\} \equiv \mathcal{D}^i.$$

Since the first period incentive constraints do not depend on the voting rule, and recalling $b^* \equiv \frac{c_H - pc_L}{1-p}$, we have that for any $\mathcal{D}$:

$$\beta(\mathcal{D}) - \beta(\mathcal{D}^i) = \delta(1-p)\int_b \big[\tau(b,\mathcal{D}^i) - \tau(b,\mathcal{D})\big]\,(b_i - b^*)\; dF(b). \tag{1}$$

Setting aside the constant, and recognizing that $\tau(b, \mathcal{D}^i)$ takes the value 1 if $b_i \geq b^*$, and zero otherwise, for any $\mathcal{D} \neq \mathcal{D}^i$, the difference (1) is:

$$\int_{\{b\,:\,b_i \geq b^*\}} \big[1 - \tau(b,\mathcal{D})\big]\,(b_i - b^*)\; dF(b) - \int_{\{b\,:\,b_i < b^*\}} \tau(b,\mathcal{D})\,(b_i - b^*)\; dF(b) > 0. \tag{2}$$

It is trivial the inequality holds, weakly. To see that the inequality is strict, recognize that since $\underline{b} < b^* < \overline{b}$, there is a positive probability realization of benefits in which either (1) all principals other than i prefer separation at date 2, but i favors pooling: $b_j < b^* < b_i$ for all $j \neq i$, or (2) all principals other than i prefer the pooling offer at date 2, but i favors separation: $b_i < b^* < b_j$ for all $j \neq i$. For these benefits realizations, i's losses from any $\mathcal{D} \neq \mathcal{D}^i$ are strictly positive. We therefore have the following observation.

Result 1. *For any $\mathcal{D} \neq \mathcal{D}^i$: $\beta(\mathcal{D}) > \beta(\mathcal{D}^i)$. In words: a principal's incentive to separate the agent in period 1 is strictly higher under any voting rule $\mathcal{D}$ than under her dictatorship $\mathcal{D}^i$.*

A collective principal prefers to screen the agent more aggressively than a singular principal. The reasoning is that screening the agent reduces the principals' future conflicts, which matters to today's principal when she is not assured of being decisive over future offers.

Suppose principal i's period-2 benefit is low (i.e,. $b_i^2 < b^*$), but that period's decisive principal j has a high benefit, $b_j^2 > b^*$. Under the prior, high-benefit principal j prioritizes agreement with the agent in period 2 by making the pooling offer c_H. This offer is excessively generous, from i's perspective. A period-1 separating offer may reveal that the agent has a low cost, c_L. This reduces high-benefit j's most preferred offer, to low-benefit i's advantage.

Conversely, suppose principal i's period-2 benefit is high (i.e., $b_i^2 > b^*$), but that period's decisive principal j has a low benefit, $b_j^2 < b^*$. Under the prior, low-benefit j prefers to gamble that the agent has a low cost of agreement by making the separating offer c_L. This offer risks that no agreement is reached the agent, which i prioritizes. Revealing that the agent has a high cost leads a future low-benefit decisive principal *not* to gamble with a low offer in period 2, to the advantage of high-benefit principal i.

Concentrating Power. Recall that the principals can select the voting rule governing how they make offers to the agent in each period. The principals select the rule in period 1, after their initial benefits are realized and before they vote their initial offer to the agent. They choose the rule under the status quo voting rule, $\mathcal{D}^0$, which we presume to be simple majority. We show that there is a positive probability that a decisive coalition of principals voluntarily cedes decision-making power, and opts to concentrate authority in a minority of principals—possibly, a single principal.

To see how this might arise, let $\overline{\beta}$ denote a principal's smallest possible pooling threshold under the prior:

$$\overline{\beta} = \min\left\{\beta(\mathcal{D}) : \mathcal{D} \neq \mathcal{D}^i\right\}. \tag{3}$$

Assume $\underline{b} < \overline{\beta}$, so that the costs of separating the agent are so expensive that every principal prefers to pool in the first period. Consider the positive probability event—illustrated in Figure 2—in which the benefits realization b^1 is such that:

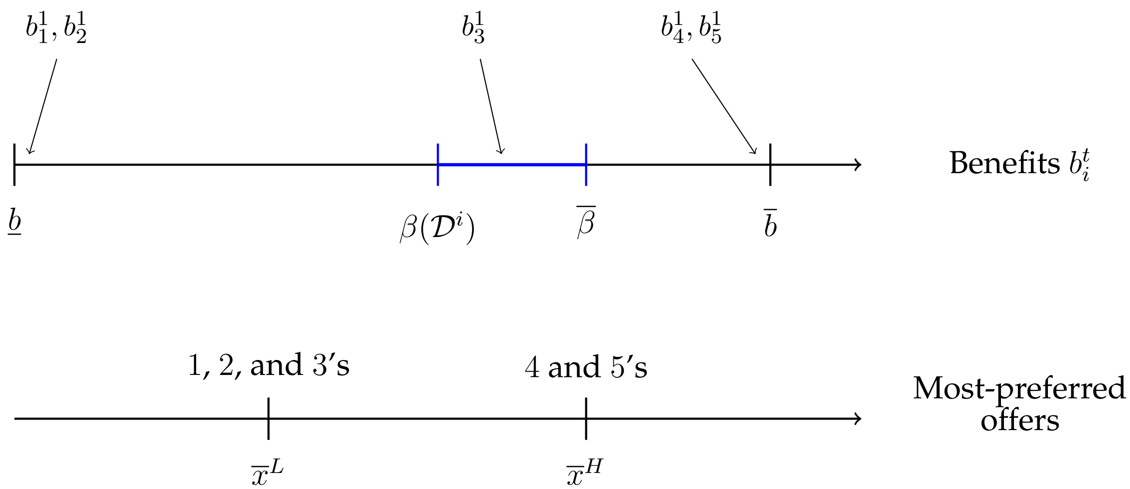

Figure 2 – The realization of principals' period-1 benefits described in text.

(i) b_1^1 and b_2^1 lie in a neighborhood of $\underline{b}$,

(ii) b_4^1 and b_5^1 lie in a neighborhood of $\overline{b}$, and

(iii) b_3^1 lies in $(\beta(\mathcal{D}^i), \overline{\beta})$.

Part (i) states that principals 1 and 2 prefer to screen the agent, but part (ii) states that principal 4 and 5 prefer the pooling offer. Part (iii) states that principal 3 prefers the pooling offer if she is a dictator; under any other voting rule, she prefers to screen the agent (*Result 1*). It follows that under a simple majority voting rule, the principals make the separating offer in period 1.

If δ is small enough—or if $\overline{b}$ is large enough—high-benefit principals 4 and 5 prioritize an agreement with the agent in period 1. Under the inherited simple majority rule, they cannot secure the pooling offer. Is there another voting rule that (1) guarantees the pooling offer will be made in period 1, and that (2) a majority of principals would prefer to simple majority? The answer is *yes*: a majority of principals strictly prefer a dictatorship of principal 3, $\mathcal{D}^3$, to simple majority.

To see why, recognize that since principals 4 and 5 prioritize agreement today, they strictly benefit from any change in the voting rule that guarantees a period-1 pooling offer. Since $\underline{b} < b_3^1 < \beta(\mathcal{D}^3)$, principal 3 favors the pooling offer if and only if she has sole authority to decide the period2 offer. And, principal 3 is trivially better off in both periods 1 and 2 when she is made a dictator. We obtain that a decisive coalition of today's principals—3, 4 and 5—strictly prefer

3's dictatorship to any voting rule that does not induce the pooling offer with probability one.

We conclude that for δ not too large, after this benefits realization, a voting rule lies in the core of the principals' collective choice at the start of period 1 *only if* it induces the pooling offer in that period. A dictatorship is not the only rule that achieves this, however. Recognizing the inevitability of a period-1 pooling offer, principals 1 and 2 could offer principals 4 or 5 an alternative procedure that establishes this commitment: namely, an oligarchy of principals 4 and 5: $\mathcal{D} = \{S \subseteq N : S \supseteq \{4,5\}\}$.

Besides a dictatorship or an oligarchy, no other voting rule guarantees the pooling offer, and thus no other voting rule commands the support of a majority.

Result 2. *If players care enough about period-1 outcomes, i.e., if δ is not too large, then there is a positive probability realization of benefits b^1 such that the only voting rules that belong to the core of the principals' collective choice are:*

1. *an oligarchy: for some $i, j \in N$, $\mathcal{D} = \{S \subseteq N : S \supseteq \{i,j\}\}$, or*
2. *a dictatorship: for some $i \in N$, $\mathcal{D} = \mathcal{D}^i$.*

Our two-period model yields two main insights. First, collective principals have stronger incentives to screen the agent, relative to singular (i.e., dictator) principals. Second, the principals may choose to concentrate the authority to select the group's proposal amongst a subset of principals

The rest of the paper extends these insights to an infinite horizon model with any (finite) number of principals and agent types. We allow the principals to reform their collective decision-making procedures at the start of every period, prior to their negotiation with the agent. We also assume that the agent's type is re-drawn with positive probability in every period, ensuring that there is always scope for learning in future periods.

We first address the robustness of Result 1 that collective principals have greater incentives to screen the agent than a single principal. In our two-type example, a separating offer fully reveals the agent's cost and fully eliminates disagreement amongst the principals. With more possible types, there are many partially-separating offers that leave residual uncertainty about the agent's cost, and thus also leave scope for disagreement amongst the principals over future offers. As a consequence, the principals' continuation values from (partially) separating offers vary with the collective choice rule, and the agent's incentive constraints associated with separating offers may also vary with the collective choice rule. The reason is that the agent must account for how information that she reveals today shapes future offers—which vary with the principals' choice rule—and thus her foregone information rent from revealing information about her cost. We verify that any wedge between a principal's incremental benefit from learning the agent's type versus any associated incremental incentive costs remains positive across different choice procedures.

The sequel also extends our substantive finding about the concentration of decision-making authority in Result 2. Recall that in our two-period example the principals can make at most one procedural reform decision; if, instead, the principals could reform their procedures more frequently, would the concentration of power stop, or would it continue, indefinitely? We provide a strong answer to this question by showing that *any* equilibrium sequence of procedures converges to the dictatorship of a single principal almost surely. That is: over time, the concentration of authority becomes absolute.

The inevitability of dictatorship will derive from three features of our model. First, the possibility that the agent's type is re-drawn in every period means that there is always a residual conflict of interest amongst the principals. Second, today's collective decision-making procedure is chosen under the inherited procedure from the previous period, which renders dictatorship absorbing. Third, we focus on settings where agents care enough about short-run outcomes—i.e., they have relatively low discount factors. This implies that principals with a large instantaneous benefit from agreement prioritize the pooling over any offer that reveals information about the agent's type but risks rejection.

Finally, our uniform random selection from the core even in the two-period model highlights how the core may be too permissive to make concrete predictions for some voting rules, such as large quotas. Rather than imposing an arbitrary selection, we model the principals' negotiations as an amendment agenda game (Duggan 2006, Austen-Smith and Banks 2005). The sequences of offers made to the agent in our noncooperative equilibria constitute Markov voting equilibria à la Roberts (2015) or Acemoglu, Egorov and Sonin (2015), and we show that the amendment agenda game refines Markov voting equilibrium to a unique prediction under any collective choice procedure.

3. Model

Main elements. A group of principals, $N \equiv \{1, \dots, n\}$, $n \geq 2$, interact with an agent, indexed 0, over an infinite number of discrete periods. In each period $t = 1, 2, \dots$, the principals can collectively make a demand to the agent, in exchange for a policy concession, x^t, chosen from a set $X \equiv [0, \hat{x}_0]$, where $\hat{x}_0 > 0$. The agent may concede to the demand, in which case we write $a^t = 1$, or not, in which case we write $a^t = 0$. If the principals choose not to make any demand to the agent (i.e., $x^t = \varnothing$), then status-quo policy 0 is implemented.

Principal i's stage payoff is $a^t\left[b_i^t - u(x^t)\right]$, where u is a convex, strictly increasing, continuously differentiable (dis)utility function on X, satisfying $u(0) = 0$; and b_i^t is a stochastic benefit chosen by Nature. We assume that each principal i's benefit from agreement is drawn at the start of every period from a c.d.f. F_i that is continuous and has full support on some interval $B \equiv [\underline{b}, \bar{b}]$, with $\underline{b} < \bar{b}$. The benefit profile's realization $b^t = (b_1^t, \dots, b_n^t)$ is publicly observed.

The agent's stage payoff is $a^t\left[u_0(x^t) - c^t\right]$, where where u_0 is a concave, strictly increasing, continuously differentiable utility function on X, satisfying $u_0(0) = 0$; and c^t is her privately observed cost from conceding to the principals' demand. This cost is initially drawn by Nature from a finite set $C \equiv \{c_1, \dots, c_K\}$, where $K \geq 2$ and $0 < c_1 < \dots < c_K < u_0(\hat{x}_0)$, according to some nondegenerate distribution $p^0 \in \Delta(C)$. We assume that p^0 satisfies a local monotone hazard rate property: for every $\underline{k} = 1, \dots, K-1$, the mapping $k \mapsto \sum_{\ell=\underline{k}}^{k} p^0(c_\ell)/p^0(c_{k+1})$ increases on $\{\underline{k}, \dots, K-1\}$.[5]

Like the principals' benefits, we allow the agent's type to change across periods. Given our focus on learning, however, we assume some persistence. For simplicity, the agent's type evolves according to a marked point process: at the end of every period, the agent's type is re-drawn from C according to p^0 (and the principals' common belief is correspondingly reset to p^0) with

[5] In fact, we only need this function not to decrease too fast. We could alternatively assume that $K = 2$ or that u is sufficiently convex, but we want to highlight that our results extend beyond the two-type case, and that they do *not* require the principals to be risk averse.

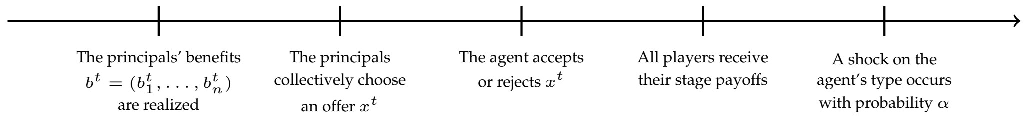

Figure 3 – Timing in each period $t = 1, 2, \ldots$.

probability $\alpha \in (0, 1)$. Otherwise, the agent's type remains unchanged.[6]

All players share a common discount factor $\delta \in (0, 1)$, and seek to maximize their average discounted payoffs.

Payoff Restrictions. First, we assume that $u_0^{-1}(c_K) < u^{-1}(\underline{b})$, so that agreement is socially efficient, regardless of the agent's type.[7] Second, players are sufficiently concerned for short-run outcomes, in the sense that $\delta < \overline{\delta}$ for some appropriately chosen $\overline{\delta} > 0$. Third, in order to guarantee some conflict of interest among the principals, we assume that $\underline{b}$ is not too large — otherwise the principals would always unanimously prefer to pool the agent's types— and that highest benefit $\overline{b}$ is not too close to $\underline{b}$. That is, we impose that $\underline{b} < \eta_1$ and $\overline{b} - \underline{b} > \eta_2$ for some appropriately chosen parameters η_1, $\eta_2 > 0$. The specific parameter thresholds $\overline{\delta}$, η_1, and η_2 are defined precisely in the appendix.

Timing. The timing is described in Figure 3.

Collective decision making. After the principals period-t benefits are realized, the principals collectively choose an offer x^t. The process of selecting an offer comprises two phases: an *organization* phase and a *negotiation phase*. Each phase is modeled as an amendment agenda game (Duggan, 2006, Austen-Smith and Banks, 2005). The agenda game is governed by a "procedure" that specifies the order in which the principals can include alternatives into the agenda, and the voting rule they use to select a winning alternative from the agenda.

Formally, let I be the set of finite sequences of proposers $\iota_1, \ldots, \iota_m$, $m \geq n$, that include all the principals (possibly with repetitions). A *procedure* consists of a probability distribution λ on I, and a collection $\mathcal{D} \subseteq 2^N \setminus \{\emptyset\}$ of decisive coalitions. We only restrict λ to belong to some (ex-

[6] We allow the agent's type to be re-drawn with positive probability at every period solely to ensure that the principals' learning process never stops.

