The Politics of Collective Principals*

Vincent Anesi† Peter Buisseret‡

September 9, 2022

Abstract

A group of principals collectively and dynamically screens an agent. The principals hold heterogeneous and evolving values from the relationship. At each date, they use a collective decision rule to determine a joint offer to the agent; the principals may also amend the procedures governing how their joint offer is chosen. Our main result shows how decisive coalitions of principals voluntarily and permanently concentrate decision-making authority in a single principal. It shows that every equilibrium sequence of procedures converges to the dictatorship of a single principal.

*We are grateful to Lars Ehlers, Daniel Garrett, Ángel Hernando-Veciana, Nicolas Klein, Antoine Loeper, Ryan Tierney, Rafael Treibich, and seminar audiences at Essex, Hong Kong Baptist University, Leicester, University of Montreal, CREST (Paris), Carlos III (Madrid), and Southern Denmark University 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

Politics is replete with negotiations. States bargain to avert costly conflict (Brito and Intriligator, 1985; Dal Bó and Powell, 2009). International organizations offer market access in exchange for reduced tariffs, or provide loans in exchange for labor market reforms (Caraway, Rickard and Anner, 2012). Central governments concede policymaking responsibilities to localities to maintain their participation in a national union (Tommasi and Weinschelbaum, 2007). Interest groups offer campaign contributions in exchange for a politician's advocacy (Grossman and Helpman, 2001). Legislatures offer agencies budgets in exchange for policy outputs (Niskanen, 2017).

These negotiations are studied theoretically as *screening* problems, in which a 'principal' makes offers to an 'agent', and the principal is uncertain about the agent's preferences. Uncertainty could concern a state's costs of fighting (Brito and Intriligator, 1985), a government's domestic opposition to free trade (Caraway, Rickard and Anner, 2012), or a bureaucrat's costs of providing a service (Banks and Weingast, 1992). Existing theoretical work conceives of the principal as a single, unitary actor. Our paper, instead, studies contexts in which the agent interacts with a *collective principal* composed of multiple actors that "come to a joint decision (according to some rule) and then enter into a single contract with an agent" (Lake and McCubbins, 2006, 361)

Collective principals are ubiquitous: legislatures consist of representatives,[1] international organizations comprise member states, and interest groups are associations of individuals and groups. Nielson and Tierney (2009) identify the collective principal as "the most common type of principal that we observe in the study of politics" (p. 5). Yet this real-world ubiquity contrasts strikingly with the absence of any formal-theoretic study. Our paper takes the first steps towards developing a theory of collective principals in a dynamic agency framework.

We focus on the principals' collective decision-making over their joint offer, and also their collective decision-making concerning the procedures that determine those offers. We do so because collective choice is central to understanding collective principals. As Tommasi and Weinschelbaum (2007) note, their defining feature is that contracts "are signed collectively through the aggregation of some actions of principals such as voting" (p. 383). Lyne and Tierney (2002) further observe that the rules governing aggregation are also within the principals' determination: the conclude that "[t]he key problem... is how to model the decision rules that determine

[1] As Gailmard (2012) puts it: *Congress is a "they," not an "it"*.

how the members of the collective principal will come to a joint decision" (p. 59).

Problems of joint decision-making are as practical as they are conceptual. Kiewiet and McCubbins (1991) worry that collective principals may be vulnerable to "social choice instability", that "a collective principal may be unable to announce a single preference", and that "[a] subset of the membership may strategically manipulate the decision-making process" (p. 27). Strom (2000) similarly views "preference aggregation" as the central challenge facing collective principals (p. 268). And, since "[c]ontracting with the agent is contingent on mutual agreement among members of the collective" (Graham, 2015, 168), conflicts amongst the principals are likely to shape their common interaction with the agent.

To address these issues, our framework studies a long-run (infinite-horizon) relationship between a group of principals and an agent. 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 the agent's concession, but they are uncertain about the agent's cost of conceding. For example, the principals could be member states of a customs union, offering a non-member (the agent) reduced tariffs in exchange for market access. The union's members may be uncertain about the non-member's domestic opposition to trade liberalization. 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.

Our analysis therefore incorporates collective choice between the principals into a dynamic principal-agent setting. Our main focus is on how the principals' procedures evolve in the long-run. Throughout, we assume that agents are sufficiently motivated by short-run outcomes, in the sense that their discount factors are not too high.

We first characterize negotiation outcomes for a given procedure and (common) belief about the agent. The principals choose from a finite set of offers—one offer for each possible type of agent. The most generous offer extracts surplus from the agent with the highest-possible cost. Since all types accept this offer, we call this the *pooling* offer; any other *partially separating* offer is rejected by some agent-types. We verify that the principals' induced preferences over this set satisfy a form of single-peakedness, yielding a non-empty core of their collective-choice problem. The amendment agenda game yields a unique selection from the core for any procedure, and thus a unique prediction about the offer the principals put to the agent.

We then characterize how the collective decision-making context shapes the principals' preferences over offers. In each period, for any belief and procedure, we obtain a cut-off benefit from an agreement above which a principal favors the pooling contract over any partially separating contract. This cut-off broadly reflects a principal's incentive to screen the agent. We compare the cut-off under any procedure that assures a principal of her most-preferred offer in every period—a 'dictatorship'—to any procedure in which she is *not* a dictator. We show that a principal's pooling cut-off is always *higher* in the class of procedures where she is not a dictator.

