



From equals to despots: The dynamics of repeated decision making in partnerships with private information [☆]

Vinicius Carrasco ^a, William Fuchs ^{b,c}, Satoshi Fukuda ^{d,*}

^a Department of Economics, PUC-Rio, Brazil

^b McCombs School of Business, UT Austin, USA

^c Universidad Carlos III Madrid, Spain

^d Department of Decision Sciences and IGIER, Bocconi University, Italy

Received 7 October 2009; accepted 9 March 2019

Available online 16 April 2019

Abstract

This paper considers an optimal renegotiation-proof dynamic Bayesian mechanism in which two privately informed players repeatedly have to take a joint action without resorting to side-payments. We provide a general framework which accommodates as special cases committee decision and collective insurance problems. Thus, we formally connect these separate strands of literature. We show: (i) first-best values can be arbitrarily approximated (but not achieved) when the players are sufficiently patient; (ii) our main result, the provision of intertemporal incentives necessarily leads to a dictatorial mechanism: in the long run the optimal scheme converges to the adoption of one player's favorite action. This can entail one agent becoming a permanent dictator or a possibility of having sporadic "regime shifts."

© 2019 The Author(s). Published by Elsevier Inc. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

JEL classification: D82; D86; D70

[☆] We have benefited from conversations with Suren Basov, Elchanan Ben-Porath, Simon Board, Alessandra Casella, Simone Cerreia-Vioglio, Martin Cripps, Sergio Firpo, Drew Fudenberg, Simone Galperti, Yingni Guo, Christian Hellwig, Johannes Hörner, Andreas Kleiner, Narayana Kocherlakota, Paolo Leonetti, Ferdinando Monte, Roger Myerson, Tymofiy Mylovanov, Ezra Oberfield, Aniko Öry, Alessandro Pavan, Phil Reny, Leo Rezende, Yuliy Sannikov, Mahyar Sefidgaran, Andy Skrzypacz, Nancy Stokey, and Bálažs Szentes. We would like to thank two helpful anonymous referees and an Associate Editor. Fuchs gratefully acknowledges support from the ERC Grant 681575.

* Corresponding author.

E-mail address: vnc@econ.puc-rio.br (V. Carrasco).

Keywords: Repeated collective decision making; Dictatorial mechanism; Dynamic Bayesian mechanism design; Renegotiation-proofness; Regime shifts

1. Introduction

There are many situations in which, repeatedly, a group of players have to take a common action, cannot resort to side-payments, and each period are privately informed about their preferences. Examples include many supranational organizations such as a monetary union or a common market. In the former, monetary policy must be jointly taken and, in the latter, a common tariff with the outside world must be adopted each period. At the national level, political coalitions must jointly and repeatedly decide on policy issues. Within firms, managers of different divisions often have to make a joint decision repeatedly.

We show players can improve the efficiency of their repeated collective decisions by using rules that result in time varying decision rights. This captures an informal notion of “political capital,” i.e., a player who pushes hard on a given decision spends his political capital and as a result can exert less influence on future decisions. Although the influence of different players vary over time, in the long run, one player ends up having all the decision rights, effectively dictating the actions to be taken by the group. The identity of the “dictator” can be constant over time or there can be “regime shifts” in which the identity of the player who dictates the actions changes. These “regime shifts” are infrequent but importantly very little time is spent in the transition.

Formally, we study the long-run properties of an optimal renegotiation-proof dynamic Bayesian mechanism without side-payments for a collective decision problem between two players.¹ Our setting accommodates any strictly concave utility function in joint action. Especially, our model can capture the dynamic insurance problem as a special case by interpreting the transfers as the joint action. Thus, an additional contribution of the paper is to bring the collective decision problem and the dynamic insurance problem under a common framework. We unify these two, previously disjoint, strands of literature.

We first show that efficiency can be arbitrarily approximated, but never attained, when players are sufficiently patient. In a repeated setting, as shown in the seminal work of Fudenberg et al. (1994), the promise of (equilibrium) continuation values play a similar role to the one side-payments play in static mechanism design problems. Indeed, in order to prove the approximate efficiency result, we construct continuation values that replicate the expected payments proposed by Arrow (1979) and d’Aspremont and Gérard-Varet (1979), that guarantee efficiency in standard static Bayesian mechanisms *with* transfers.