[7] Alternatively, we could assume that $F_i\left[u(u_0^{-1}(c_K))\right]$ is sufficiently small for all i. We discuss this further in the paper's concluding remarks.

ogenously given) finite subset Λ of $\Delta(I)$; and $\mathcal{D}$ to be monotonic (e.g., $C \in \mathcal{D}$ and $C \subseteq C'$ imply $C' \in \mathcal{D}$) and proper ($C, C' \in \mathcal{D}$ implies $C \cap C' \neq \emptyset$) — e.g., Austen-Smith and Banks (1999). In what follows, we refer to any such a collection $\mathcal{D}$ as a *voting rule.* The family of procedures that satisfy these conditions is denoted by $\mathcal{P}$, with generic element $\wp = (\lambda, \mathcal{D})$.

Figure 4 illustrates the collective decision-making process. We describe each phase in detail.

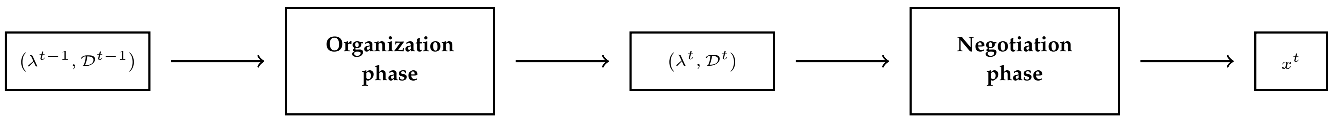

Figure 4 – The principals' collective decision-making process.

Organization Phase. In period t, the principals begin with a procedure $\wp^{t-1} = (\lambda^{t-1}, \mathcal{D}^{t-1})$ inherited from the previous period—the procedure $\wp^0$ that prevails at the start of the first period is exogenously given. A finite sequence of proposers $\iota_1, \ldots, \iota_m$, $m \geq n$, is first drawn from I using λ^{t-1}. The proposers can then suggest, in that order, amendments to $\wp^{t-1}$; let $\wp_j$ be the procedure suggested by the j^{th} proposer. The collective's final choice is determined by applying an amendment agenda to the resulting set of proposals: $\wp_m$ is pitted against $\wp_{m-1}$, the winner is then pitted against $\wp_{m-2}$, and so on, with the last remaining proposal $\wp_1$ pitted against the status quo, $\wp_0 = \wp^{t-1}$. In each round $j = 1, \ldots, m$ of the agenda, the principals vote sequentially (in an arbitrary order) either for $\wp_{m-j+1}$ or for $\wp_{m-j}$. The outcome of each pairwise vote is decided by the ongoing voting rule $\mathcal{D}^{t-1}$.

Following Duggan (2006), we assume that procedural ties—situations in which none of the proposals in a pairwise vote is supported by a decisive coalition—are resolved in favor of the proposal made earlier. As a consequence, $\wp_{m-j}$ beats $\wp_{m-j+1}$ in the jth round if and only if a blocking coalition of principals—i.e., a coalition S such that $N \setminus S \notin \mathcal{D}^{t-1}$—votes for $\wp_{m-j}$.

Let $\wp^t = (\lambda^t, \mathcal{D}^t)$ denote the outcome of the organization phase. The principals next move to the negotiation phase.

Negotiation Phase. A new sequence of proposers $\jmath_1, \ldots, \jmath_{m'}$, $m' \geq n$, is drawn from I using λ^t. Then, the same process as in the previous phase repeats, except that proposals are now policies in X, and pairwise votes in the amendment agenda are decided by the newly adopted voting rule $\mathcal{D}^t$. The winner of the agenda, denoted x^t, is the offer submitted by the principals to the agent.

Equilibrium. We study (pure-strategy) Markov perfect Bayesian equilibria of this game. Let Δ_{p^0} denote the set of probability distributions in $\Delta(C)$ that can be obtained from p^0 by Bayes updating, i.e.,

$$\Delta_{p^0} \equiv \left\{ p \in \Delta(C) \colon \exists C_0 \in 2^C \setminus \{\emptyset\} \text{ such that } p(c) = \frac{p^0(c)\mathbf{1}_{C_0}(c)}{\sum_{c' \in C_0} p^0(c')}, \forall c \in C \right\} ;$$

for every $p \in \Delta_{p^0}$, we define Δ_p in like manner. Equilibrium belief systems are required to satisfy the usual "no-signaling-what-you-don't-know condition," and to update any $p \in \Delta_{p^0}$ within Δ_p. Henceforth, we will refer to any Markov perfect Bayesian equilibrium that satisfies these restrictions more succinctly as an *equilibrium.*

4. Preliminary Results

Our goal is to extend the two results from our two-period example to the present framework. As a preliminary step, Lemma 1 establishes equilibrium existence, and Lemma 2 characterizes the outcome of any negotiation phase for a given period-t procedure. Lemma 3 then identifies the equilibrium offers generated by the principals' collective choice procedure. This allows us to (i) generalize the index $\beta(\mathcal{D})$ that captures a principal's incentive to screen the agent, and (ii) identify a "decisive" principal with whom a principal i might disagree.

Lemma 1. *An equilibrium exists.*

All equilibria of the negotiation phase have a simple structure.[8]

Lemma 2. *Let ϕ be any equilibrium. For any negotiation phase that begins with a procedure $\wp$ and a belief $p \in \Delta_{p^0}$, having support $\{c^1, \ldots, c^m\}$, $m \leq K$, there exist $\overline{x}^1 < \cdots < \overline{x}^m = u_0^{-1}(c^m)$ such that:*[9]

(i) regardless of the principals' benefits and the sequence of proposers, the principals' offer $x \in X$ must belong to $\{\overline{x}^1, \cdots, \overline{x}^m\}$; and

(ii) the type-c^ℓ agent accepts $\overline{x}^k$ if and only if $c^\ell \leq c^k$.

[8] Using a different refinement of PBE than Markov perfection, Acharya and Ortner (2017) obtain a similar equilibrium characterization for their single-principal framework. We stress that our proof, unlike theirs, relies on our restriction to small discount factors.

[9] To lighten notation, we omit the dependency of the $\overline{x}_k$'s on the equilibrium ϕ, procedure $\wp$, and belief p.

The principals select from a finite set of strictly increasing offers—one for each agent-type in their common belief's support. The largest offer $\overline{x}^m$ fully extracts surplus from the agent with the greatest possible cost; because the offer is accepted by all agent types, we call this the *pooling* offer. For each remaining $k = 1, ..., m-1$, offer $\overline{x}^k$ separates agent-types $\{c_1, ..., c_k\}$ from types $\{c_{k+1}, ..., c_{m-1}\}$. The agent's dynamic incentive constraints reflect that the principals' beliefs determine their future preferred offers, as well as the procedures the principals use to select from amongst those offers.

Which of the offers identified in Lemma 2 is chosen? Fix an equilibrium ϕ, and let $V_i^\phi(p; \lambda, \mathcal{D})$ denote principal i's continuation payoff at the start of every period that begins with belief p, and procedure (λ, $\mathcal{D}$). Lemma 2 yields that for any realization of the principals' benefit from an agreement $b = (b_1, \ldots, b_n)$, the negotiation phase induces a collective choice problem amongst the principals from the finite set of feasible alternatives $\{\overline{x}^1, \cdots, \overline{x}^m\}$. Principal i's preferences over this set are given by the utility function

$$U_i^\phi(\overline{x}^k \mid p, b_i, \lambda, \mathcal{D}) \equiv (1-\delta)\big[b_i - u(\overline{x}^k)\big] \sum_{\ell=1}^{k} p(c^\ell) + \delta \mathbb{E}[V_i^\phi(\tilde{p}; \lambda, \mathcal{D})]\,, \tag{4}$$

for each $k = 1, \ldots, m$, where $\tilde{p}$ is a random variable corresponding to the principals' belief at the start of the next period. The core $\mathcal{K}^\phi(p, b, \lambda, \mathcal{D})$ of this collective-choice problem can then be defined in the usual way: it is the subset of alternatives in $\{\overline{x}^1, \cdots, \overline{x}^m\}$ that cannot be defeated in a pairwise vote under the voting rule $\mathcal{D}$ (e.g., Austen-Smith and Banks 2005). In the Appendix, we verify that the principals' induced preferences defined in (4) are single-peaked for almost all $b_i \in B$, yielding that the core is non-empty.

Building on this observation, our next lemma has two parts. First, it identifies the outcome of the negotiation phase, i.e., it identifies which offer the principals actually make. Second—for future reference—it identifies a necessary and sufficient condition for principal i to prefer the pooling offer.

Lemma 3. *Let ϕ be any equilibrium, let $p \in \Delta_{p^0}$ and $(\lambda, \mathcal{D}) \in \mathcal{P}$, and let $\overline{x}^1, \ldots, \overline{x}^m$ be defined as in Lemma 2. Then, in any negotiation phase that begins with belief p and procedure $(\lambda, \mathcal{D})$:*

(i) for almost all $b \in B^n$ and all $\iota \in I$, the principals' offer when their realized benefits are b and the proposal sequence is ι solves

$$\max_x U^{\phi}_{\iota_1}(x \mid p, b, \lambda, \mathcal{D}), \text{ subject to } x \in \mathcal{K}^{\phi}(p, b, \lambda, \mathcal{D})\ ; \tag{5}$$

(ii) for every $i \in N$, there exists threshold $\beta^{\phi}_i(p; \lambda, \mathcal{D}) \in (\underline{b}, \bar{b})$ such that

$$\overline{x}^m = \arg\max_{x \in X} U^{\phi}_i(x \mid p, b, \lambda, \mathcal{D}) \tag{6}$$

if and only if $b_i > \beta^{\phi}_i(p; \lambda, \mathcal{D})$.

Recalling that ι_1 identifies the first proposer in the negotiation phase, Lemma 3 states that the principals select the first proposer's preferred offer from amongst the core alternatives of the collective choice problem.

The lemma also establishes an interior threshold on each principal i's benefit such that her ideal offer—regardless of whether it lies in the core—is the pooling offer if and only if her benefit realization exceeds that threshold. In our earlier example with two types of agent, the principals choose whether to offer a separating contract, or a pooling contract. With $K \geq 3$ agent types, there are potentially many ways to partially separate the agent. The threshold $\beta^{\phi}_i(\cdot)$ can be interpreted as a heuristic that reflects a principal's incentives to pursue any learning about the agent's type, instead of pursuing agreement by making an offer that all types accept: if $b_i > \beta^{\phi}_i(\cdot)$, the principal prefers to make a pooling offer that every agent accepts; but single-peakedness implies that for any $b_i < \beta^{\phi}_i(\cdot)$, there exists $m \geq 1$ such that a principal strictly prefers a contract that separates types $\{1, \ldots, m\}$ from the highest type $K-m$ types in the support of the principals' beliefs.

We now define a dictatorship in our framework.

Definition 1.

(1) Procedure $(\lambda, \mathcal{D})$ is a *formal dictatorship* if the voting rule $\mathcal{D}$ is dictatorial, i.e., if there is some principal i such that $\mathcal{D} = \{S \subseteq N \colon S \ni i\} \equiv \mathcal{D}^i$.

(2) Procedure $(\lambda, \mathcal{D})$ is an *informal dictatorship* if there is some $i \in \bigcap \mathcal{D}$ who proposes first with probability one under λ.

A procedure is a *dictatorship* if either (1) or (2) holds; otherwise, it is a *non-dictatorship*.

The first definition is standard: it identifies a unique individual that belongs to every decisive coalition and it corresponds to our two-period model. Nonetheless, a complete description of a "procedure" in our non-cooperative amendment agenda formulation includes not only a voting rule, but also the order in which proposers are recognized. Correspondingly, Lemma 3 suggests another way that procedures can concentrate authority. The lemma states that the first principal recognized in the negotiation phase secures her preferred offer from amongst the alternatives in the core. Moreover, any veto player's preferred offer lies in the core. So, a procedure that gives a veto player first-proposer rights ensures her most-preferred offer, even if the voting rule does not explicitly make her a dictator.

While the specific definition of an informal dictatorship is closely tied to the details of our amendment agenda game, it more broadly captures real-world decision-making contexts in which veto power is jointly vested with agenda-setting power, or where formal rules grant outsized privileges to some individuals. For example, Ali, Bernheim and Fan (2019) show that predictability about the order of future proposers in the Baron-Ferejohn legislative bargaining framework ensures that the first proposer is tantamount to a dictator, while Bernheim, Rangel and Rayo (2006) obtain that the last proposer has pre-eminent decision-making power in the context of an evolving default option.

5. Collective versus Individual Screening Incentives

Lemma 3 identifies a cut-off benefit $\beta_i^{\phi}(p; \wp)$ such that principal i prefers the pooling offer if and only if her realized benefit b_i exceeds $\beta_i^{\phi}(p; \wp)$. This cutoff can be loosely interpreted as reflecting a principal i's incentive to learn the agent's type. Our earlier Result 1 from our two-period example highlighted that a principal's benefit from learning the agent's type was higher under any rule that did not make her a dictator, relative to a rule that made her a dictator. We show that this result extends.

Proposition 1. *Let $\wp$ be any procedure in which principal i is not a dictator, and let $\wp^i$ be any dictatorship in which i is a dictator. Then for any equilibria ϕ and φ, we have*

$$\beta_i^{\phi}(p, \wp^i) < \beta_i^{\varphi}(p, \wp) \; ,$$

for all non-degenerate $p \in \Delta_{p^0}$.

Note that the comparison is strong, in the sense that it holds across *any* equilibria under either protocol. The intuition for i's benefit of learning about the agent under a non-dictatorship $\varphi \neq \varphi^i$ is the same as in our two-period example: screening the agent reduces conflict between the principals, and therefore insures i against the risks from not being decisive over future offers.

In our infinite horizon setting, however, there may also be *costs* of screening. To see why, suppose that the principals' beliefs place positive probability on K types of agent. For any period-t procedure, suppose the principals' period-t offer separates types $\{c_1, \ldots, c_{K-1}\}$ from $\{c_K\}$. Routine arguments establish that this offer, $\overline{x}^{K-1}$, is determined by type c_{K-1}'s binding incentive constraint:

$$(1-\delta)\left[u_0(\overline{x}^{K-1}) - c_{K-1}\right] + \delta \begin{bmatrix} \text{type } c_{K-1}\text{'s expected} \\ \text{continuation payoff} \\ \text{from accepting } \overline{x}^2 \end{bmatrix} = (1-\delta)\times 0 + \delta \begin{bmatrix} \text{type } c_{K-1}\text{'s expected} \\ \text{continuation payoff} \\ \text{from rejecting } \overline{x}^{K-1} \end{bmatrix},$$

so that her period-t rent is

$$(1-\delta)\left[u_0(\overline{x}^{K-1}) - c_{K-1}\right] = \delta \begin{bmatrix} \text{Expected difference in type } c_{K-1}\text{'s} \\ \text{continuation payoffs from} \\ \text{rejecting and accepting } \overline{x}^{K-1} \end{bmatrix}. \qquad (7)$$

Recognize that any shock to the agent's type between periods t and $t+1$ has no bearing on the type c_2 agent's period-t incentive constraint. The reason is that the shock resets the principals' common period-$t+1$ belief to p^0, and the period-t procedure persists at period $t+1$. Thus, the agent's period-$t+1$ continuation value after a shock at the end of the previous period is independent of her acceptance decision. It follows that the incentive constraint is:

$$(1-\delta)\left[u_0(\overline{x}^{K-1}) - c_{K-1}\right] = \delta(1-\alpha)\begin{bmatrix} \text{Expected difference in type } c_{K-1}\text{'s continuation} \\ \text{payoffs from rejecting and accepting } \overline{x}^{K-1} \\ \text{conditional on no shock between } t \text{ and } t+1 \end{bmatrix}.$$

In fact, the variation in the bracketed expression on the RHS across procedures is $O(\delta)$. To see why, observe that

(1) if type c_{K-1} *accepts* $\overline{x}^{K-1}$, then conditional on no shock between t and $t+1$, hers is the highest possible type in the support of the principals' beliefs in $t+1$. Standard arguments yield that she obtains zero rent. This observation is invariant across procedures.

(2) If type c_{K-1} *rejects* $\overline{x}^{K-1}$, then conditional on no shock between t and $t+1$, the principals assign probability one to c_K, and offer $u_0^{-1}(c_K)$ in $t+1$. This observation, again, is invariant across procedures, since the principals unanimously prefer this offer.

Hence, any wedge in the type c_{K-1} agent's continuation value from accepting versus rejecting a period-t partially separating offer under different procedures happens *no sooner* than period $t+2$. Any such wedge—and therefore any incremental cost to the principals across procedures—is scaled by δ^2 in the agent's period-t incentive constraints. The principals' learning benefit is instead scaled by δ, since it accrues immediately from period $t+1$. We conclude that so long as δ is not too large, the incremental costs of learning are second-order to the benefits of learning.