To see why, recognize that if the principals offer a pooling contract, today, the agent's acceptance reveals nothing about her type. Tomorrow, a decisive high-benefit principal that wants to secure the agent's agreement may impose the pooling offer, again. This harms a low-benefit principal who prefers to gamble on an agreement with fewer concessions. Alternatively, a decisive low-benefit principal may prefer to gamble on the agent's willingness to accept fewer concessions. This harms a high-benefit principal if the agent subsequently rejects. Suppose, instead, the principals make an offer that all agent types but the very highest accept. If the agent accepts, the principals learn she does *not* have high costs. This reduces the most generous (pooling) offer *any* future principal wants to make, regardless of her benefit from agreement. In turn, this protects a future low-benefit principal against a decisive high-benefit principal. If the agent instead rejects, the principals learn her type, ensuring future agreements with an appropriately targeted offer. This protects a future high-benefit principal against a decisive low-benefit principal.

In sum: collective learning reduces the principals' mis-alignment over offers, insuring today's principal against her lack of future decision-making power and allowing her to influence outcomes indirectly, even when she is not in a decisive coalition of principals. The collective choice setting therefore intensifies incentives to screen the agent.

Screening the agent more aggressively nonetheless raises the risk that negotiations fail. This is always inefficient. We therefore study how the principals' decision-making processes evolve in order to mitigate this risk. Recall that the principals can amend the inherited procedure in every period before they interact with the agent. One important class of procedures is a dictatorship, which guarantees that some principal always imposes her preferred offer to the agent on the remaining principals. Dictatorships can be 'formal', via a voting rule that explicitly endows a singular principal the right to approve offers. However, we can also have 'informal' dictatorships; these procedures may not appear to concentrate power—for example, they may feature simple majority rules with veto power, or quota rules. When combined with appropriately specified agenda control, nonetheless, they ensure that one principal always secures her preferred outcome.

Our main result is Theorem 1: any equilibrium sequence of procedures converges to either formal or informal dictatorship almost surely. To see why, suppose that in a given period, a group of high-benefit principals prioritize agreement with the agent but cannot unilaterally impose the pooling offer under the inherited standing procedure. This also means that they cannot unilaterally change the collective-decision making process. Suppose, however, that there is another marginal principal who would support the pooling offer if she were a dictator, but prefers partial separation under any other procedure that does not ensure her preferred outcome in that and all future periods. If the high-benefit principals together with this marginal principal can amend the procedure, the former may support concentrating power in the latter to obtain their preferred outcome, today.

Our result does not imply an immediate transition to dictatorship. Some low-benefit principals may suggest other non-dictatorial rules that nonetheless assure the high-benefit principals of imposing the pooling offer, today. We illustrate some of the possible dynamics through which power gradually concentrates over time. These dynamics nonetheless tend inexorably towards dictatorship.

Other Related Work. Ours is not the first paper to integrate the principal-agent framework with the political economy of collective decision-making (e.g., Laffont, 2000 and Grossman and Helpman, 2001). Existing theoretical work on delegation with multiple principals nonetheless exclusively focuses on common agency environments (Bernheim and Whinston, 1986, Grossman and Helpman, 1994 and Gailmard, 2009) in which the principals non-cooperatively offer distinct and competing contracts to a single agent, or multiple agents (Prat and Rustichini, 2003). 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): in every period, no decisive coalition of principals would be better off selecting a different offer, taking into consideration the dynamic consequences of that deviation. 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 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 insufficient 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.

Finally, our paper studies the structure of authority in organizations. Our main result res-

[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.

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

onates with Robert Michels' organizational dictum—his 'Iron Law'—that "the formation of oligarchies within the various forms of democracy is the outcome of organic necessity" (Michels, 1959, 418). Concentration of authority arises in our setting not to improve coordination or communication, but as a commitment to make a principal the residual claimant of their common information.

2. 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 $a^t\big[b_i^t - u(x^t)\big]$, 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 $a^t\big[u_0(x^t) - c^t\big]$, 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 $0 < c_1 < \cdots < c_K < u_0(\hat{x}_0)$, 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 $\{\underline{k}, \dots, K-1\}$.[4]

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

[4] 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.

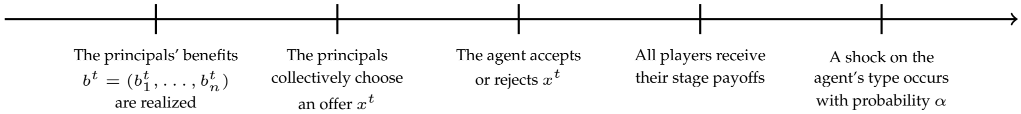

Figure 1 – Timing in each period $t = 1, 2, \ldots$.

from C according to p^0 with probability $\alpha \in (0, 1)$. Otherwise, the agent's type remains unchanged.

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.[5] 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 1.

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 (exogenously 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

[5] Alternatively, we could assume that $F_i\left[u(u_0^{-1}(c_K))\right]$ is sufficiently small for all i.

satisfy these conditions is denoted by $\mathcal{P}$, with generic element $\wp = (\lambda, \mathcal{D})$.

Figure 2 illustrates the collective decision-making process. We describe each phase in detail.