6. Evolution of Collective Choice Procedures

In our two-period setting, Result 2 unearthed a positive probability that the principals move from the status quo majority rule to a concentration of power either in the hands of an oligarchy or a dictator. In that example, however, the principals only chose their procedure once—at the start of the first period. This raises an obvious question: what can be said about the long-run evolution of decision-making when the principals can amend the status quo procedure in every period, and how does the answer depend on the initial rule at the start of their interaction with the agent?

Proposition 2. *Every equilibrium sequence of procedures* $\{(\lambda^t, \mathcal{D}^t)\}$ *converges to a dictatorship almost surely.*

To illustrate the theorem, we can extend our earlier example with five principals, in which the ongoing procedure at the start of period t is simple majority rule, the period-t belief is p^0 (e.g., a shock to the agent's type resets beliefs).

Fix an equilibrium, and let E denote the event "*the sequence of procedures starting in period t does not converge to a dictatorship.*" Suppose, contrary to Theorem 2, that $\Pr(E) > 0$, where probabilities are calculated according to the equilibrium strategies, and the distributions of principal benefits and shocks on the agent's type. Let $\mathcal{P}_E$ denote the set of procedures that the principals use in event E, and $P(\lambda, \mathcal{D})$ denote a lower bound (to be determined) on the probability that the principals adopt a dictatorship conditional on the arrival of a shock to the agent's type, given inherited procedure $(\lambda, \mathcal{D})$. Finally, let $\underline{P} \equiv \min\{P(\lambda, \mathcal{D}) : (\lambda, \mathcal{D}) \in \mathcal{P}_E\}$.[10]

To verify that $\underline{P} > 0$, let $\overline{\beta}_i$ denote principal i's smallest possible pooling threshold at belief p_0 in the event E, i.e.,

$$\overline{\beta}_i \equiv \min\left\{\beta_i^{\phi}(p^0; \lambda, \mathcal{D})\colon (\lambda, \mathcal{D}) \text{ is a non-dictatorship}\right\}, \tag{8}$$

where $\beta_i^{\phi}(p^0; \lambda, \mathcal{D})$ is defined in Lemma 3, and is the analogue of $\overline{\beta}$ defined in (3) of our earlier example. Proposition 1 yields that $\overline{\beta}_i > \underline{\beta}_i$, where $\underline{\beta}_i$ is i's pooling threshold when she is a dictator. Let F_1 denote the event described in our earlier example and highlighted in earlier Figure 2, in which

(i) b_1^t and b_2^t lie in a neighborhood of $\underline{b}$,

(ii) b_4^t and b_5^t lie in a neighborhood of $\overline{b}$, and

[10] $\underline{P}$ is well-defined because $\mathcal{P}_E$ is finite.

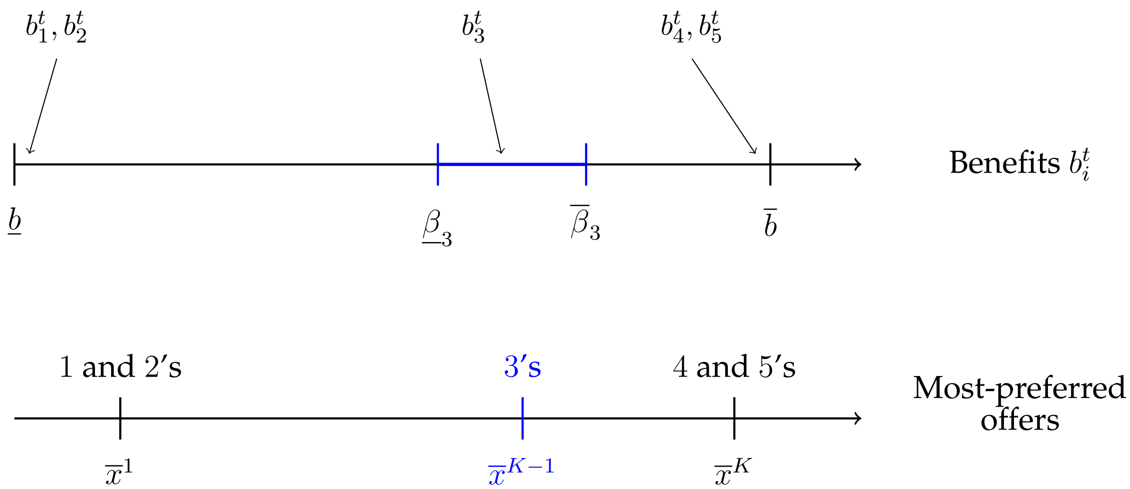

Figure 5 – The realization of principals' period-t benefits in event F_1.

(iii) b_3^t lies in $(\underline{\beta}_3, \overline{\beta}_3)$.

Figure 5 replicates Figure 2, but extends the principals' induced preferences to account for $K \geq 3$ agent-types. Nonetheless, all the intuition from that example extends: principals 4 and 5 prioritize short-term agreement and are therefore willing to make principal 3 a dictator—a procedure that 3 clearly welcomes. Nonetheless, we pointed out in our earlier example that 3's dictatorship is not the sole procedure that commits the principals to a period-t pooling offer: Lemma 3 yields that the pooling offer is assured if and only if it is the first proposer's preference from amongst alternatives in the core. In fact, there are three classes of procedures $\wp^t$ that satisfy this requirement:

Class A: either principal 3, 4 or 5 is a dictator, i.e., $\mathcal{D}^t = \mathcal{D}^i$ for $i \in \{3, 4, 5\}$,

Class B: principals 4 and 5 are oligarchs, i.e., $\mathcal{D}^t = \big\{S \subseteq N \colon S \supseteq \{4, 5\}\big\}$,

Class C: principals 4 and 5 are only blocking, i.e., $\{1, 2, 3\}$, $\{4, 5\} \notin \mathcal{D}^t$, and λ^t ensures that the first proposer is drawn from $\{4, 5\}$ with probability one, i.e., $\iota_1 \in \{4, 5\}$.

Note that Class C procedures were absent from our two-period example because we assumed that the principals could only choose the voting rule. If the principals adopt a procedure from class A, we set $P(\lambda^{t-1}, \mathcal{D}^{t-1}) = \Pr(F_1) > 0$.

Suppose, instead, the period-t organization phase yields a procedure from either classes B

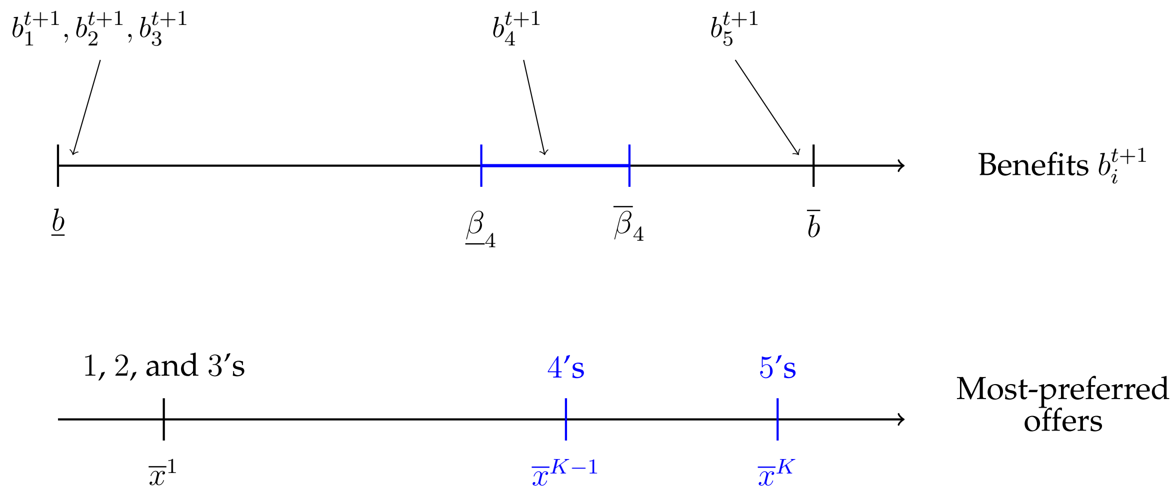

Figure 6 – The realization of principals' benefits in period $t+1$ in event F_2.

or C. Recognizing the inevitability of a period-t pooling offer, principals 1 and 2 might prefer to offer principals 4 or 5 a procedure that establishes this commitment without reverting immediately to a full-blown dictatorship. Suppose, for concreteness, that the principals adopt a class-B procedure in period-t's organization phase, and which therefore persists to period $t+1$. Since the period-t negotiation phase yields the pooling offer, the principals hold belief p^0 at period $t+1$ regardless of whether there is a shock to the agent's type.

Define the event F_2—illustrated in Figure 6—to be the conjunction of event F_1 in period t, followed by the following realization of benefits in period $t+1$:

(i) b_1^{t+1}, b_2^{t+1} and b_3^{t+1} lie in a neighborhood of $\underline{b}$,

(ii) b_5^{t+1} lies in a neighborhood of $\overline{b}$, and

(iii) b_4^{t+1} lies in $(\underline{\beta}_4, \overline{\beta}_4)$.

By a similar logic to the previous case, oligarch principals 4 and 5 are assured of a procedure that guarantees a period-$t+1$ pooling offer. Now, however, any such procedure *must* make either 4 or 5 a dictator. We can therefore set $P(\lambda^t, \mathcal{D}^t) = \Pr(F_2) > 0$. Notice that the final possible class C procedure the principals could adopt at period t follows a similar logic: while 4 and 5 are not oligarchs, whichever of these principals is recognized in the period-$t+1$ organization phase to propose first can propose her ideal rule and then vote for it. We can again set $P(\lambda^t, \mathcal{D}^t) = \Pr(F_2) > 0$.

Since there are infinitely many shocks to the agent's type in event E, and each shock is followed by the adoption of dictatorship with probability at least $\underline{P} = \min\{P(\lambda, \mathcal{D}) : (\lambda, \mathcal{D}) \in \mathcal{P}_E\} > 0$, we obtain a contradiction that $\Pr(E) = 0$, and thus obtain our result.

While our example supposed that the principals initially operate under a simple majority rule, Proposition 2 verifies that our argument also applies to any other inherited rule. To make this point concrete, notice that if the principals inherit a unanimous rule we can amend the event F_1 in Figure 2 to the positive probability event in which *all* the principals' benefits except for principal 3's are in a neighborhood of $\bar{b}$. By the same logic as our earlier analysis under majority rule, the high-benefit principals prioritize the agent's agreement. Since $b_3 \in (\underline{\beta}_3, \overline{\beta}_3)$, making principal 3 a dictator switches her induced preference for a partially separating offer to a pooling offer. Since principal 3 is strictly better off when made a dictator, and the remaining principals are strictly better off from the pooling offer than any other, the only outcome of the organization phase is *some* procedure that ensures the pooling offer at the negotiation phase. But since the organization phase operates under unanimity rule, the *only* shift in procedures that commits the principals to the pooling offer is 3's dictatorship. We therefore obtain the complete concentration of decision-authority in principal 3, which persists through all future periods. This example highlights that reverting to a dictatorship can be Pareto-improving for the principals.

7. Final Remarks

We introduce a framework to study collective screening by a group of principals, and study how the principals' static and dynamic internal conflicts shape their incentives to screen the agent. We further ask how the principals structure their internal bargaining processes in order to shape their external negotiations with the agent. We show how any non-dictatorial procedure encourages a principal to the agent more aggressively than under a dictatorial (i.e., single principal) benchmark. We also provide sufficient conditions such that, over time, short-run considerations lead to the inexorable concentration of power in the hands of a single principal.

We see a number of interesting questions for future research. The most immediate one concerns how information increases or instead decreases conflict between principals. In our framework, better information about the agent's cost reduces disagreement between the principals.

This is plausible in many negotiations contexts—for example, a union can calibrate its wage demands more effectively with better information about management's preferences. However, there may be other settings in which more information increases conflict between the principals.

To see how this phenomenon could arise in an extension of our model, return to our leading example but instead of presuming that $c_H < \underline{b}$, suppose instead that $c_L < \underline{b} < c_H < \bar{b}$. The assumption that $\underline{b} < c_H$ implies that for some realizations of the principals' benefits, an agreement with the high-cost agent is not efficient. Let $\tau^p(b, \mathcal{D})$ denote the probability that the period-2 offer is c_H when the principals' assign probability p that the agent's type is c_L. In the Appendix, we generalize expression (1) by showing that a principal i's net incentive to learn the agent's type under non-dictatorship is:

$$\beta(\mathcal{D}) - \beta(\mathcal{D}^i) = \delta(1-p) \int_b \left[\tau^p(b, \mathcal{D}^i) - \tau^p(b, \mathcal{D}) + \tau^0(b, \mathcal{D}) - \tau^0(b, \mathcal{D}^i)\right](b_i - b^*)\, dF(b). \quad (9)$$

Our benchmark with $c_H < \underline{b}$ corresponds to $\tau^0(b, \mathcal{D}) = \tau^0(b, \mathcal{D}^i) = 1$. So long as the probability of a benefits realization for which the decisive principal's benefit is below c_H isn't too large, (9) is strictly positive, yielding that i still has a greater incentive to screen the agent in a collective.

References

Acemoglu, Daron, Georgy Egorov and Konstantin Sonin. 2012. "Dynamics and stability of constitutions, coalitions, and clubs." *American Economic Review* 102(4):1446–76.

Acemoglu, Daron, Georgy Egorov and Konstantin Sonin. 2015. "Political economy in a changing world." *Journal of Political Economy* 123(5):1038–1086.

Acemoglu, Daron, Georgy Egorov and Konstantin Sonin. 2021. Institutional change and institutional persistence. In *The Handbook of Historical Economics*. Elsevier pp. 365–389.

Acharya, Avidit and Juan Ortner. 2017. "Progressive learning." *Econometrica* 85(6):1965–1990.

Ali, S Nageeb, B Douglas Bernheim and Xiaochen Fan. 2019. "Predictability and power in legislative bargaining." *The Review of Economic Studies* 86(2):500–525.

Anesi, Vincent and John Duggan. 2018. “Existence and indeterminacy of Markovian equilibria in dynamic bargaining games.” *Theoretical Economics* 13(2):505–525.

Anesi, Vincent and T Renee Bowen. 2021. “Policy experimentation in committees: A case against veto rights under redistributive constraints.” *American Economic Journal: Microeconomics* 13(3):124–62.

Austen-Smith, David and Jeffrey Banks. 1999. *Positive Political Theory I: Collective Preference*. University of Michigan Press.

Austen-Smith, David and Jeffrey Banks. 2005. *Positive Political Theory II: Strategy and Structure*. University of Michigan Press.

Bernheim, B Douglas, Antonio Rangel and Luis Rayo. 2006. “The power of the last word in legislative policy making.” *Econometrica* 74(5):1161–1190.

Bernheim, B Douglas and Michael D Whinston. 1986. “Menu auctions, resource allocation, and economic influence.” *The Quarterly Journal of Economics* 101(1):1–31.

Bowen, Renee, Ilwoo Hwang and Stefan Krasa. 2022. “Personal Power Dynamics in Bargaining.” *Journal of Economic Theory* p. 105530.

Bowen, T Renee, Ying Chen, Hülya Eraslan and Jan Zápal. 2017. “Efficiency of flexible budgetary institutions.” *Journal of Economic Theory* 167:148–176.

Buisseret, Peter and Dan Bernhardt. 2017. “Dynamics of policymaking: Stepping back to leap forward, stepping forward to keep back.” *American Journal of Political Science* 61(4):820–835.

Crawford, David L and Douglas H Blair. 2002. Labor union objectives and collective bargaining. In *The economics of labor unions/edited by Alison L. Booth*. Edward Elgar pp. 311–330.

Dessein, Wouter. 2002. “Authority and communication in organizations.” *The Review of Economic Studies* 69(4):811–838.

Diermeier, Daniel and Razvan Vlaicu. 2011. “Parties, coalitions, and the internal organization of legislatures.” *American Political Science Review* 105(2):359–380.

Duggan, John. 2006. "Endogenous voting agendas." *Social Choice and Welfare* 27(3):495–530.

Duggan, John and César Martinelli. 2020. "Electoral Accountability and Responsive Democracy." *The Economic Journal* 130(167):675–715.

Dziuda, Wioletta and Antoine Loeper. 2018. "Dynamic pivotal politics." *American Political Science Review* 112(3):580–601.

Eraslan, Hülya, Kirill S Evdokimov and Jan Zápal. 2022. "Dynamic legislative bargaining." *Bargaining* pp. 151–175.

Farber, Henry S. 1978. "Individual preferences and union wage determination: the case of the united mine workers." *Journal of Political Economy* 86(5):923–942.

Freer, Mikhail, César Martinelli and Siyu Wang. 2020. "Collective experimentation: A laboratory study." *Journal of Economic Behavior & Organization* 175:365–379.

Gieczewski, Germán and Svetlana Kosterina. 2020. "Endogenous Experimentation in Organizations." *Working Paper* .

Grossman, Gene M and Elhanan Helpman. 1994. "Protection for Sale." *The American Economic Review* pp. 833–850.

Grossman, Gene M and Elhanan Helpman. 2001. *Special Interest Politics*. MIT press.

Laffont, Jean-Jacques. 2000. *Incentives and Political Economy*. OUP Oxford.

Lagunoff, Roger. 2009. "Dynamic stability and reform of political institutions." *Games and Economic Behavior* 67(2):569–583.

Lefebvre, Perrin and David Martimort. 2020. ""When Olson Meets Dahl": From Inefficient Groups Formation to Inefficient Policy Making." *The Journal of Politics* 82(3):1026–1043.

Liu, Shuo and Dimitri Migrow. 2022. "When does centralization undermine adaptation?" *Journal of Economic Theory* 205:105533.

Michels, Robert. 1959. *Political parties: A sociological study of the oligarchical tendencies of modern democracy*. Dover.

Nunnari, Salvatore. 2021. "Dynamic legislative bargaining with veto power: Theory and experiments." *Games and Economic Behavior* 126:186–230.

Prat, Andrea and Aldo Rustichini. 2003. "Games played through agents." *Econometrica* 71(4):989–1026.

Rantakari, Heikki. 2008. "Governing adaptation." *The Review of Economic Studies* 75(4):1257–1285.

Riker, William H. 1980. "Implications from the Disequilibrium of Majority Rule for the Study of Institutions." *American Political Science Review* 74(2):432–446.

Roberts, Kevin. 2015. "Dynamic voting in clubs." *Research in Economics* 69(3):320–335.

Strulovici, Bruno. 2010. "Learning while voting: Determinants of collective experimentation." *Econometrica* 78(3):933–971.

Tommasi, Mariano and Federico Weinschelbaum. 2007. "Centralization vs. decentralization: A principal-agent analysis." *Journal of Public Economic Theory* 9(2):369–389.

Williamson, Oliver E. 1988. "The logic of economic organization." *Journal of Law, Economics, and Organization* 4:65.

Online Appendix

A. Proofs of Lemmas 1-3

We set $\overline{\delta} \equiv \min\{\overline{\delta}_0, \overline{\delta}_1, \overline{\delta}_2, \overline{\delta}_3, \overline{\delta}_4, \overline{\delta}_5, \overline{\delta}_6\}$, where the $\overline{\delta}_\ell$'s are upper bounds for the discount factor, defined below. We begin by establishing some notation and preliminary results. For each $k \in \{1, \dots, K\}$, let

$$y_k^-(\delta) \equiv u_0^{-1}\left(c^k - \frac{\delta(1-\alpha)}{1-\delta}u_0(\hat{x}_0)\right)$$

and

$$y_k^+(\delta) \equiv u_0^{-1}\left(c^k + \frac{\delta(1-\alpha)}{1-\delta}u_0(\hat{x}_0)\right).$$

Moreover, for every $p \in \Delta_{p^0}$, and each $c_k \in \operatorname{supp}(p)$, let $S_k^- \equiv \{c_1, \dots, c_k\} \cap \operatorname{supp}(p)$ and $S_k^+ \equiv \{c_{k+1}, \dots, c_K\} \cap \operatorname{supp}(p)$; let $p^{k-} \in \Delta_p$ be defined by

$$p^{k-}(c) \equiv \begin{cases} p(c)/p(S_k^-) & \text{if } c \in S_k^-, \\ 0 & \text{otherwise;} \end{cases}$$

let $p^{k+} \in \Delta_p$ be defined by

$$p^{k+}(c) \equiv \begin{cases} p(c)/p(S_k^+) & \text{if } c \in S_k^+, \\ 0 & \text{otherwise,} \end{cases}$$

where $p(S_k^-) \equiv \sum_{c \in S_k^-} p(c)$ and $p(S_k^+) \equiv \sum_{c \in S_k^+} p(c)$. For every nondegenerate $p \in \Delta_{p^0}$, whose support is denoted $\{c^1, \dots, c^m\}$, let $\beta_p \colon \{1, \dots, m-1\} \to \mathbb{R}$ be defined by

$$\beta_p(k) \equiv u(x^{k+1}) + \left[u(x^{k+1}) - u(x^k)\right] \frac{\sum_{\ell=1}^k p(c^\ell)}{p(c^{k+1})},$$

for all $k \in \{1, \dots, m-1\}$, where $x^\ell \equiv u_0^{-1}(c^\ell)$. This is the cutoff value of b_i that leaves each principal i indifferent between separating types All we need to ensure some conflict of interest among the principals (for low δ) is that $\underline{b} < \beta_p(k) < \overline{b}$, for some nondegenerate p and k. Without loss of generality, we will assume throughout that $\underline{b} < \min_p \beta_p(1) \equiv \underline{\beta}$ and $\overline{\beta} \equiv \max_p \beta_p(m-1) < \overline{b}$, where the minimum and the maximum are calculated over the nondegenerate type distributions in Δ_{p^0}. As β_p is strictly increasing function (see Lemma A1 below), this is achieved by setting $\eta_1 \equiv \underline{\beta}$ and $\eta_2 \equiv \overline{\beta} - \underline{\beta}$.

Finally, we say that a function $f\colon \{0, 1, \ldots, K\} \to \mathbb{R}$ is *quasi-single-peaked* if: (i) $\left|\arg\max_k f(k)\right| \leq 2$; (ii) if $k, \ell \in \arg\max_k f(k)$, then $\ell \in \{k-1, k, k+1\}$; and (iii) $\ell_1 < \ell_2 \leq \min\arg\max_k f(k)$ implies $f(\ell_1) < f(\ell_2)$, and $\max\arg\max_k f(k) \leq \ell_2 < \ell_1$ also implies $f(\ell_1) < f(\ell_2)$. In words, f is quasi-single-peaked if it has a single maximizer and is single-peaked; or if it has two maximizers, which must be adjacent, and it is increasing "below" the maximizers and decreasing "above" them.

Lemma A1. For every nondegenerate $p \in \Delta_{p^0}$, with support $\{c^1, \ldots, c^m\}$, the function β_p is strictly increasing on $\{1, \ldots, m-1\}$.

Proof. Take any nondegenerate $p \in \Delta_{p^0}$, and let $\underline{k} \equiv \min \operatorname{supp}(p)$. For each $k = 1, \ldots, m-2$, we have

$$\begin{aligned}\beta_p(k+1) - \beta_p(k) &= \left[u(x^{k+2}) - u(x^{k+1})\right]\left(1 + \frac{\sum_{\ell=1}^{k+1} p(c^\ell)}{p(c^{k+2})}\right) - \left[u(x^{k+1}) - u(x^k)\right]\frac{\sum_{\ell=1}^{k} p(c^\ell)}{p(c^{k+1})} \\ &= \left[u(x^{k+2}) - u(x^{k+1})\right]\left(1 + \frac{\sum_{\ell=\underline{k}}^{k-\underline{k}+2} p^0(c_\ell)}{p^0(c_{k-\underline{k}+3})}\right) - \left[u(x^{k+1}) - u(x^k)\right]\frac{\sum_{\ell=1}^{k-\underline{k}+1} p^0(c_\ell)}{p^0(c_{k-\underline{k}+2})},\end{aligned}$$

so that β_p is strictly increasing if

$$\frac{\sum_{\ell=1}^{k-\underline{k}+1} p^0(c_\ell)/p^0(c_{k-\underline{k}+2})}{1 + [\sum_{\ell=\underline{k}}^{k-\underline{k}+2} p^0(c_\ell)/p^0(c_{k-\underline{k}+3})]} < \frac{u(x_{k+2}) - u(x_{k+1})}{u(x_{k+1}) - u(x_k)}\ .$$

By convexity of u, the ratio on the right-hand side is greater than or equal to one; and by the local monotone hazard rate property, the ratio on the left-hand side is strictly less than one. $\square$

Lemma A2. There is $\overline{\delta}_0 > 0$ such that the following holds for all $\delta < \overline{\delta}_0$. Let $p \in \Delta_{p^0}$ be a belief whose support is denoted by $\{c^1, \ldots, c^m\}$, $1 \leq m \leq K$. Then, for each $i \in N$, every $b_i \in B$, every mapping $W_i\colon \Delta_p \to \left[\underline{b} - u(\hat{x}_0), \overline{b}\right]$ and $W_{i,0} \in \left[\underline{b} - u(\hat{x}_0), \overline{b}\right]$, and every $(\overline{x}_1, \ldots, \overline{x}_m) \in X^m$ such that

$\overline{x}_k \in \left[y_k^-(\delta), y_k^+(\delta)\right]$ for all $k = 1 \ldots, m$, the mapping $U_i(\cdot \mid b_i)\colon \{0, 1, \ldots, m\} \to \mathbb{R}$, defined by

$$\begin{aligned} U_i(0 \mid b_i) &\equiv \delta\big[(1-\alpha)W_i(p) + \alpha W_{i,0}\big] \ , \\ U_i(k \mid b_i) &\equiv (1-\delta)\big[b_i - u(\overline{x}_k)\big]p(S_k^-) + \delta\big[W_i(p^{k-})p(S_k^-) + W_i(p^{k+})p(S_k^+)\big] \\ &\qquad + \delta\alpha W_{i,0},\ k \neq 0, m \ , \\ U_i(m \mid b_i) &\equiv (1-\delta)\big[b_i - u(x_m)\big] + \delta\big[(1-\alpha)W_i(p) + \alpha W_{i,0}\big] \ , \end{aligned}$$

is quasi-single-peaked. Moreover, it is single-peaked for almost all $b_i \in B$.

Proof. Fix $p \in \Delta_{p^0}$. Consider first the mapping $U^p\colon \{0, 1, \ldots, m\} \times B \to \mathbb{R}$, defined by $U^p(0 \mid b) \equiv 0$, and $U^p(k \mid b) \equiv \big[b - u(x_k)\big]p(S_k^-)$, for all $k \in \{1, \ldots, m\}$ and $b \in B$. By definition, for any $k \in \{1, \ldots, m-1\}$, we have $U^p(k \mid b) \leq U^p(k+1 \mid b)$ if and only if $b \geq \beta(k)$ (and $U^p(k \mid b) > U^p(0 \mid b)$). As $\beta_p(k)$ is increasing in k (Lemma A1), the mapping $U^p(\cdot \mid b)$ is quasi-single-peaked, for all $b \in B$; and it is single-peaked for all $b \notin \big\{\beta_p(1), \ldots, \beta_p(m)\big\}$.

Now, let

$$\beta_k^-(\delta) \equiv p(c_{k+1})^{-1}\left[u\big(y_{k+1}^-(\delta)\big)p(S_{k+1}^-) - u\big(y_k^+(\delta)\big)p(S_k^-) - \frac{\delta(1-\alpha)}{1-\delta}u(\hat{x}_0)\right]$$

and

$$\beta_k^+(\delta) \equiv p(c_{k+1})^{-1}\left[u\big(y_{k+1}^+(\delta)\big)p(S_{k+1}^-) - u\big(y_k^-(\delta)\big)p(S_k^-) + \frac{\delta(1-\alpha)}{1-\delta}u(\hat{x}_0)\right] ;$$

and let $\overline{\beta}_k(\delta)$ be implicitly defined by $U_i\big(k \mid \overline{\beta}_k(\delta)\big) \equiv U_i\big(k+1 \mid \overline{\beta}_k(\delta)\big)$ for each $k \in \{1, \ldots, m-1\}$ — if $U(k \mid b_i) < U(k+1 \mid b_i)$ for all $b_i \in B$, then we set $\overline{\beta}_k(\delta) \equiv \underline{b}$; and if $U(k \mid b_i) > U(k+1 \mid b_i)$ for all $b_i \in B$, then $\overline{\beta}_k(\delta) \equiv \bar{b}$. By construction, for each k, $\overline{\beta}_k(\delta) \in \big[\beta_k^-(\delta), \beta_k^+(\delta)\big]$ and $\beta_k^-(\delta), \beta_k^+(\delta) \to \beta_p(k)$ as $\delta \to 0$. Hence, there exists $\overline{\delta}_p > 0$ such that $\overline{\beta}_k(\delta)$ is increasing in k and belongs to $(\underline{b}, \bar{b})$ whenever $\delta < \overline{\delta}_p$. This in turn implies that the mapping $U_i(\cdot \mid b_i)$ is quasi-single-peaked for all $b_i \in B$, whenever $\delta < \overline{\delta}_p$. Moreover, it is single-peaked for almost all $b_i \in B$, since indifference only occurs if b_i is equal to one of the $\overline{\beta}_k(\delta)$'s. As Δ_{p^0} is a finite set, we obtain the lemma by setting $\overline{\delta}_0 \equiv \min_{p \in \Delta_{p^0}} \overline{\delta}_p$. □

For any set of alternatives $\{0, 1, \ldots, m\}$, $1 \leq m \leq K$, and any profile of utility functions $f =$

$(f_1, \ldots, f_n)$ on $\{0, 1, \ldots, m\}$, we denote by $\text{Core}(m, f)$ the core of the corresponding collective-choice problem. Given a sequence of proposers ι, let $\mathcal{A}(m, f, \iota)$ denote the (one-shot) amendment-agenda game in which the set of alternatives is $\{0, 1, \ldots, m\}$, alternative 0 is the status quo, and the principals' payoffs are given by f. The following lemma is a variant on Duggan's (2006) Theorem 6.

Lemma A3. Let $f = (f_1, \ldots, f_n)$ be a profile of single-peaked functions on $\{0, 1, \ldots, m\}$, $1 \leq m \leq K$. Then, any Markovian equilibrium outcome of the amendment-agenda game $\mathcal{A}(m, f, \iota)$ is a maximizer of f_{ι_1} on $\text{Core}(m, f)$, for every realization of ι_1.

Proof. Consider any amendment-agenda game $\mathcal{A}(m, f, \iota)$. From the singlepeakedness of the f_i's, $\text{Core}(m, f)$ is nonempty, and all the alternatives in $\text{Core}(m, f)$ must be adjacent. It follows that each principal i has a unique ideal alternative in $\text{Core}(m, f)$, denoted $\hat{k}_i$. Suppose towards a contradiction that there is an equilibrium in which the chosen alternative, say k^*, is not $\hat{k}_{\iota_1}$. Then, the first proposer prefers k^* to $\hat{k}_{\iota_1}$; otherwise, she could profitably deviate from her equilibrium strategy by proposing $\hat{k}_{\iota_1}$, which would then be implemented — recall that procedural ties are resolved in favor of the alternatives proposed earlier. This in turn implies that k^* lies outside $\text{Core}(m, f)$. There must therefore exist an alternative $k \in \{0, 1, \ldots, m\}$ and a decisive coalition S such that all members of S prefer k to k^*. Recall that all principals have an opportunity to propose. None of the members of S can propose before k^* is included in the agenda (on the equilibrium path); otherwise she could profitably deviate from the equilibrium by proposing k as soon as it is her turn to propose. Now consider the proposal by a member of S, say j, when k^* is the provisionally selected alternative. As the equilibrium is Markovian, she and all the other members of S know that k^* will be implemented if k^* remains the provisionally selected alternative after this round — at the start of any new round, the number of remaining rounds and the provisionally selected alternative are the only payoff-relevant variables. All the members of S would therefore be strictly better off accepting proposal k, and therefore, proposing k is a profitable deviation for proposer j; a contradiction. $\square$

A.1. Proof of Lemma 1

Let $\overline{\delta}_0$ be defined as in Lemma A2. Observe that there exists $\overline{\delta}_1 > 0$ such that

$$\frac{2\delta(1-\alpha)}{1-\delta}u_0(\hat{x}_0) \leq \min_{k\in\{1,\ldots,K-1\}}(c_{k+1}-c_k)\,,$$

for all $\delta < \overline{\delta}_1$. The upper bound $\overline{\delta}$ is chosen to be smaller than or equal to $\min\{\overline{\delta}_0, \overline{\delta}_1\}$, so that $\delta < \min\{\overline{\delta}_0, \overline{\delta}_1\}$.