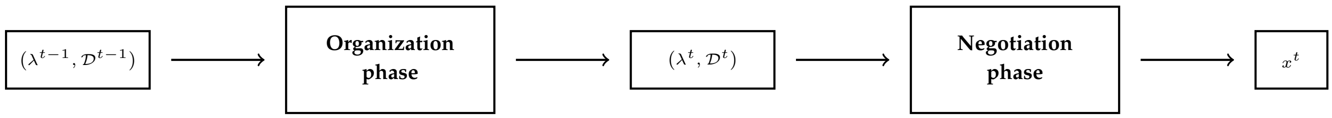

Figure 2 – The principals' collective decision-making process.

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.*

3. Preliminary Results

We begin with some useful preliminary results. These results establish equilibrium existence, as well as characterizing the outcome of any negotiation phase, for a given period-t procedure.

Lemma 1. *An equilibrium exists.*

Our first characterization result shows that all equilibria of the negotiation phase have a simple structure.

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:*[6]

(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 $c^\ell \leq c^k$.

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.

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

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 $b=(b_1,\dots,b_n)$, 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

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

for each $k=1,\dots,m$, 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 (1) 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 $\overline{x}^1,\dots,\overline{x}^m$ 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

$$\max_x U_{\iota_1}^{\phi}(x\mid p,b,\lambda,\mathcal{D}), \text{ subject to } x\in\mathcal{K}^{\phi}(p,b,\lambda,\mathcal{D})\,; \tag{2}$$

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

$$\overline{x}^m=\arg\max_{x\in X}U_i^{\phi}(x\mid p,b,\lambda,\mathcal{D}) \tag{3}$$

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

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.

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. Nonetheless, 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, the preferred offer of any principal that is made a veto player under voting rule $\mathcal{D}$ 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.

4. Collective versus Individual Incentives to Learn

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 β_i. Cut-off β_i can be loosely interpreted as reflecting a principal i's incentive to learn the agent's type. We now show how different procedures shape this incentive. We proceed by way of an example in which the principals' common belief p places positive probability on three agent-types, $C = \{c_1, c_2, c_3\}$. Suppose, initially, that the interaction proceeds over two periods: 1 and 2.[7]

Lemma 2 states that in every period and under any procedure and belief the principals' highest offer leaves the highest possible type c^m zero rents: $\overline{x}^m = u_0^{-1}(c_m)$. At belief p, principal i prefers the pooling offer to an offer that only agent types $\{c_1, c_2\}$ accept if and only if

$$(1-\delta)\big[b_i - u(\overline{x}^3)\big] + \delta V_i^{\phi}(p;\wp) \geq (1-\delta)(1-p(c_3))[b_i - u(\overline{x}^2(p,\phi,\wp))] + \delta\alpha V_i^{\phi}(p;\wp)$$
$$+\, \delta(1-\alpha)\big[(1-p(c_3))V_i^{\phi}(p_2^-;\wp) + p(c_3)V_i^{\phi}(p_2^+;\varphi)\big], \qquad (4)$$

where

$$p_2^-(c) \equiv \begin{cases} \frac{p(c)}{p(c_1)+p(c_2)} & \text{if } c \leq c_2 \\ 0 & \text{if } c = c_3 \end{cases} \quad \text{and} \quad p_2^+(c) \equiv \begin{cases} 1 & \text{if } c = c_3 \\ 0 & \text{if } c < c_3. \end{cases}$$

The LHS of (4) is principal i's payoff from the pooling offer: all agent-types accept, and belief p persists to the second period, regardless of whether there is any shock to the agent's type. The RHS of (4) is i's payoff from an offer $\overline{x}^2$ that separates types $\{c_1, c_2\}$ from $\{c_3\}$. With probability α, the agent's type is redrawn, and with probability $1-\alpha$, the agent's type persists to the second period. If the agent accepted the period-1 offer $\overline{x}^2$, with probability $1 - p(c_3)$, the principals' period-2 belief shifts to p_2^-. With probability $p(c_3)$ the agent rejects the offer, and the principals learn that the agent's type is c_3.

Condition (4) is necessary for principal i to prefer the pooling offer. Because the principals' preferences over offers are quasi-single-peaked, condition (4) is also sufficient for i to prefer the

[7] The finite horizon simplifies our example by ensuring that the agent's terminal period-2 (static) incentive constraints are invariant across all procedures and beliefs.

pooling offer to any other offer. So, threshold $\beta_i^\phi(p;\wp)$ solves (4) with equality:

$$\begin{aligned}\beta_i^\phi(p;\wp) = &\frac{1}{p(c_3)}\big[u(\overline{x}^3) - u(\overline{x}^2(p,\phi,\wp))(1-p(c_3))\big] \\ &+ \frac{\delta}{1-\delta}\frac{1-\alpha}{p(c_3)}\big[p(c_3)V_i^\phi(p_2^+;\wp) + (1-p(c_3))V_i^\phi(p_2^-;\wp) - V_i^\phi(p;\wp)\big].\end{aligned} \tag{5}$$

We compare β_i derived in expression (5) under two different procedures: a dictatorship $\wp^j$ of principal j versus a dictatorship $\wp^i$ of principal i. For any equilibria ϕ and φ, (5) yields:

$$\beta_i^\phi(p;\wp^j) - \beta_i^\varphi(p;\wp^i) \propto (1-\delta)(1-p(c_3))\big[u(\overline{x}^2(p,\varphi,\wp^i)) - u(\overline{x}^2(p,\phi,\wp^j))\big] \tag{6}$$

$$+\,\delta(1-\alpha)\big[p(c_3)V_i^\phi(p_2^+;\wp^j) + (1-p(c_3))V_i^\phi(p_2^-;\wp^j) - V_i^\phi(p;\wp^j)\big] \tag{7}$$

$$-\,\delta(1-\alpha)\big[p(c_3)V_i^\phi(p_2^+;\wp^i) + (1-p(c_3))V_i^\varphi(p_2^-;\wp^i) - V_i^\varphi(p;\wp^i)\big]. \tag{8}$$

The first line (6) is the difference in the principals' period-1 cost of separating types $\{c_1, c_2\}$ from $\{c_3\}$ under i's dictatorship $\wp^i$ versus j's dictatorship $\wp^j$. These incentive costs may differ because the agent anticipates different future rents depending on which principal is a dictator in the future. The second and third line capture the difference in principal i's expected period-2 benefit from partly screening the agent under $\wp^j$, versus $\wp^i$.