Let $\mathfrak{D}$ be the set of monotonic, proper voting rules $\mathcal{D}$, and let $L \equiv |\Lambda \times \mathfrak{D}| < \infty$. We can thus label the set of feasible procedures $\big\{(\lambda_1, \mathcal{D}_1), \ldots, (\lambda_L, \mathcal{D}_L)\big\}$. Let $\mathcal{V} \equiv \big[0, u_0(\hat{x}_0) - c_1\big]^L \times \big[0, u_0(\hat{x}_0) - c_K\big]^L \times \big[\underline{b} - u(\hat{x}_0), \overline{b}\big]^{nL}$. In what follows, a typical element of $\mathcal{V}$ will be denoted $(\nu_0, \nu_1, \ldots, \nu_n)$, where $\nu_0 = (\nu_{0,1}, \ldots, \nu_{0,K})$ with $\nu_{0,k} \in \big[0, u_0(\hat{x}_0) - c_k\big]^L$, for each $k = 1, \ldots, K$; and $\nu_i \in \big[\underline{b} - u(\hat{x}_0), \overline{b}\big]^L$, for each $i \in N$. We will think of $\nu_{0,k}$ as the L-dimensional vector whose ℓth component, $\nu_{0,k,\ell}$, describes the continuation payoff of the type-c_k agent at the start of period that begins with procedure $(\lambda_\ell, \mathcal{D}_\ell)$ and belief p^0. The vector ν_i and its components, the $\nu_{i,\ell}$'s, will be interpreted in like manner.

Fix a degenerate belief p that assigns probability one to some type c_k, $k = 1, \ldots, K$. For each procedure $(\lambda_\ell, \mathcal{D}_\ell)$, we define the game $\mathcal{G}^p(\lambda_\ell, \mathcal{D}_\ell \mid \nu_0, \nu_1, \ldots, \nu_n)$ among the principals as follows. Each period $t = 1, 2, \ldots$ begins with an ongoing procedure, say $(\lambda_l, \mathcal{D}_l)$. Then, events unfold as follows (if the game has not ended yet):

(1) The principals' benefit profile b^t is drawn according to the $F_i's$, and the sequence of proposers ι^t according to λ_k.

(2) The organizational phase takes place as in the main game. Let $(\lambda_{l'}, \mathcal{D}_{l'})$ denote the resulting procedure.

(3) A shock on the agent's type occurs with probability α.

(4) If a shock occurred in the previous stage, then the game ends, and each principal i receives a payoff of $(1-\delta)\big[b_i^t - u(x_k)\big] + \delta\nu_{i,l'}$; otherwise, she receives a stage-payoff of $(1-\delta)\big[b_i^t - u(x_k)\big]$, and the game transitions to period $t+1$, which begins with procedure $(\lambda_{l'}, \mathcal{D}_{l'})$.

The (exogenously given) initial procedure at the start of period 1 is $(\lambda_\ell, \mathcal{D}_\ell)$. All principals seek to maximize their average discounted payoffs. This is a noisy stochastic game, in which action sets are finite, the noise component of the state (i.e., the principals' benefits) is generated by the continuous distributions $F_1, \dots, F_n$ in every period, and the standard component (i.e., all the other payoff-relevant parameters) belongs to a finite set. It therefore admits a (possibly mixed) stationary Markov perfect equilibrium (Duggan, 2012). Let $V_i^p(\lambda_\ell, \mathcal{D}_\ell \mid \nu_0, \nu_1, \dots, \nu_n)$ denote principal i's equilibrium payoff. For future reference, we also define $V_{0,k}^p(\lambda_\ell, \mathcal{D}_\ell \mid \nu_0, \nu_1, \dots, \nu_n)$ as the corresponding expected payoff of the passive type-c_k agent.

Now fix $m = 2, \dots, K$. Suppose that for every $p' \in \Delta_{p^0}$ with $\left|\mathrm{supp}(p')\right| \leq m-1$, we have defined a game $\mathcal{G}^{p'}(\lambda_\ell, \mathcal{D}_\ell \mid \nu_0, \nu_1, \dots, \nu_n)$, $\ell = 1, \dots, L$, and corresponding continuation payoffs $V_i^{p'}(\lambda_\ell, \mathcal{D}_\ell \mid \nu_0, \nu_1, \dots, \nu_n)$ and $V_{0,k}^{p'}(\lambda_\ell, \mathcal{D}_\ell \mid \nu_0, \nu_1, \dots, \nu_n)$, as above. Consider a belief $p \in \Delta_{p^0}$ such that $\left|\mathrm{supp}(p)\right| = m$. For (and only for) expositional ease, suppose that $\mathrm{supp}(p) = \{c_1, \dots, c_m\}$. Observe that for every $k = 1, \dots, m-1$, $\left|\mathrm{supp}(p^{k-})\right| \leq m-1$ and $\left|\mathrm{supp}(p^{k+})\right| \leq m-1$ and therefore, $V_i^{p^{k-}}(\lambda_\ell, \mathcal{D}_\ell \mid \nu_0, \nu_1, \dots, \nu_n)$, $V_{0,k'}^{p^{k-}}(\lambda_\ell, \mathcal{D}_\ell \mid \nu_0, \nu_1, \dots, \nu_n)$, $V_i^{p^{k+}}(\lambda_\ell, \mathcal{D}_\ell \mid \nu_0, \nu_1, \dots, \nu_n)$, and $V_{0,k'}^{p^{k+}}(\lambda_\ell, \mathcal{D}_\ell \mid \nu_0, \nu_1, \dots, \nu_n)$, are well-defined for all i, k', and ℓ. This allows us to (implicitly) define the policy $\chi_k(\lambda_\ell, \mathcal{D}_\ell \mid \nu_0, \nu_1, \dots, \nu_n)$ as the unique solution x to

$$(1-\delta)\big[u_0(x) - c_k\big]p(S_k^-) + \delta(1-\alpha)V_{0,k}^{p^{k-}}(\lambda_\ell, \mathcal{D}_\ell \mid \nu_0, \nu_1, \dots, \nu_n)$$

$$= \delta(1-\alpha)V_{0,k}^{p^{k+}}(\lambda_\ell, \mathcal{D}_\ell \mid \nu_0, \nu_1, \dots, \nu_n)\,,$$

for each $k \leq m-1$, and $\chi_m(\lambda_\ell, \mathcal{D}_\ell \mid \nu_0, \nu_1, \dots, \nu_n) \equiv x_m$. Observe that $\overline{x}_k$, $k < m$, is defined in such a way that the type-c_k is indifferent between revealing that her type belongs to S_k^- and pretending that her type belongs to S_k^+, given the continuation values obtained for the "continuation games" above.

Next, for each procedure $(\lambda_\ell, \mathcal{D}_\ell)$, we define the game $\mathcal{G}^p(\lambda_\ell, \mathcal{D}_\ell \mid \nu_0, \nu_1, \dots, \nu_n)$ among the principals as follows. Each period $t = 1, 2, \dots$ begins with an ongoing procedure, say $(\lambda_l, \mathcal{D}_l)$. Then, events unfold as follows (if the game has not ended yet):

(1) The principals' benefit profile b^t is drawn according to the $F_i's$, and the sequence of proposers ι^t according to λ_k.

(2) The organizational phase takes place as in the main game. Let $(\lambda_{l'}, \mathcal{D}_{l'})$ denote the resulting procedure.

(3) The negotiation phase takes place as in the main game, but the principals are constrained to choose offers from the set $\{\chi_k(\lambda_{l'}, \mathcal{D}_{l'} \mid \nu_0, \nu_1, \ldots, \nu_n)\}_{k=1,\ldots,m}$. Let $\chi_{k'}(\lambda_{l'}, \mathcal{D}_{l'} \mid \nu_0, \nu_1, \ldots, \nu_n)$ denote the resulting offer to the agent.

(4) A shock on the agent's type occurs with probability α.

(5) If a shock occurred in the previous stage, then the game ends, and each principal i receives a payoff of $(1-\delta)\big[b_i^t - u\big(\chi_{k'}(\lambda_{l'}, \mathcal{D}_{l'} \mid \nu_0, \nu_1, \ldots, \nu_n)\big)\big] + \delta \nu_{i,l'}$; if a shock did not occur and $k' < m$, then the game ends, and she receives a payoff of $(1-\delta)\big[b_i^t - u\big(\chi_{k'}(\lambda_{l'}, \mathcal{D}_{l'} \mid \nu_0, \nu_1, \ldots, \nu_n)\big)\big] + \delta\big[p(S_{k'}^-) V_i^{p^{k'-}}(\lambda_{l'}, \mathcal{D}_{l'} \mid \nu_0, \nu_1, \ldots, \nu_n) + p(S_{k'}^+) V_i^{p^{k'+}}(\lambda_{l'}, \mathcal{D}_{l'} \mid \nu_0, \nu_1, \ldots, \nu_n)\big]$; otherwise, she receives a stage-payoff of $(1-\delta)\big[b_i^t - u\big(\chi_{k'}(\lambda_{l'}, \mathcal{D}_{l'} \mid \nu_0, \nu_1, \ldots, \nu_n)\big)\big]$, and the game transitions to period $t+1$, which begins with procedure $(\lambda_{l'}, \mathcal{D}_{l'})$.

The (exogenously given) initial procedure at the start of period 1 is $(\lambda_\ell, \mathcal{D}_\ell)$. All principals seek to maximize their average discounted payoffs. By the same logic as above, $\mathcal{G}^p(\lambda_\ell, \mathcal{D}_\ell \mid \nu_0, \nu_1, \ldots, \nu_n)$ admits a stationary Markov perfect equilibrium, and we can define $V_i^p(\lambda_\ell, \mathcal{D}_\ell \mid \nu_0, \nu_1, \ldots, \nu_n)$ as principal i's equilibrium payoff, and $V_{0,k}^p(\lambda_\ell, \mathcal{D}_\ell \mid \nu_0, \nu_1, \ldots, \nu_n)$ as the (passive) type-c_k agent's corresponding payoff. Proceeding recursively, we thus obtain the functions $V_i^p(\cdot \mid \cdot)$ and $V_{0,k}^p(\cdot \mid \cdot)$ for $p = p^0$.

Consider the continuous function that maps every $(\nu_0, \nu_1, \ldots, \nu_n) \in \mathcal{V}$ into $\Big(\big(V_{0,k}^{p^0}(\lambda_\ell, \mathcal{D}_\ell \mid \nu_0, \nu_1, \ldots, \nu_n)\big)_{\substack{\ell=1,\ldots,L\\k=1,\ldots,K}}, \big(V_i^{p^0}(\lambda_\ell, \mathcal{D}_\ell \mid \nu_0, \nu_1, \ldots, \nu_n)\big)_{\substack{i\in N\\k=1,\ldots,K}}\Big) \in \mathcal{V}$. Applying Brouwer's fixed point theorem, we obtain a fixed point $(\nu_0^*, \nu_1^*, \ldots, \nu_n^*)$ for this function. Now, define the game Γ as follows. Each period $t = 1, 2, \ldots$ begins with a belief $p \in \Delta_{p_0}$ and a procedure $(\lambda, \mathcal{D}) \in \Lambda \times \mathfrak{D}$, inherited from the previous period. (The initial belief and procedure at the start of period 1 are as in our main game.) Then, events unfold as follows:

(1) The principals' benefit profile b^t is drawn according to the $F_i's$, and the sequence of proposers ι^t according to λ_k.

(2) The organizational phase takes place as in the main game. Let $(\lambda_{l'}, \mathcal{D}_{l'})$ denote the resulting procedure.

(3) The negotiation phase takes place as in the main game, but the principals are constrained to choose offers from the set $\{\chi_k(\lambda_\ell, \mathcal{D}_\ell \mid \nu_0, \nu_1, \ldots, \nu_n)\}_{k=1,\ldots,m}$. Let $\chi_{k'}(\lambda_\ell, \mathcal{D}_\ell \mid \nu_0, \nu_1, \ldots, \nu_n)$ denote the resulting offer to the agent.

(4) A shock on the agent's type occurs with probability α.

(5) The game transitions to period $t+1$, which begins with ongoing procedure $(\lambda_{l'}, \mathcal{D}_{l'})$. If a shock occurred in the previous stage, then the belief at the start of $t+1$ is p^0; otherwise, it is p^{k-}.

It is easy to see that prescribing the principals to play as in the equilibrium of $\mathcal{G}^p(\lambda_\ell, \mathcal{D}_\ell \mid \nu_0^*, \nu_1^*, \ldots, \nu_n^*)$ in every period that begins with belief p and procedure $(\lambda_\ell, \mathcal{D}_\ell)$, we obtain a stationary Markov perfect equilibrium ς for Γ. We now modify ς to a pure-strategy profile $\hat{\varsigma}$ as follows. Observe that the outcome of every period is a policy $\chi_{k'}(\lambda_\ell, \mathcal{D}_\ell \mid \nu_0^*, \nu_1^*, \ldots, \nu_n^*)$ in $\{\chi_k(\lambda_\ell, \mathcal{D}_\ell \mid \nu_0^*, \nu_1^*, \ldots, \nu_n^*)\colon k = 1, \ldots, K \;\&\; \ell = 1, \ldots, L\}$ and a procedure $(\lambda_\ell, \mathcal{D}_\ell) \in \Lambda \times \mathfrak{D}$, yielding a payoff $(1-\delta)\big[b_i - u\big(\chi_{k'}(\lambda_\ell, \mathcal{D}_\ell \mid \nu_0^*, \nu_1^*, \ldots, \nu_n^*)\big)\big]p(S_{k'}^-) + \delta\big[\alpha V_i^{p^0}(\lambda_\ell, \mathcal{D}_\ell \mid \nu_0^*, \nu_1^*, \ldots, \nu_n^*) + (1-\alpha)V_i^{p^{k-}}(\lambda_\ell, \mathcal{D}_\ell \mid \nu_0^*, \nu_1^*, \ldots, \nu_n^*)\big]$ to the benefit-b_i principal i. Thus, for any pair of outcomes o and o', there is a unique cutoff value of b_i, say $\beta_i(o, o')$, for which principal i is indifferent between o and o'. Given that the sets of principals and outcomes are finite (and the F_i's are continuous), the set of benefit profiles $(b_1, \ldots, b_n) \in B^n$ such that $b_i = \beta_i(o, o')$ for some principal i and outcome pair (o, o'), denoted B_0, is of measure zero. In any period that begins with a benefit profile in B_0, we modify the actions prescribed by ς to those prescribed by some pure-strategy Markov-perfect equilibrium of the corresponding one-period game, where payoffs are defined using the continuation values induced by ς. (Existence of such an equilibrium follows directly from backward induction. Note that to maintain Markov perfection in the entire Γ, one must change ς in the same way in all periods that start with the same belief, procedure, and proposer sequence.) As B_0 is a measure-zero event, those changes to ς do not affect the continuation values at the start of each period, which we obtained above. Therefore, the strategy profile thus obtained is still a Markov perfect equilibrium of Γ.

Now take any period in which the realization of the benefit profile lies outside B_0, so that no principal can be indifferent between any two possible outcomes in this period. In the final (voting) stage, if the active principal randomizes, then it must be that her choice has no impact on the final outcome — otherwise, she would not be indifferent and, consequently, would not randomize. It follows that we can replace her randomized choice by a pure one without affecting the period's outcome and, therefore, the equilibrium conditions in the other stages of the game. We can then apply the same logic recursively to the previous stage in both the organizational and negotiation phases; and repeat the same process in any such period to obtain a new pure-strategy Markovian strategy profile, $\hat{\varsigma}$. By construction, the latter is a Markov perfect equilibrium of Γ.

We are now in a position to construct a (putative) equilibrium strategy profile for our main game. We begin with principals' strategies $(\phi_1, \ldots, \phi_n)$. Fix any belief $p \in \Delta_{p^0}$, with support $\{c^1, \ldots, c^m\}$, and any ongoing procedure $(\lambda, \mathcal{D}) \in \mathcal{P}$. Given p and $(\lambda, \mathcal{D})$, $(\phi_1, \ldots, \phi_n)$ prescribes the principals to play exactly as in $\hat{\varsigma}$ in the organizational phase, for all realizations of the benefit profile and the sequence of proposers. Given the belief p, the benefit profile b, and the protocol $(\lambda', \mathcal{D}')$ inherited from the organizational phase, consider the (one-shot) amendment agenda game, in which: the set of alternatives is X; the sequence of proposers is drawn according to λ'; the voting rule is $\mathcal{D}'$; and each principal i's payoff from choosing x is given by $(1-\delta)\big[b_i - u(x)\big]p(S_k^-) + \delta(1-\alpha)V_i^{p^{k-}}(\lambda', \mathcal{D}' \mid \nu_0^*, \nu_1^*, \ldots, \nu_n^*)$, where $k = 1, \ldots, m$ is the unique integer that satisfies $x \in \big[\chi_k(\lambda', \mathcal{D}' \mid \nu_0^*, \nu_1^*, \ldots, \nu_n^*), \chi_{k+1}(\lambda', \mathcal{D}' \mid \nu_0^*, \nu_1^*, \ldots, \nu_n^*)\big)$. (If $x \geq \chi_m(\lambda', \mathcal{D}' \mid \nu_0^*, \nu_1^*, \ldots, \nu_n^*)$, then $k = m$.) It follows from Zermelo's theorem that this game has pure-strategy subgame-perfect equilibria; it is readily checked that in one of them, the principals make the same offers as those prescribed by $\hat{\varsigma}$ in the negotiation phase. Strategies $(\phi_1, \ldots, \phi_n)$ prescribe the same behavior as that equilibrium in the corresponding negotiation phase.