We begin with the difference of the second and third line. In (terminal) period 2 an agent type c_k accepts offer $\overline{x}^k = u_0^{-1}(c_k)$. If the principals made a period-1 pooling offer, principal i's period-2 payoff from an offer $\overline{x}^k$ is given by expression (1), evaluated at belief p and $\delta = 0$:[8]

$$U_i(u_0^{-1}(c_k) \mid p, b_i) \equiv \big[b_i - u(\overline{x}^k)\big]\sum_{\ell=1}^{k} p(c_\ell). \tag{9}$$

If the principals instead made a period-1 partially separating offer that the agent rejected—and absent any shock—i's period-2 payoff from offer $\overline{x}^k$ for $k \leq 2$ is:

$$U_i(u_0^{-1}(c_k) \mid p_2^-, b_i) \equiv \big[b_i - u(\overline{x}^k)\big]\frac{\sum_{\ell=1}^{k} p(c_\ell)}{1-p(c_3)} = \frac{U_i(u_0^{-1}(c_k) \mid p, b_i)}{1-p(c_3)}. \tag{10}$$

Comparing (10) with (9), we see that principal i's period-2 preferences over offers $\overline{x}^1$ and $\overline{x}^2$

[8] We omit references to equilibrium ϕ and procedure (λ, $\mathcal{D}$) since i's period-t preferences depend on these only through continuation payoffs, which vanish with $\delta = 0$.

coincide under beliefs p and $\tilde{p}$. Let $B^i_k(p)$ denote the set of benefits b_i such that $k \in \{1,2,3\}$ maximizes $U_i(\overline{x}^k \mid p, b_i)$. We conclude that in any period-2 event $b_i \in B^i_k(p)$ and $b_j \in B^j_l(p)$ for $k,l \leq 2$, the period-2 negotiation outcome under any procedure coincide at both p and p_2^-. In other words, given dictatorship $\wp^r$ in which principal $r \in \{i,j\}$ is a dictator, differences in beliefs p and p_2^- only impact period-2 outcomes in the event that $b_r \in B^r_3(p)$. Finally, notice that in every equilibrium under any procedure, the principals fully extract surplus from the highest-possible type in the support of their beliefs. This implies that $V_i^\phi(p_2^+; \wp^j)$ and $V_i^\phi(p_2^+; \wp^i)$ coincide. We can therefore re-write (7)

$$(1-\alpha)\delta \Pr(b_j \in B^j_3(p))\mathbb{E}_i\big[U_i(u_0^{-1}(c_2) \mid p, b_i) - U_i(u_0^{-1}(c_3) \mid p, b_i)\big],$$

and likewise (8):

$$(1-\alpha)\delta \int_{b_i \in B^i_3(p)} \big[U_i(u_0^{-1}(c_2) \mid p, b_i) - U_i(u_0^{-1}(c_3) \mid p, b_i)\big]\, dF(b_i),$$

so that the difference of (7) and (8) is (up to a constant)

$$\begin{aligned}\Pr\big(b_j \in B^j_3(p)\big) &\int_{b_i \notin B^i_3(p)} \big[U_i(u_0^{-1}(c_2) \mid p, b_i) - U_i(u_0^{-1}(c_3) \mid p, b_i)\big]\, dF(b_i) \\ &- \Pr\big(b_j \notin B^j_3(p)\big) \int_{b_i \in B^i_3(p)} \big[U_i(u_0^{-1}(c_2) \mid p, b_i) - U_i(u_0^{-1}(c_3) \mid p, b_i)\big]\, dF(b_i) > 0. \qquad (11)\end{aligned}$$

The first integral is positive and the second integral is negative under quasi-single-peakedness of i's preferences over offers. We conclude that principal i's learning benefit from screening the agent under principal j's dictatorship strictly exceeds i's corresponding benefit when she is the dictator.

The intuition is that screening out high-cost agent types reduces future mis-alignment between the principals. In the event that $b_j \in B^j_3(p)$ and $b_i \notin B^i_3(p)$, principal j's main priority is to secure the agent's agreement by making the pooling offer that all types accept; principal i instead prefers to take her chance with a less generous offer that nonetheless risks the agent's rejection. When j has decision-making power, she therefore imposes an offer that is too generous from i's perspective. By screening out the highest-cost type, i insures herself against the imposition of j's preferred offer by lowering the most generous offer j would be prepared to make in the future.

We showed that principal i's value of learning the agent's type is higher under $\wp^j$ than $\wp^i$, but we must also account for the possible costs of learning under these different procedures. These incentive costs are reflected in (6). While these costs could rise, we show that they nonetheless do not increase faster than the benefits. To do so, we revert to the infinite horizon, replacing "periods 1 and 2" in our example with "periods t and $t+1$." For any period-t procedure, suppose the principals' period-t offer separates types $\{c_1, c_2\}$ from $\{c_3\}$. Routine arguments establish that this offer, $\overline{x}^2$, is determined by type c_2's binding incentive constraint in equilibrium:

$$(1-\delta)\big[u_0(\overline{x}^2)-c_2\big]+\delta\begin{bmatrix}\text{type } c_2\text{'s expected}\\ \text{continuation payoff}\\ \text{from accepting } \overline{x}^2\end{bmatrix}=(1-\delta)\times 0+\delta\begin{bmatrix}\text{type } c_2\text{'s expected}\\ \text{continuation payoff}\\ \text{from rejecting } \overline{x}^2\end{bmatrix},$$

so that her period-t rent is

$$(1-\delta)\big[u_0(\overline{x}^2)-c_2\big]=\delta\begin{bmatrix}\text{Expected difference in type } c_2\text{'s}\\ \text{continuation payoffs from}\\ \text{rejecting and accepting } \overline{x}^2\end{bmatrix}. \tag{12}$$

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:

$$(1-\delta)\big[u_0(\overline{x}^2)-c_2\big]=\delta(1-\alpha)\begin{bmatrix}\text{Expected difference in type } c_2\text{'s continuation}\\ \text{payoffs from rejecting and accepting } \overline{x}^2\\ \text{conditional on no shock between } t \text{ and } t+1\end{bmatrix}.$$

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

(1) if type c_2 *accepts* $\overline{x}^2$, 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_2 *rejects* $\overline{x}^2$, then conditional on no shock between t and $t+1$, the principals assign probability one to c_3, and offer $u_0^{-1}(c_3)$ in $t+1$. This observation, again, is invariant across procedures, since the principals unanimously prefer this offer.

Hence, any wedge in the type c_2 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.

The following proposition generalizes the insights from this example. It shows that "j's dictatorship $\wp^j$" in our example can be replaced with "any procedure $\wp$ in which i is not a dictator". It is easy to see that in any equilibrium of a continuation game that begins under some principal i's dictatorship, she remains a dictator in all future periods—possibly under different procedures.

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.

5. Evolution of Collective Choice Procedures

Earlier, we characterized the principals' joint offer to the agent in the negotiation phase, given the inherited procedure from the organization phase. We now show which procedures emerge, over time. Our main result unearths a striking tendency towards the unfettered concentration of power within the collective principals.

Theorem 1. *Every equilibrium sequence of procedures $\left\{(\lambda^t, \mathcal{D}^t)\right\}$ converges to a dictatorship almost surely.*

The theorem highlights sequences of decisive coalition voluntarily cede decision-making authority in a process that inexorably tends towards the complete concentration of power in a single principal.

We illustrate the theorem with an example in which the principals are $N = \{1, 2, 3, 4, 5\}$. 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 1, 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 $\underline{P} \equiv \min\{P(\lambda, \mathcal{D}) : (\lambda, \mathcal{D}) \in \mathcal{P}_E\}$.[9]

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.,

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

where $\beta_i^{\phi}(p^0; \lambda, \mathcal{D})$ is defined in Lemma 3. Proposition 1 yields that $\overline{\beta}_i > \underline{\beta}_i$, where $\underline{\beta}_i$ is i's pooling threshold when she is a dictator. For illustration, suppose the ongoing procedure at the start of period t is simple majority rule, and a shock to the agent's type yields period-t belief p^0. Let F_1 be the event—illustrated in Figure 3—in which b^t's realization is such that

(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

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

Part (i) yields that principals 1 and 2's short-term preference is to partially screen the agent, but part (ii) implies that principal 4 and 5's short-term preference is the pooling offer. Following Proposition 1, part (iii) states that principal 3 prefers the pooling offer if she is a dictator; otherwise, she too prefers to partially screen the agent.

In period t's organization phase, a simple majority can amend the procedure to one that commits the principals to the pooling offer in the negotiation phase. If $\delta > 0$ is not too large, princi-

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

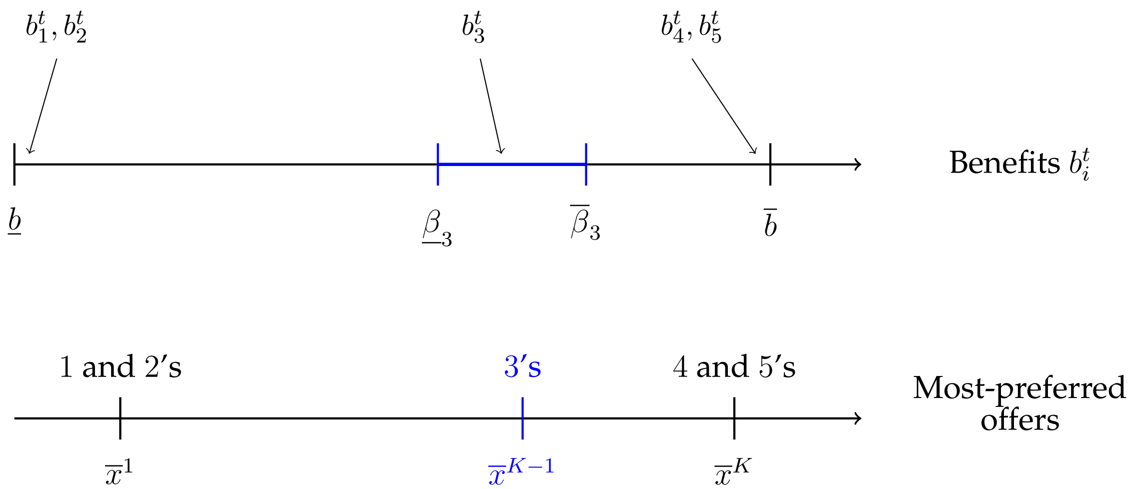

Figure 3 – The realization of principals' period-t benefits in event F_1.