We now turn to the agent's strategy, σ. Given any belief $p \in \Delta_{p^0}$, with support $\{c^1, \ldots, c^m\}$, and any ongoing procedure $(\lambda, \mathcal{D}) \in \mathcal{P}$, the type-$c_l$ accepts an offer $x \in [\overline{x}_k, \overline{x}_{k+1})$ if and only if

$$\delta(1-\alpha)V_{0,l}^{p^{k+}}(\lambda, \mathcal{D} \mid \nu_0^*, \nu_1^*, \ldots, \nu_n^*) \leq (1-\delta)\big[u_0(x) - c_l\big] + \delta(1-\alpha)V_{0,l}^{p^{k-}}(\lambda, \mathcal{D} \mid \nu_0^*, \nu_1^*, \ldots, \nu_n^*)\,;$$

she accepts any offer $x \geq x_m$, and rejects any offer $x \in \big[0, \chi_1(\lambda, \mathcal{D} \mid \nu_0^*, \nu_1^*, \ldots, \nu_n^*)\big)$. Finally, beliefs are updated as follows: if the principals make no offer, or if they make an offer $x \in \big[0, \chi_1(\lambda, \mathcal{D} \mid \nu_0^*, \nu_1^*, \ldots, \nu_n^*)\big)$, then the belief remains equal to p, irrespective of the agent's response; and for each $k = 1, \ldots, m-1$, if they make an offer $x \in \big[\chi_k(\lambda, \mathcal{D} \mid \nu_0^*, \nu_1^*, \ldots, \nu_n^*), \chi_{k+1}(\lambda, \mathcal{D} \mid \nu_0^*, \nu_1^*, \ldots, \nu_n^*)\big)$, then their belief becomes p^{k+} if the offer is accepted by the agent, and it becomes

p^{k-} if it is rejected.

To complete the proof of the lemma, it remains to verify that the strategy profile and belief system constructed in the previous paragraph is an equilibrium of our main game. By construction (and the induction hypothesis), we can focus on periods that begin with belief p. First, optimality of the principals' choices follows by construction — if a principal i had a profitable deviation from ϕ_i in this game, then she would also have a profitable deviation in one of the equilibria constructed for the other games above. Moreover, it follows from the definition of the strategy profile that the type-c_k agent's equilibrium value function at belief p and procedure $(\lambda_\ell, \mathcal{D}_\ell)$ is given by $V_{0,\ell}(\cdot \mid c_k) \equiv V^p_{0,k}(\lambda_\ell, \mathcal{D}_\ell \mid \nu^*_0, \nu^*_1, \ldots, \nu^*_n)$. Therefore, it follows immediately from the definition of her strategy and the principals' belief-updating rule that deviations are unprofitable.

Finally, we must verify that the principals' belief-updating rule is consistent with Bayes' rule (whenever possible). Take any belief $p \in \Delta_{p^0}$, with support $\{c^1, \ldots, c^m\}$, and any procedure $(\lambda_\ell, \mathcal{D}_\ell) \in \Lambda \times \mathfrak{D}$; and for notational ease, let $\overline{x}_k \equiv \chi_k(\lambda_\ell, \mathcal{D}_\ell \mid \nu^*_0, \nu^*_1, \ldots, \nu^*_n)$, for each $k = 1, \ldots, m$. Observe first that by definition of the $\overline{x}_k$'s, the type-c_k agent accepts the offer $\overline{x}_k$ from the principals in equilibrium. As her continuation values from accepting or rejecting any $x \in (\overline{x}_k, \overline{x}_{k+1})$ are equal to those from accepting or rejecting $\overline{x}_k$, and u_0 is an increasing function, she also accepts any $x \in (\overline{x}_k, \overline{x}_{k+1})$. This in turn implies that for all $c < c_k$, we have

$$\begin{aligned}(1-\delta)\big[u_0(x) - c\big] &+ \delta(1-\alpha)\big[V_{0,\ell}(p^{k-} \mid c) - V_{0,\ell}(p^{k+} \mid c)\big] \\ &\geq (1-\delta)\big[u_0(x) - c\big] + \delta(1-\alpha)\big[V_{0,\ell}(p^{k-} \mid c) - V_{0,\ell}(p^{k+} \mid c)\big] \\ &\quad - \Big[(1-\delta)\big[u_0(x) - c_k\big] + \delta(1-\alpha)\big[V_{0,\ell}(p^{k-} \mid c_k) - V_{0,\ell}(p^{k+} \mid c_k)\big]\Big] \\ &\geq (1-\delta)(c_k - c) - 2\delta(1-\alpha)u_0(\hat{x}_0) > 0\,,\end{aligned}$$

where the last inequality follows from $\delta < \overline{\delta} \leq \overline{\delta}_1$. Thus, all types $c \leq c_k$ accept any $x \in (\overline{x}_k, \overline{x}_{k+1})$. Moreover, for all $c > c_k$, the type-c agent's continuation value from accepting any $x \in (\overline{x}_k, \overline{x}_{k+1})$ is zero, conditional on no shock occurring on the path. As $(1-\delta)\big[u_0(x) - c\big] < 0 \leq \delta(1-\alpha)V_0(p^{k+} \mid c)$, her strategy then prescribes her to reject x. We conclude that the updating rule is consistent Bayes' rule following any offer $x \in (\overline{x}_k, \overline{x}_{k+1})$, $k = 1, \ldots, m-1$. By the same logic, it is also consistent Bayes' rule following offers in $[0, \overline{x}_1) \cap [x_m, \hat{x}_0]$. It is readily checked that principals' beliefs must belong to Δ_{p^0}, and that they satisfy the no-signaling-what-you-don't-know condition. This proves that the strategy profile and belief system constructed above constitute an equilibrium of the main game.

A.2. Proof of Lemma 2

Let $\overline{\delta}_1 > 0$ be defined as in the proof of Lemma 1. As $\delta \to 0$, $y_k^-(\delta), y_k^+(\delta) \to x_k \equiv u_0^{-1}(c_k)$. Therefore, there exists $\overline{\delta}_2 > 0$ such that $y_k^+(\delta) < y_{k+1}^-(\delta)$ for all $k = 1, \ldots, K-1$, whenever $\delta < \overline{\delta}_2$. For each $k \in \{1, \ldots, m-1\}$, let $\overline{\beta}_k(\delta)$ be defined as in the proof of Lemma A2. As we saw in that proof, $\overline{\beta}_k(\delta) \to \beta_p(k)$ as $\delta \to 0$. It follows that there exists a sufficiently small $\overline{\delta}_3 > 0$ such that $\overline{\beta}_{k+1}(\delta) - \overline{\beta}_k(\delta) \geq \big[\beta_p(k+1) - \beta_p(k)\big]/2$ for all $k \in \{1, \ldots, m-1\}$, whenever $\delta < \overline{\delta}_3$. We set $\overline{\delta} < \min\{\overline{\delta}_1, \overline{\delta}_2, \overline{\delta}_3\}$ and, henceforth, assume that $\delta < \overline{\delta}$.

Take any equilibrium, and let $p \in \Delta_{p^0}$. For notational ease, and without any loss of generality, assume that the support of p is $\{c_1, \ldots, c_m\}$, where $1 \leq m \leq K$. If the principals hold belief p and they make an offer that all agent types accept, then this offer must be x_m. To see this, observe first that as the type-c_{m-1} agent accepts any offer greater than or equal to $y_{m-1}^+(\delta) < y_m^-(\delta) < x_m$ (where the first inequality follows from $\delta < \overline{\delta} \leq \overline{\delta}_2$), she must accept any offer $x \geq x_m$. As we showed in the proof of Lemma 1, $\delta < \overline{\delta} \leq \overline{\delta}_1$ then implies that all types $c < c_{m-1}$ also accept any such offer. This in turn implies that type c_m must accept any offer $x > x_m$ in equilibrium: if she rejected x, thus revealing her type to the principals, then she would receive a payoff of zero until the arrival of the next shock, as the principals would trivially offer her x_m in every period. Accepting x (thus receiving a positive payoff) would be a profitable deviation. Now suppose that the principals make an offer $x > x_m$ that is accepted by all agent types in equilibrium. The proposer who successfully proposed x in that period could then profitably deviate by proposing some $x' \in (x_m, x)$ instead. That policy would still be accepted by all agent types; all the principals' stage-payoffs would be increased; and their continuation values would remain unchanged, as the belief would remain the same. This is a contradiction, showing that an equilibrium offer that is accepted by all agent types must be x_m. Note in passing that this also shows that the principals never make an offer above x_m in equilibrium and, consequently, that the payoff to the highest type in the support of p must be zero until the arrival of the next shock.

Let $\sigma(p,\lambda,\mathcal{D},x \mid c_k) \in \{0,1\}$ be the type-c_k agent's response to an offer $x \in X$ when the principals hold belief p and the ongoing procedure is $(\lambda,\mathcal{D})$. As $\delta < \overline{\delta} \le \overline{\delta}_2$, we have $y_\ell^+(\delta) < y_m^-(\delta)$, for all $\ell < m$. Hence, there exist offers that are accepted by all agent types but type c_m, i.e., the set $\{x \in X\colon \sigma(p,\lambda,\mathcal{D},x \mid c_{m-1}) = 1 - \sigma(p,\lambda,\mathcal{D},x \mid c_m) = 1\}$ is nonempty. Let $\overline{x}^{m-1}(p,\lambda,\mathcal{D}) \equiv \inf\{x \in X\colon \sigma(p,\lambda,\mathcal{D},x \mid c_{m-1}) = 1 - \sigma(p,\lambda,\mathcal{D},x \mid c_m) = 1\}$. Observe that $\overline{x}^{m-1}(p,\lambda,\mathcal{D})$ belongs to $\left[y_{m-1}^-(\delta), y_{m-1}^+(\delta)\right]$ and therefore, $\overline{x}^{m-1}(p,\lambda,\mathcal{D}) < \overline{x}^{m}(p,\lambda,\mathcal{D}) \equiv x_m$. By the same logic as in the previous paragraph, if the principals hold belief p and they make an offer that separates agent types in $\{c_1,\ldots,c_{m-1}\}$ from c_m, then this offer must be $\overline{x}_{m-1}(p,\lambda,\mathcal{D})$ — otherwise, it would have to be strictly higher than $\overline{x}_{m-1}(p,\lambda,\mathcal{D})$, and at least one principal could profitably deviate by inducing a slightly lower offer. Proceeding recursively, we define $\overline{x}_k(p,\lambda,\mathcal{D})$ for every $k = 1,\ldots,m-2$, in like manner.

To complete the proof of Lemma 2, it remains to establish that for each $k = 1,\ldots,m-1$, the principals separate agent types in $\{c_1,\ldots,c_k\}$ from those in $\{c_{k+1},\ldots,c_m\}$, and that they pool agent types (with a successful offer), with positive probability in equilibrium. As $\delta < \overline{\delta} \le \overline{\delta}_3$, the open intervals $\left(\overline{\beta}_{k-1}(\delta), \overline{\beta}_k(\delta)\right)$ (or $(\overline{\beta}_{m-1}(\delta), \overline{b})$) are nonempty. For realizations $(b_1,\ldots,b_n)$ of the principals' benefit profile such that $b_i \in \left(\overline{\beta}_{k-1}(\delta), \overline{\beta}_k(\delta)\right)$ (an event that arises with positive probability), the principals unanimously agree that separating $\{c_1,\ldots,c_k\}$ from $\{c_{k+1},\ldots,c_m\}$ is the best option, and must therefore do so in equilibrium by offering policy $\overline{x}_k(p,\lambda,\mathcal{D})$. Similarly, when all the principals' benefits belongs to $(\overline{\beta}_{m-1}(\delta), \overline{b})$, they all agree that pooling all the agent's types is the best option, so that the only possible outcome of the amendment-agenda game must be the offer x_m.

A.3. Proof of Lemma 3

The first part of the lemma is an immediate corollary of Lemmas 2, A2, and A3. The second part is directly obtained by defining $\beta_i^{\phi}(p,\lambda,\mathcal{D})$ as $\overline{\beta}_{m-1}(\delta)$ in the proof of Lemma A2 for the case where $W_i(p)$ is principal i's continuation value at belief p and ongoing procedure $(\lambda,\mathcal{D})$ under the equilibrium ϕ.

B. Proof of Proposition 1

For every equilibrium ϕ, let $V_i^\phi \colon \Delta_{p^0} \times \Lambda \times \mathfrak{D} \to \mathbb{R}$ be the value function of principal i induced by ϕ — i.e., for all $p \in \Delta_{p^0}$ and $(\lambda, \mathcal{D}) \in \Lambda \times \mathfrak{D}$, $V_i^\phi(p; \lambda, \mathcal{D})$ is i's expected continuation payoff at the start of any period that begins with belief p and procedure $(\lambda, \mathcal{D})$ (before the realization of the principals' benefit profile). Moreover, we denote by Γ the main game with endogenous procedures and for each $i \in N$, by Γ^i the benchmark game in which principal i is an (exogenously given) permanent dictator. For every equilibrium ϕ^i of the latter game, we denote by $W_i^{\phi^i}(p)$ dictator i's equilibrium continuation value at belief $p \in \Delta_{p^0}$. We begin by establishing a useful lemma.

Lemma B1. There exist $\kappa > 0$ and $\overline{\delta}_4 > 0$ such that the following holds for every $\delta < \overline{\delta}_4$, $i \in N$, and non-dictatorship $(\lambda, \mathcal{D})$. Let ϕ and ϕ^i be any equilibria of Γ and Γ^i, respectively; and let $p \in \Delta_{p^0}$ be a belief whose support is denoted by $\{c^1, \ldots, c^m\}$. Then,

$$W_i^{\phi^i}(p) - V_i^\phi(p; \lambda, \mathcal{D}) - p(S_{m-1}^-)\big[W_i^{\phi^i}(p_k^-) - V_i^\phi(p_{m-1}^-; \lambda, \mathcal{D})\big]$$

$$-p(S_{m-1}^+)\big[W_i^{\phi^i}(p_{m-1}^+) - V_i^\phi(p_{m-1}^+; \lambda, \mathcal{D})\big] > \kappa\,.$$

Proof. Take any principal $i \in N$, non-dictatorship $(\lambda, \mathcal{D})$, and nondegenerate belief $p \in \Delta_{p^0}$, whose support is denoted by $\{c^1, \ldots, c^m\}$. Consider a period of game Γ that begins with belief p and procedure $(\lambda, \mathcal{D})$; and suppose for the time being that $\delta = 0$. For every $b_j \in B$, the payoff to the benefit-b_j principal j from offering policy $x_k \equiv u_0^{-1}(c_k)$, $k = 1, \ldots, m$, to the agent is given by $U^p(k \mid b_j)$, as defined in the proof of Lemma A2. It follows that if the principals do not amend the ongoing procedure $(\lambda, \mathcal{D})$ in the organizational phase, the offer made to the agent will be the ideal of the first proposer ι_1 in the core induced by $(\lambda, \mathcal{D})$. Moreover, since the shortsighted principals' payoffs are independent of the ongoing procedure, it follows from the definition of the core that no procedure that would induce a different outcome may result from the organizational phase (in which $(\lambda, \mathcal{D})$ is the status quo).

For each $k = 1, \ldots, m$, let B_k be the set of realizations of the benefits and proposer sequences (at the start of the period) for which x_k is ι_1's ideal in the core, and let $\widehat{B}_k^i$ be those for which k is principal i's ideal in $\{1,\ldots,m\}$. We then have

$$\sum_{k=1}^{m}\Pr(\widehat{B}_k^i)\mathbb{E}\big[U^p(k\mid\tilde{b}_i)\mid\widehat{B}_k^i\big]=\sum_{k=1}^{m}\sum_{\ell=1}^{m}\Pr(\widehat{B}_k^i\cap B_\ell)\mathbb{E}\big[U^p(k\mid\tilde{b}_i)\mid\widehat{B}_k^i\cap B_\ell\big]\ ,$$

and

$$\sum_{\ell=1}^{m}\Pr(B_\ell)\mathbb{E}\big[U^p(\ell\mid\tilde{b}_i)\mid B_\ell\big]=\sum_{k=1}^{m}\sum_{\ell=1}^{m}\Pr(\widehat{B}_k^i\cap B_\ell)\mathbb{E}\big[U^p(\ell\mid\tilde{b}_i)\mid\widehat{B}_k^i\cap B_\ell\big]\ .$$

Let $\Delta_{k,\ell}\equiv\mathbb{E}\big[U^p(k\mid\tilde{b}_i)-U^p(\ell\mid\tilde{b}_i)\mid\widehat{B}_k^i\cap B_\ell\big]$. Since $\Delta_{k,\ell}=0$ whenever $k=\ell$, we have

$$\sum_{k=1}^{m}\Pr(\widehat{B}_k^i)\mathbb{E}\big[U^p(k\mid\tilde{b}_i)\mid\widehat{B}_k^i\big]-\sum_{\ell=1}^{m}\Pr(B_\ell)\mathbb{E}\big[U^p(\ell\mid\tilde{b}_i)\mid B_\ell\big]=\sum_{k=1}^{m}\sum_{\ell\neq k}\Pr(\widehat{B}_k^i\cap B_\ell)\Delta_{k,\ell}\ .$$

Note that since the decision-making procedure $(\lambda,\mathcal{D})$ is not a dictatorship (and the F_i's have full support), there exist different k and ℓ such that $\Pr(\widehat{B}_k^i\cap B_\ell)>0$.

Next, let $\Delta^-_{k,\ell}\equiv\mathbb{E}\big[U^{p^{(m-1)-}}(k'\mid\tilde{b}_i)-U^{p^{(m-1)-}}(\vec{\ell}\mid\tilde{b}_i)\mid\widehat{B}_k^i\cap B_\ell\big]$, where k' is a (random) maximizer of $U^{p^{(m-1)-}}(\cdot\mid\tilde{b}_i)$ — as above, we can ignore the measure-zero event in which i has two ideal alternatives — and $\vec{\ell}$ is the (random) alternative that satisfies $\phi(p^{(m-1)-},\tilde{b})=\overline{x}^{\vec{\ell}}$ (conditional on $\widehat{B}_k^i\cap B_\ell$). Observe that $U^{p^{(m-1)-}}(k,b)=U^p(k,b)/p(S^-_{m-1})$, for all $k\in\{1,\ldots,m-1\}$ and $b\in B$. Thus, if $k,\ell\geq m-1$, then $k'=\vec{\ell}=m-1$ and therefore, $\Delta^-_{k,\ell}=0$; if $k,\ell<m-1$, then $k'=k$ and $\vec{\ell}=\ell$, so that

$$\Delta^-_{k,\ell}=\mathbb{E}\big[U^p(k\mid\tilde{b}_i)-U^p(\ell\mid\tilde{b}_i)\mid\widehat{B}_k^i\cap B_\ell\big]p(S^-_{m-1})^{-1}\ ;$$

if $k<m-1\leq\ell$, then $k'=k$ and $\vec{\ell}=m-1$, so that

$$\Delta^-_{k,\ell}\equiv\mathbb{E}\big[U^p(k\mid\tilde{b}_i)-U^p(m-1\mid\tilde{b}_i)\mid\widehat{B}_k^i\cap B_\ell\big]p(S^-_{m-1})^{-1}\ ;$$

and, conversely, if $\ell<m-1\leq k$, then

$$\Delta^-_{k,\ell}\equiv\mathbb{E}\big[U^p(m-1\mid\tilde{b}_i)-U^p(\ell\mid\tilde{b}_i)\mid\widehat{B}_k^i\cap B_\ell\big]p(S^-_{m-1})^{-1}\ .$$

Hence, for all $k, \ell \in \{1, \ldots, m\}$ such that $k \neq \ell$, we have

$$\Delta_{k,\ell} - p(S^-_{m-1})\Delta^-_{k,\ell} = \begin{cases} \mathbb{E}\big[U^p(k \mid \tilde{b}_i) - U^p(\ell \mid \tilde{b}_i) \mid \widehat{B}^i_k \cap B_\ell\big] > 0 & \text{if } k, \ell \geq m-1\,, \\ \mathbb{E}\big[U^p(m-1 \mid \tilde{b}_i) - U^p(\ell \mid \tilde{b}_i) \mid \widehat{B}^i_k \cap B_\ell\big] > 0 & \text{if } k < m-1 \leq \ell\,, \\ \mathbb{E}\big[U^p(k \mid \tilde{b}_i) - U^p(m-1 \mid \tilde{b}_i) \mid \widehat{B}^i_k \cap B_\ell\big] > 0 & \text{if } \ell < m-1 \leq k\,, \\ 0 & \text{otherwise,} \end{cases}$$

where the inequalities follow from quasi-single-peakedness and the fact that by continuity of the F_i's, principal i can only be indifferent between two offers with probability zero. Hence, there is a sufficiently small $\kappa^i_p(\lambda, \mathcal{D}) > 0$ such that

$$\sum_{k=1}^{m} \sum_{\ell \neq k} \Pr(\widehat{B}^i_k \cap B_\ell)\big[\Delta_{k,\ell} - p(S^-_{m-1})\Delta^-_{k,\ell}\big] > \kappa^i_p(\lambda, \mathcal{D})\,.$$

Now let $\Delta^+_{p_0}$ be the subset of nondegenerate probability distributions in Δ_{p^0}; and let $\kappa \equiv \min\big\{\kappa^i_p(\lambda, \mathcal{D})\colon p \in \Delta^+_{p^0}, i \in N, (\lambda, \mathcal{D}) \in \mathcal{P}\big\} > 0$. As the principals' continuation payoffs are (uniformly) bounded over all possible outcomes, and $\beta^-_k(\delta), \beta^+_k(\delta) \to \beta_p(k)$ as $\delta \to 0$ (so that the probability measure of benefit profiles for which dynamic preferences differ from static ones converges to zero), there exists a sufficiently small $\overline{\delta}_p > 0$ such that whenever $\delta < \overline{\delta}_p$, $\big|W^{\phi^i}_i(p) - V^\phi_i(p; \lambda, \mathcal{D}) - \sum_{k=1}^m \sum_{\ell \neq k} \Pr(\widehat{B}^i_k \cap B_\ell)\Delta_{k,\ell}\big| < \kappa/2$ and $\big|W^{\phi^i}_i(p^{(m-1)-}) - V^\phi_i(p^{(m-1)-}; \lambda, \mathcal{D}) - \sum_{k=1}^m \sum_{\ell \neq k} \Pr(\widehat{B}^i_k \cap B_\ell)\Delta^-_{k,\ell}\big| < \kappa/2$, for any $i \in N$ and any equilibria ϕ and ϕ^i of Γ and Γ^i. Let $\overline{\delta}_4 \equiv \min\{\overline{\delta}_p\colon p \in \Delta^+_{p^0}\}$.

Trivially, $W^{\phi^i}_i(p^{(m-1)+}) - V^\phi_i(p^{(m-1)+}; \lambda, \mathcal{D}) = 0$ — all principals agree on the best offer to the agent when their common belief is degenereate. Therefore, for any equilibria ϕ and ϕ^i of Γ and Γ^i, we have

$$\begin{aligned} &W^{\phi^i}_i(p) - V^\phi_i(p; \lambda, \mathcal{D}) - p(S^-_{m-1})\big[W^{\phi^i}_i(p^-_k) - V^\phi_i(p^-_{m-1}; \lambda, \mathcal{D})\big] \\ &\qquad\qquad\qquad - p(S^+_{m-1})\big[W^{\phi^i}_i(p^+_{m-1}) - V^\phi_i(p^+_{m-1}; \lambda, \mathcal{D})\big] \\ &\geq \sum_{k=1}^{m} \sum_{\ell \neq k} \Pr(\widehat{B}^i_k \cap B_\ell)\big[\Delta_{k,\ell} - p(S^-_{m-1})\Delta^-_{k,\ell}\big] - \kappa > 0\,, \end{aligned}$$

as desired. □

We now return to the proof of the main proposition. Let $\delta < \overline{\delta} < \min\{\overline{\delta}_1, \overline{\delta}_2, \overline{\delta}_3, \overline{\delta}_4\}$. Take any principal $i \in N$, non-dictatorship $(\lambda, \mathcal{D})$, and nondegenerate belief $p \in \Delta_{p^0}$, whose support is denoted by $\{c^1, \ldots, c^m\}$. Let ϕ and ϕ^i be equilibria of Γ and Γ^i, respectively.

Consider first a negotiation phase of Γ, in which the principals hold belief p and use procedure $(\lambda, \mathcal{D})$. Given the equilibrium ϕ, any principal i prefers separating the agent types in $\{c^1, \ldots, c^{m-1}\}$ from type m to pooling all types in this period if and only if

$$(1-\delta)\big[b_i - u(\overline{x}^m)\big] + \delta(1-\alpha)V_i^{\phi}(p;\lambda,\mathcal{D}) \le (1-\delta)\big[b_i - u(\overline{x}^{m-1})\big]p(S_{m-1}^-) \\ + \delta(1-\alpha)\big[p(S_{m-1}^-)V_i^{\phi}(p_{m-1}^-;\lambda,\mathcal{D}) + p(S_{m-1}^+)V_i^{\phi}(p_{m-1}^+;\lambda,\mathcal{D})\big]\ ,$$

where $\overline{x}^{m-1}$ and $\overline{x}^m$ denote the equilibrium offers characterized in Lemma 2. In fact, by quasi-single-peakedness of continuation payoffs (Lemma A2), she prefers any separation of types to pooling all types if and only this inequality holds. It follows that

$$\beta_i^{\phi}(p;\lambda,\mathcal{D}) \equiv \big[(1-\delta)p(c^m)\big]^{-1}\Big[(1-\delta)\big[u(\overline{x}^m) - u(\overline{x}^{m-1})p(S_{m-1}^-)\big] \\ + \delta(1-\alpha)\big[p(S_{m-1}^-)V_i^{\phi}(p_{m-1}^-;\lambda,\mathcal{D}) + p(S_{m-1}^+)V_i^{\phi}(p_{m-1}^+;\lambda,\mathcal{D})\big] - V_i^{\phi}(p;\lambda,\mathcal{D})\Big]\ .$$

By the same logic, given the equilibrium ϕ^i of Γ^i, we can define $\hat{\beta}_i^{\phi^i}(p)$ as

$$\hat{\beta}_i^{\phi^i}(p) \equiv \big[(1-\delta)p(c^m)\big]^{-1}\Big[(1-\delta)\big[u(\overline{x}^m) - u(\hat{x}^{m-1})p(S_{m-1}^-)\big] \\ + \delta(1-\alpha)\big[p(S_{m-1}^-)W_i^{\phi^i}(p_{m-1}^-) + p(S_{m-1}^+)W_i^{\phi^i}(p_{m-1}^+)\big] - W_i^{\phi^i}(p)\Big]\ ,$$

where $\hat{x}^{m-1}$ is the policy offered by dictator i when she seeks to separate the agent types in $\{c^1, \ldots, c^{m-1}\}$ from type m in ϕ^i. It then follows from Lemma B1 (and the fact that $\overline{x}^m = \hat{x}^m = u_0^{-1}(c^m)$) that $\hat{\beta}_i^{\phi^i}(p) < \beta_i^{\phi}(p)$ if

$$(1-\delta)\big[u(\overline{x}^{m-1}) - u(\hat{x}^{m-1})\big] < \delta(1-\alpha)\kappa\ . \tag{B1}$$

Let $V_0^{\phi}(\cdot \mid c_{m-1})$ and $W_0^{\phi^i}(\cdot \mid c_{m-1})$ be the type-c_{m-1} agent's continuation values induced by ϕ and ϕ^i. Observe that $\overline{x}^{m-1}$ is the unique solution to

$$(1-\delta)\big[u_0(\overline{x}^{m-1}) - c_{m-1}\big] + \delta(1-\alpha)V_0^{\phi}(p^{(m-1)-} \mid c_{m-1}) = \delta(1-\alpha)V_0^{\phi}(p^{(m-1)+} \mid c_{m-1})$$

or, equivalently,

$$(1-\delta)\big[u_0(\overline{x}^{m-1}) - c_{m-1}\big] = \delta(1-\alpha)\big[V_0^{\phi}(p^{(m-1)+} \mid c_{m-1}) - V_0^{\phi}(p^{(m-1)-} \mid c_{m-1})\big] ,$$

where, for notational ease, we omit the dependency of $V_0^{\phi}(\cdot \mid c_{m-1})$ on $(\lambda, \mathcal{D})$. To see why this equation must hold in equilibrium, suppose towards a contradiction that the type-c_{m-1} agent is strictly better off accepting offer $\overline{x}_{m-1}$. By continuity of u_0, this implies that there exists a sufficiently small $\varepsilon > 0$ such that

$$(1-\delta)\big[u_0(\overline{x}^{m-1} - \varepsilon) - c_{m-1}\big] > \delta(1-\alpha)\big[V_0^{\phi}(p^{(m-1)+} \mid c_{m-1}) - V_0^{\phi}(p^{(m-1)-} \mid c_{m-1})\big] .$$

As $\overline{x}^{m-1} - \varepsilon < y_m^-(\delta)$, the type-$c_m$ agent would reject the offer $\overline{x}^{m-1} - \varepsilon$, so that the principals' updated beliefs would assign a probability of zero to types $c \geq c_m$ after observing a rejection of $\overline{x}^{m-1} - \varepsilon$. Hence, the type-$c_{m-1}$ agent would be strictly better off accepting $\overline{x}^{m-1} - \varepsilon$ than rejecting it, so that all the principals would be better off offering her $\overline{x}^{m-1} - \varepsilon$ rather than $\overline{x}^{m-1}$; a contradiction. By the same logic, $\hat{x}^{m-1}$ must satisfy

$$(1-\delta)\big[u_0(\hat{x}^{m-1}) - c_{m-1}\big] = \delta(1-\alpha)\big[W_0^{\phi^i}(p^{(m-1)+} \mid c_{m-1}) - W_0^{\phi^i}(p^{(m-1)-} \mid c_{m-1})\big] .$$

Let $v_0 \equiv u_0^{-1}$, $\overline{\Delta} \equiv V_0^{\phi}(p^{(m-1)+} \mid c_{m-1}) - V_0^{\phi}(p^{(m-1)-} \mid c_{m-1})$, and $\widehat{\Delta} \equiv W_0^{\phi^i}(p^{(m-1)+} \mid c_{m-1}) - W_0^{\phi^i}(p^{(m-1)-} \mid c_{m-1})$. Using the agent's incentive constraints above, we obtain:

$$\begin{aligned}
u(\overline{x}^{m-1}) - u(\hat{x}^{m-1}) &\leq u'(\overline{x}^{m-1})(\overline{x}^{m-1} - \hat{x}^{m-1}) \\
&= u'(\overline{x}^{m-1}) \left[v_0\left(c_{m-1} + \frac{\delta(1-\alpha)}{1-\delta}\overline{\Delta}\right) - v_0\left(c_{m-1} + \frac{\delta(1-\alpha)}{1-\delta}\widehat{\Delta}\right)\right] \\
&\leq u'(\overline{x}^{m-1}) v_0'\left(c_{m-1} + \frac{\delta(1-\alpha)}{1-\delta}\overline{\Delta}\right)\frac{\delta(1-\alpha)}{1-\delta}(\overline{\Delta} - \widehat{\Delta}) ,
\end{aligned}$$

where the inequalities follow from the convexity of u and v_0. Thus, if $\overline{\Delta} \leq \widehat{\Delta}$, condition B1 holds and we obtain the proposition.