pals $\{3,4,5\}$ would prefer to do so by making 3 a dictator. To see why, recognize that principal 3 strictly benefits from securing her preferred outcome in every future period, while principal 4's and 5's main priority is to secure the agent's agreement, today. We conclude that if δ is small enough, the outcome of the period-t negotiation phase *must* be the pooling offer—otherwise, one of the principals 3, 4, or 5 would have a profitable deviation at the organization phase.

Nonetheless, a dictatorship of principal 3 is not the *only* procedure that commits the principals to a period-t pooling offer. According to Lemma 3, 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., $\iota_1 \in \{4,5\}$.

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 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

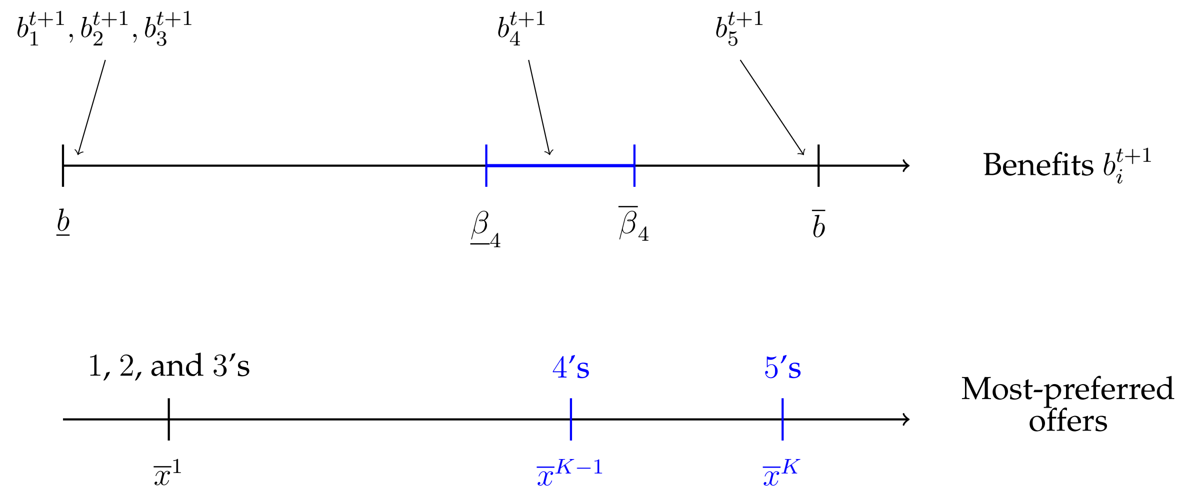

Figure 4 – The realization of principals' benefits in period $t+1$ in event F_2.

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 4—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 $\bar{b}$, 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 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, our argument also applies if the inherited rule is unanimity. To make this point concrete, we can amend the event F_1 in Figure 3 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.

6. Concluding Comments

We study the politics of collective principals, who collectively make a common offer to the agent. While collective principals abound in real-world political and economic contexts, our paper is the first to study them, theoretically. We asked how conflicts between principals shape their interaction with the agent. And, we asked how the principals self-govern over the course their interaction with the agent.

We unearth a tendency towards excessive collective learning that emerges across *all* non-dictatorial procedures. We then show that decisive coalitions of principals mitigate this tendency by successive rule changes that ultimately concentrate all decision-power in the hands of a single principal. This outcome obtains regardless of the initial procedure: it holds even if the principals initially operate under unanimity rule. Our results apply to settings where the principals and the agent care enough about short-term outcomes, i.e., where the discount factor is not too large. We believe that this case is relevant in our motivating applications.

We hope our framework spurs further work on collective principals. While we focused on screening, other principal-agent environments may be relevant. Collective principals that need to monitor or sanction the agent may face free-riding incentives; changes in the identity of decisive principals across periods may also impede their commitment power. How these issues might be mitigated or exacerbated by collective decision-making procedures is left to future research.

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.

Banks, Jeffrey S and Barry R Weingast. 1992. "The political control of bureaucracies under asymmetric information." *American Journal of Political Science* pp. 509–524.

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.

Brito, Dagobert L and Michael D Intriligator. 1985. "Conflict, war, and redistribution." *American Political Science Review* 79(4):943–957.

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.

Caraway, Teri L, Stephanie J Rickard and Mark S Anner. 2012. "International negotiations and domestic politics: The case of IMF labor market conditionality." *International Organization* 66(1):27–61.

Dal Bó, Ernesto and Robert Powell. 2009. "A model of spoils politics." *American Journal of Political Science* 53(1):207–222.

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.

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

Gailmard, Sean. 2009. "Multiple principals and oversight of bureaucratic policy-making." *Journal of Theoretical Politics* 21(2):161–186.

Gailmard, Sean. 2012. "Accountability and principal-agent models." *Chapter prepared for the Oxford Handbook of Public Accountability* .

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

Graham, Erin R. 2015. "Money and multilateralism: how funding rules constitute IO governance." *International Theory* 7(1):162–194.