Now suppose that $\overline{\Delta} > \widehat{\Delta}$, so that $(1-\delta)\big[u(\overline{x}^{m-1}) - u(\hat{x}^{m-1})\big] \leq \delta(1-\alpha)u'(\hat{x}_0)v_0'(\hat{x}_0)(\overline{\Delta} - \widehat{\Delta})$; and condition B1 holds whenever $u'(\hat{x}_0)v_0'(\hat{x}_0)(\overline{\Delta} - \widehat{\Delta}) < \kappa$. By definition, $p^{(m-1)+}$ is the degenerate probability distribution that assigns probability one to type c^m. When the principals hold such a belief, they unanimously agree that the best offer to the agent $\overline{x}^m = u_0^{-1}(c^m)$. It follows that starting from belief $p^{(m-1)+}$, this is the offer that must be made in the current period — and, as long as no shock occurs, in every future period — regardless of the procedures in place. As this is the best offer that the agent can receive in the continuation game, it is always optimal for her to accept it. It follows that $\big|V_0^{\phi}(p^{(m-1)+} \mid c_{m-1}) - W_0^{\phi^i}(p^{(m-1)+} \mid c_{m-1})\big| \leq \delta u_0(\hat{x}_0)$. Moreover, in any equilibrium (of either game), the offer made to the agent must be lower than or equal to $x_{m-1} \equiv u_0^{-1}(c_{m-1})$ (so that her stage-payoff is zero) when the principals hold belief $p^{(m-1)-}$. This implies that $\big|W_0^{\phi^i}(p^{(m-1)-} \mid c_{m-1}) - V_0^{\phi}(p^{(m-1)-} \mid c_{m-1})\big| \leq \delta u_0(\hat{x}_0)$. Therefore, $\overline{\Delta} - \widehat{\Delta} = V_0^{\phi}(p^{(m-1)+} \mid c_{m-1}) - W_0^{\phi^i}(p^{(m-1)+} \mid c_{m-1}) + W_0^{\phi^i}(p^{(m-1)-} \mid c_{m-1}) - V_0^{\phi}(p^{(m-1)-} \mid c_{m-1}) \leq 2\delta u_0(\hat{x}_0)$. We conclude that condition B1 holds whenever $\delta < \overline{\delta} \leq \overline{\delta}_5 \equiv \kappa/\big[2u_0(\hat{x}_0)u'(\hat{x}_0)v_0'(\hat{x}_0)\big]$.

Finally, observe that in any equilibrium φ of a continuation game of Γ that begins under some principal i's dictatorship, she remains a dictator in all future periods—possibly under different procedures. It follows that $V_i^{\varphi}(p;\wp^i) \equiv W_i^{\phi^i}(p)$, and therefore $\beta_i^{\varphi}(p;\wp^i) \equiv \hat{\beta}_i^{\phi^i}(p)$, for every procedure $\wp^i$ under which principal i is a dictator. This completes the proof of the proposition.

C. Proof of Proposition 2

We begin with some useful observations. First, for every $\lambda \in \Lambda$, let $q_i(\lambda)$ be the probability that principal $i \in N$ proposes first under λ; and let $q \equiv \min\{q_i(\lambda)\colon i \in N, \lambda \in \Lambda, q_i(\lambda) > 0\}$. Then, there exists a sufficiently small $\hat{\delta}_{6,1} > 0$ such that

$$(1-\delta)q\big[\overline{b} - u\big(y_{K-1}^{+}(\delta)\big)\big]p^0(S_{K-1}^{-}) + \delta\overline{b} < (1-\delta)q\big[\overline{b} - u(x_K)\big]\;,$$

for all $\delta < \hat{\delta}_{6,1}$. Given any equilibrium ϕ, let $\overline{x}^{K-1}$ be defined as in Lemma 2 for $p = p^0$; and observe that $\overline{x}^{K-1} \leq y_{K-1}^{+}(\delta)$ (otherwise, the type-c_{K-1} agent would have a profitable deviation when offered $\overline{x}^{K-1}$). It follows from the inequality above that in any period t, any principal whose period-t benefit is $\overline{b}$ strictly prefers pooling all agent types with certainty to separating those in $\{c_1, \ldots, c_{K-1}\}$ from c_K with a probability greater than or equal to q, regardless of what happens from period $t+1$ onward. By continuity, this also holds for all benefits $b \in (\overline{b} - \varepsilon_1, \overline{b}]$, for some small enough $\varepsilon_1 > 0$. Similarly, there exist sufficiently small $\hat{\delta}_{6,2}, \varepsilon_2 > 0$ such that whenever $\delta < \hat{\delta}_{6,2}$, any principal whose benefit belongs to $[\underline{b}, \underline{b} + \varepsilon_2)$ strictly prefers separating type c_1 from those in $\{c_2, \ldots, c_K\}$ to making the pooling offer, regardless of future play. Let $\varepsilon \equiv \min\{\varepsilon_1, \varepsilon_2\}$. Moreover, by the same logic as in the proof of Lemma A2, there exists a sufficiently small $\overline{\delta}_6 > 0$, lower than $\min\{\hat{\delta}_{6,1}, \hat{\delta}_{6,2}\}$, such that $\beta^{+}_{K-1}(\delta) < \overline{b} - \varepsilon$, for all $\delta < \overline{\delta}_6$. Henceforth, we assume that $\delta < \overline{\delta} \leq \overline{\delta}_6$.

Now suppose towards a contradiction that there is an equilibrium ϕ of the extended game in which the sequence of procedures adopted by the principals does not converge almost surely to a dictatorship. In any period t, if $(\lambda^t, \mathcal{D}^t)$ is a dictatorship, then either $(\lambda^{t+1}, \mathcal{D}^{t+1}) = (\lambda^t, \mathcal{D}^t)$, or $(\lambda^{t+1}, \mathcal{D}^{t+1})$ is another dictatorship with the same dictator as in t. Therefore, the set of stochastic sequences of principal benefits, shocks on the agent's types, and proposer sequences for which the principals never adopt a dictatorship in equilibrium constitutes an event that occurs with positive probability. We denote this event by E. Thus, by Proposition 1, at every history in the period-t negotiation phase that is consistent with E, if the belief p^{t-1} is nondegenerate, then the equilibrium pooling cutoff of each principal i, $\beta^{\phi}_i(p^{t-1}; \wp^t)$, must be be lower than her pooling cutoff when she is a dictator, which we denote by $\hat{\beta}_i(p^{t-1})$.

Let $\mathcal{P}_E$ be the set of procedures that may prevail on paths consistent with E. Our next step is to define for every $(\lambda, \mathcal{D}) \in \mathcal{P}_E$, a lower bound $P(\lambda, \mathcal{D})$ on the probability that the principals adopt a dictatorship as their decision-making procedure if a shock on the agent's type occurs while the ongoing procedure is $(\lambda, \mathcal{D})$. (Markov perfection ensures that this probability only depends on $(\lambda, \mathcal{D})$ and the principals' belief, which must be p^0 after a shock.) For each principal i, let $\overline{\beta}_i(p^0) \equiv \min\left\{\beta^{\phi}_i(p^0; \lambda, \mathcal{D}) \colon (\lambda, \mathcal{D}) \in \mathcal{P}_E\right\}$. Take any $(\lambda, \mathcal{D}) \in \mathcal{P}_E$; pick an arbitrary minimal decisive coalition S_1 in $\mathcal{D}_1$ and a principal i_1 in S_1 who may propose first with positive probability (if such a coalition does not exist, add a first proposer to some minimal decisive coalition); and let F_1 be the positive-probability event: "$b_{i_1} \in \left(\hat{\beta}_{i_1}(p^0), \overline{\beta}_{i_1}(p^0)\right)$, $b_j \in (\overline{b} - \varepsilon, \overline{b}]$ for all $j \in S_1 \setminus \{i_1\}$, and $b_j \in [\underline{b}, \underline{b}+\varepsilon)$ for all $j \in N \setminus S_1$." We claim that at any history (consistent with E) with ongoing procedure $(\lambda, \mathcal{D})$ that ends with a shock on the agent's type, followed by F_1, one of the following procedural changes must occur in equilibrium: either (i) some member of S_1 is made a (formal or informal) dictator; or (ii) some subcoalition of $S_1 \setminus \{i_1\}$ is made minimal decisive; or (iii) some subcoalition of $S_1 \setminus \{i_1\}$ is made blocking, but not decisive, and the first proposer belongs to that subcoalition with probability one. Moreover, the offer made to the agent must be x_K — so that the belief at the start of the next period must still be p^0. To see this, observe first that if i_1 is made a dictator, it will be optimal for her to pool all the agent types, since $b_{i_1} > \hat{\beta}_{i_1}(p^0)$. As $b_j \in (\overline{b}-\varepsilon, \overline{b}]$ for all the other members j of S_1 (and $\delta < \overline{\delta}_6$), this is also their ideal offer, regardless of the prevailing procedure. It follows that in the organizational phase, the only possible outcomes are procedures that induce the pooling offer as the outcome of the ensuing negotiation phase — otherwise, at least one member of the decisive coalition S_1 would have a profitable deviation during the former phase — since making i_1 a dictator guarantees that coalition's ideal outcome. Finally, observe that for offer x_K to be made with certainty in equilibrium of the negotiation phase, one of the following must be true: x_K is the only alternative in the core (leaving the first proposer no other option), i.e., either case (i) or case (ii) above hold; or case (iii) holds, so that x_K belongs to the core and the first proposer always selects it. If some member of S_1 becomes a dictator after F_1, then we set $P(\lambda, \mathcal{D}) \equiv \Pr(F_1) > 0$; otherwise, we denote by $(\lambda_2, \mathcal{D}_2)$ the new ongoing procedure, by S_2 the relevant subcoalition of $S_1 \setminus \{i_1\}$, and we proceed recursively as explained below.

Fix $k = 2, \ldots, |S_1|-1$. Suppose that we have defined F_ℓ for each $\ell = 1, \ldots, k-1$ (and therefore, S_ℓ for each $\ell = 1, \ldots, k$), but $P(\lambda, \mathcal{D})$ is not yet defined. Fixing $i_k \in S_k$ — when S_k is blocking but not decisive, i_k must be one of the members of S_k who may propose first — we then define the positive-probability event F_k as follows: "events $F_1, \ldots, F_{k-1}$ have successively occurred in the previous $k-1$ periods; $b_{i_k} \in \left(\hat{\beta}_{i_k}(p^0), \overline{\beta}_{i_k}(p^0)\right)$, $b_j \in (\overline{b}-\varepsilon, \overline{b}]$ for all $j \in S_k \setminus \{i_k\}$, and $b_j \in [\underline{b}, \underline{b}+\varepsilon)$ for all $j \in N \setminus S_k$." (Note that by construction, in cases where S_k is not decisive, the first proposer ι_1 must be i_k.) Repeating the same arguments as in the previous paragraph, we obtain that in equilibrium, one of the following procedural changes must occur after F_k: either (i) some member of S_k is made a dictator; or (ii) some subcoalition of $S_k \setminus \{i_k\}$ is made minimal decisive; or (iii) some subcoalition of $S_k \setminus \{i_k\}$ is made blocking, but not decisive, and the first proposer belongs to that subcoalition with probability one. (Note that even when coalition S_k is not decisive, its ideal outcome can still be guaranteed by making i_k a dictator. The coalition being decisive, the pooling offer must belong to the core and be selected by the first proposer, who must be one of its members by construction.) If some member of S_k becomes a dictator after F_k, then we set $P(\lambda, \mathcal{D}) \equiv \Pr(F_k) > 0$; otherwise, we denote by $(\lambda_{k+1}, \mathcal{D}_{k+1})$ the new ongoing procedure, by S_{k+1} the relevant subcoalition of $S_k \setminus \{i_k\}$, and repeat the same process.

Observe that this process must end with a dictatorship after at most $|S_1|$ iterations. We can then conclude that in event E, the probability that the principals adopt a dictatorship after a shock on the agent's type is bounded from below by $\min\left\{P(\lambda, \mathcal{D})\colon (\lambda, \mathcal{D}) \in \mathcal{P}_E\right\} > 0$. As an infinite number of such shocks must occur on any path, this in turn implies that $\Pr(E) = 0$, yielding the desired contradiction.

D. Derivation of Equation (9)

This section derives equation (9). Suppose $c_L < \underline{b} < c_H < \bar{b}$. For every benefits profile $b \in [\underline{b}, \bar{b}]^5$ and rule $\mathcal{D}$, let $\tau^{\text{sep}}(b, \mathcal{D})$ denote the equilibrium probability that the principals offer c_H in period 2, given that they learned that $c = c_H$ in period 1; and let $\tau^{\text{pool}}(b, \mathcal{D})$ denote the probability that they offer c_H in period 2, given that they pooled in period 1. Then, principal i's continuation value from a period-1 separating offer is now

$$W^{\text{sep}}(\mathcal{D}) \equiv p \int_b (b_i - c_L)\, dF(b) + (1-p) \int_b \tau^{\text{sep}}(b, \mathcal{D})(b_i - c_H)\, dF(b)\,,$$

and her continuation value from a period-1 pooling offer is

$$\begin{aligned} W^{\text{pool}}(\mathcal{D}) \equiv p \int_b \Big[b_i - \tau^{\text{pool}}(b, \mathcal{D})c_H - \big(1 - \tau^{\text{pool}}(b, \mathcal{D})\big)c_L\Big]\, dF(b) \\ + (1-p) \int_b \tau^{\text{pool}}(b, \mathcal{D})(b_i - c_H)\, dF(b)\,. \end{aligned}$$

Hence,

$$\Delta(\mathcal{D}) \equiv W^{\text{sep}}(\mathcal{D}) - W^{\text{pool}}(\mathcal{D})$$

$$= p \int_b \tau^{\text{pool}}(b, \mathcal{D})(c_H - c_L)\, dF(b) + +(1-p) \int_b \big[\tau^{\text{sep}}(b, \mathcal{D}) - \tau^{\text{pool}}(b, \mathcal{D})\big](b_i - c_H)\, dF(b) \; ;$$

and the difference in continuation values, $\Delta(\mathcal{D}) - \Delta(\mathcal{D}^i)$, is equal to

$$\begin{aligned}
&\int_b \big[\tau^{\text{pool}}(b, \mathcal{D}) - \tau^{\text{pool}}(b, \mathcal{D}^i)\big]\big[c_H - pc_L - (1-p)b_i\big]\, dF(b) \\
&\qquad\qquad\qquad + \int_b \big[\tau^{\text{sep}}(b, \mathcal{D}) - \tau^{\text{sep}}(b, \mathcal{D}^i)\big](b_i - c_H)\, dF(b) \\
&= (1-p) \int_b \big[\tau^{\text{pool}}(b, \mathcal{D}^i) - \tau^{\text{pool}}(b, \mathcal{D})\big](b_i - b^*)\, dF(b) \\
&\qquad\qquad\qquad + (1-p) \int_b \big[\tau^{\text{sep}}(b, \mathcal{D}) - \tau^{\text{sep}}(b, \mathcal{D}^i)\big](b_i - c_H)\, dF(b) \; ,
\end{aligned}$$

where the equality follows from the definition of b^*.

Next, we turn to the agent's incentive-compatibility constraints. If the principals choose to separate, then the low-type agent's (binding) constraint under rule $\mathcal{D}$ is

$$\overline{x}^L - c_L + \delta \times 0 = 0 + \delta(c_H - c_L) \int_b \tau^{\text{sep}}(b, \mathcal{D})\, dF(b) \; ,$$

which allows us to define the first-period offer

$$\overline{x}^L(\mathcal{D}) \equiv c_L + \delta(c_H - c_L) \int_b \tau^{\text{sep}}(b, \mathcal{D})\, dF(b) \; .$$

It follows that for any rule $\mathcal{D}$, the net value of separation to principal i is given by

$$\varphi(\mathcal{D}) \equiv p\big[b_i - \overline{x}^L(\mathcal{D})\big] - (b_i - c_H) + \delta\Delta(\mathcal{D}) \; .$$

This in turn implies that

$$\begin{aligned}
\varphi(\mathcal{D}) - \varphi(\mathcal{D}^i) &= \delta(1-p) \int_b \big[\tau^{\text{pool}}(b, \mathcal{D}^i) - \tau^{\text{pool}}(b, \mathcal{D})\big](b_i - b^*)\, dF(b) \\
&\qquad + \delta(1-p) \int_b \big[\tau^{\text{sep}}(b, \mathcal{D}) - \tau^{\text{sep}}(b, \mathcal{D}^i)\big] \left(b_i - \frac{c_H - pc_L}{1-p}\right) dF(b) \\
&= \delta(1-p) \int_b \big[\tau^{\text{pool}}(b, \mathcal{D}^i) - \tau^{\text{pool}}(b, \mathcal{D})\big](b_i - b^*)\, dF(b)
\end{aligned}$$

$$
\begin{aligned}
&+ \delta(1-p) \int_b \left[\tau^{\text{sep}}(b, \mathcal{D}) - \tau^{\text{sep}}(b, \mathcal{D}^i)\right](b_i - b^*)\, dF(b) \\
&= \delta(1-p) \int_b \Big[\big(\tau^{\text{pool}}(b, \mathcal{D}^i) - \tau^{\text{pool}}(b, \mathcal{D})\big) + \big(\tau^{\text{sep}}(b, \mathcal{D}) - \tau^{\text{sep}}(b, \mathcal{D}^i)\big)\Big](b_i - b^*)\, dF(b)\,,
\end{aligned}
$$

as desired.