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.

Kiewiet, D Roderick and Mathew D McCubbins. 1991. *The logic of delegation*. University of Chicago 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.

Lake, David A and Mathew D McCubbins. 2006. *The logic of delegation to international organizations*. Cambridge University Press Cambridge.

Lyne, Mona and Michael Tierney. 2002. Variation in the structure of principals: conceptual clarification for research on delegation and agency control. In *Memo presented at the Conference on Delegation and International Organizations*. pp. 3–4.

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

Nielson, Daniel L and Michael J Tierney. 2009. "Principals and Interests: Common Agency and Multilateral Development Bank Lending." Prepared for the 2010 Political Economy of International Organizations meeting, Washington, D.C., January 28-29, 2010.

Niskanen, William A. 2017. *Bureaucracy & representative government*. Routledge.

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.

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

Strom, Kaare. 2000. “Delegation and accountability in parliamentary democracies.” *European Journal of Political Research* 37(3):261–290.

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.

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 $k \in \{1, \ldots, K\}$, 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 $S_k^- \equiv \{c_1, \ldots, c_k\} \cap \operatorname{supp}(p)$ and $S_k^+ \equiv \{c_{k+1}, \ldots, c_K\} \cap \operatorname{supp}(p)$; 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 $\{c^1, \ldots, c^m\}$, let $\beta_p \colon \{1, \ldots, m-1\} \to \mathbb{R}$ 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 $k \in \{1, \ldots, m-1\}$, 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 $\underline{b} < \beta_p(k) < \bar{b}$, 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 $\overline{\beta} \equiv \max_p \beta_p(m-1) < \bar{b}$, 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

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

$$+ \delta\alpha W_{i,0},\, k \neq 0, m\ ,$$

$$U_i(m \mid b_i) \equiv (1-\delta)\big[b_i - u(x_m)\big] + \delta\big[(1-\alpha)W_i(p) + \alpha W_{i,0}\big]\ ,$$

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 $U^p\colon \{0,1,\dots,m\} \times B \to \mathbb{R}$, 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 $k \in \{1,\dots,m\}$ and $b \in B$. By definition, for any $k \in \{1,\dots,m-1\}$, 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 $b \notin \{\beta_p(1),\dots,\beta_p(m)\}$.

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 $k \in \{1,\dots,m-1\}$ — 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 $\{0,1,\dots,m\}$, $1 \leq m \leq K$, and any profile of utility functions $f = (f_1,\dots,f_n)$ on $\{0,1,\dots,m\}$, we denote by $\mathrm{Core}(m,f)$ 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 $\{0,1,\dots,m\}$, 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 $\mathrm{Core}(m, f)$, 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, $\mathrm{Core}(m, f)$ is nonempty, and all the alternatives in $\mathrm{Core}(m, f)$ must be adjacent. It follows that each principal i has a unique ideal alternative in $\mathrm{Core}(m, f)$, 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 $\mathrm{Core}(m, f)$. 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. □

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 $\mathcal{V} \equiv \big[0, u_0(\hat{x}_0) - c_1\big]^L \times$

$\left[0, u_0(\hat{x}_0) - c_K\right]^L \times \left[\underline{b} - u(\hat{x}_0), \overline{b}\right]^{nL}$. 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 $\nu_{0,k} \in \left[0, u_0(\hat{x}_0) - c_k\right]^L$, for each $k = 1, \ldots, K$; and $\nu_i \in \left[\underline{b} - u(\hat{x}_0), \overline{b}\right]^L$, 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 $(1-\delta)\left[b_i^t - u(x_k)\right] + \delta \nu_{i,l'}$; otherwise, she receives a stage-payoff of $(1-\delta)\left[b_i^t - u(x_k)\right]$, 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 $F_1, \ldots, F_n$ 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 $V_i^p(\lambda_\ell, \mathcal{D}_\ell \mid \nu_0, \nu_1, \ldots, \nu_n)$ denote principal i's equilibrium payoff. For future reference, we also define $V_{0,k}^p(\lambda_\ell, \mathcal{D}_\ell \mid \nu_0, \nu_1, \ldots, \nu_n)$ as the corresponding expected payoff of the passive type-c_k agent.

Now fix $m = 2, \ldots, K$. Suppose that for every $p' \in \Delta_{p^0}$ with $\left|\mathrm{supp}(p')\right| \leq m-1$, we have defined a game $\mathcal{G}^{p'}(\lambda_\ell, \mathcal{D}_\ell \mid \nu_0, \nu_1, \ldots, \nu_n)$, $\ell = 1, \ldots, L$, and corresponding continuation payoffs $V_i^{p'}(\lambda_\ell, \mathcal{D}_\ell \mid \nu_0, \nu_1, \ldots, \nu_n)$ and $V_{0,k}^{p'}(\lambda_\ell, \mathcal{D}_\ell \mid \nu_0, \nu_1, \ldots, \nu_n)$, 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 $\mathrm{supp}(p) = \{c_1, \ldots, c_m\}$.

Observe that for every $k = 1, \ldots, m-1$, $\left|\mathrm{supp}(p^{k-})\right| \leq m-1$ and $\left|\mathrm{supp}(p^{k+})\right| \leq m-1$ and therefore, $V_i^{p^{k-}}(\lambda_\ell, \mathcal{D}_\ell \mid \nu_0, \nu_1, \ldots, \nu_n)$, $V_{0,k'}^{p^{k-}}(\lambda_\ell, \mathcal{D}_\ell \mid \nu_0, \nu_1, \ldots, \nu_n)$, $V_i^{p^{k+}}(\lambda_\ell, \mathcal{D}_\ell \mid \nu_0, \nu_1, \ldots, \nu_n)$, and $V_{0,k'}^{p^{k+}}(\lambda_\ell, \mathcal{D}_\ell \mid \nu_0, \nu_1, \ldots, \nu_n)$, are well-defined for all i, k', and ℓ. This allows us to (implicitly) define the policy $\chi_k(\lambda_\ell, \mathcal{D}_\ell \mid \nu_0, \nu_1, \ldots, \nu_n)$ as the unique solution x to

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

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

for each $k \leq m-1$, and $\chi_m(\lambda_\ell, \mathcal{D}_\ell \mid \nu_0, \nu_1, \ldots, \nu_n) \equiv x_m$. 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 $\mathcal{G}^p(\lambda_\ell, \mathcal{D}_\ell \mid \nu_0, \nu_1, \ldots, \nu_n)$ 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) The negotiation phase takes place as in the main game, but the principals are constrained to choose offers from the set $\big\{\chi_k(\lambda_{l'}, \mathcal{D}_{l'} \mid \nu_0, \nu_1, \ldots, \nu_n)\big\}_{k=1,\ldots,m}$. 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 $\big\{\chi_k(\lambda_\ell, \mathcal{D}_\ell \mid \nu_0, \nu_1, \ldots, \nu_n)\big\}_{k=1,\ldots,m}$. 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 $\big\{\chi_k(\lambda_\ell, \mathcal{D}_\ell \mid \nu_0^*, \nu_1^*, \ldots, \nu_n^*)\colon k = 1, \ldots, K \;\&\; \ell = 1, \ldots, L\big\}$ 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 $V_{0,\ell}(\cdot \mid c_k) \equiv V_{0,k}^p(\lambda_\ell, \mathcal{D}_\ell \mid \nu_0^*, \nu_1^*, \ldots, \nu_n^*)$. 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

$$(1-\delta)\big[u_0(x) - c\big] + \delta(1-\alpha)\big[V_{0,\ell}(p^{k-} \mid c) - V_{0,\ell}(p^{k+} \mid c)\big]$$

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

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 $\delta < \overline{\delta} \leq \overline{\delta}_2$, 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 $\left\{x \in X \colon \sigma(p, \lambda, \mathcal{D}, x \mid c_{m-1}) = 1 - \sigma(p, \lambda, \mathcal{D}, x \mid c_m) = 1\right\}$ is nonempty. Let $\overline{x}^{m-1}(p, \lambda, \mathcal{D}) \equiv \inf\left\{x \in X \colon \sigma(p, \lambda, \mathcal{D}, x \mid c_{m-1}) = 1 - \sigma(p, \lambda, \mathcal{D}, x \mid c_m) = 1\right\}$. 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 $\delta < \overline{\delta} \leq \overline{\delta}_3$, 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 $\widehat{B}^i_k$ be those for which k is principal i's ideal in $\{1, \ldots, m\}$. We then have

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

and

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

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

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

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 $\Pr(\widehat{B}^i_k \cap B_\ell) > 0$.

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

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

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

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

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

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

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. $\square$

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

$$\begin{aligned}(1-\delta)\big[b_i - u(\overline{x}^m)\big] + \delta(1-\alpha)V_i^{\phi}(p;\lambda,\mathcal{D}) \leq (1-\delta)\big[b_i - u(\overline{x}^{m-1})\big]p(S^-_{m-1}) \\ + \delta(1-\alpha)\big[p(S^-_{m-1})V_i^{\phi}(p^-_{m-1};\lambda,\mathcal{D}) + p(S^+_{m-1})V_i^{\phi}(p^+_{m-1};\lambda,\mathcal{D})\big]\ ,\end{aligned}$$

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

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

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

$$\begin{aligned}\hat{\beta}_i^{\phi^i}(p) \equiv & \big[(1-\delta)p(c^m)\big]^{-1}\Big[(1-\delta)\big[u(\overline{x}^m) - u(\hat{x}^{m-1})p(S^-_{m-1})\big] \\ & + \delta(1-\alpha)\big[p(S^-_{m-1})W_i^{\phi^i}(p^-_{m-1}) + p(S^+_{m-1})W_i^{\phi^i}(p^+_{m-1})\big] - W_i^{\phi^i}(p)\Big]\ ,\end{aligned}$$

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 $\left|V_0^{\phi}(p^{(m-1)+} \mid c_{m-1}) - W_0^{\phi^i}(p^{(m-1)+} \mid c_{m-1})\right| \leq \delta u_0(\hat{x}_0)$. 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 $\left|W_0^{\phi^i}(p^{(m-1)-} \mid c_{m-1}) - V_0^{\phi}(p^{(m-1)-} \mid c_{m-1})\right| \leq \delta u_0(\hat{x}_0)$. 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 $\delta < \overline{\delta} \leq \overline{\delta}_5 \equiv \kappa/\left[2u_0(\hat{x}_0)u'(\hat{x}_0)v_0'(\hat{x}_0)\right]$.

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 Theorem 1

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 $q \equiv \min\left\{q_i(\lambda) \colon i \in N, \lambda \in \Lambda, q_i(\lambda) > 0\right\}$. Then, there exists a sufficiently small $\hat{\delta}_{6,1} > 0$ such that

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

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.