The difference between side-payments and continuation values is that the latter can only imperfectly transfer utility across players. In particular, to transfer continuation utility from one player to another in any period, joint actions for any future periods must be altered. When players are sufficiently patient (i.e., their (common) discount factor is close to one), their weighted current payoffs become insignificant relative to the continuation values. Hence, in order to guarantee truth-telling in the current period, continuation values have to vary only minimally, and the

¹ While the introduction of dynamics to the collective decision making enhances efficiency, the improved efficiency may be supported by inefficient future decision rules. Thus, an optimal mechanism may be subject to renegotiation. Among various concepts of renegotiation-proofness in infinite horizon games, we impose *internal renegotiation-proofness*, following Ray (1994). We restrict attention to two players as renegotiation-proofness among more than two players raises an issue of renegotiation at a sub-coalitional level.

associated efficiency losses from variation in continuation values are arbitrarily small in the limit as the discount factor tends to one. The attainment of full efficiency, however, would call for no variation in continuation values, and this is at odds with the provision of incentives needed for an efficient action to be taken. Hence, full efficiency is not attainable.

Although the limiting efficiency result is of interest, our main focus is on understanding the long-run properties of the optimal mechanism. As a benchmark, it is useful to keep in mind what the first-best mechanism would entail in the environment in which players' private information is their favorite actions. The first-best action would call for a constant weighted average of players' favorite actions. The problem with this decision rule is that players away from the center have an incentive to exaggerate their positions. If they expect the other types to be to the left (right) of them they would have an incentive to claim to be further right (left) and in that way bring the chosen decision closer to their preferred point. As a result, absent any additional tools to provide incentives, the only way one has to prevent players from lying is by making the decision rule less sensitive to their reports.²

When decisions are made repeatedly, players care about the future so that they can trade decision rights in the current action for decision rights in the future. More extreme types are given more weight in the current decision but they pay for it by having less influence on future decisions. In this way, the optimal mechanism can be more sensitive to extreme announcements of preferences.

As known from static mechanism design theory, once incentive compatibility constraints are taken into account, players' utilities have to be adjusted to incorporate the rents derived from their private information. The adjusted utility is referred to as *virtual* utility, and the optimal mechanism design problem amounts to maximizing the (weighted) sum of players' virtual utilities for each type (Myerson, 1981, 1984, 1985). In our dynamic setting, virtual utilities also play a key role. We show that the dynamics of efficient collective decision making are fully determined by: (i) a decision rule that, at each period, maximizes the (weighted) sum of players' virtual utilities and (ii) a process that governs the evolution of the weights given to players' virtual utilities. While players' Pareto weights in the first-best problem are constant over time, we show that their weights vary over time in the second-best problem in a way such that the ratio of players' weights increases in expectation over time.

The dynamics of the collective decision making process lead to our most interesting result. Continuation values vary from period to period reflecting players' weights.³ Continuation values increase (higher future bargaining power) for a player who reports a less extreme announcement, and decrease (lower future bargaining power) for a player who reports an extreme announcement. In the limit as time goes to infinity, only the preferences of one player are taken into account in the decision process. Thus, the optimal mechanism converges to a dictatorial one in the limit. In our numerical simulation of the insurance problem, we illustrate that the identity of the dictator is fully persistent, a result akin to the immiseration results in the literature. In contrast, in our simulation of the general collective decision problem, we observe "regime shifts" where the decision making process goes from one player having almost all the say to the other having almost all the say. Nonetheless, these regimes are persistent and switches take place infrequently.

² See Carrasco and Fuchs (2009) for an analysis of the static problem in which two ex-ante symmetric players have single-peaked preferences.

³ Our approach to the analysis of the optimal mechanism in the repeated game relies on the factorization results of Abreu et al. (1990) that a player's payoff can be split into a current value and a continuation value.

1.1. Related literature

This paper relates to two strands of literature: (i) the one on collective decision making without side-payments and (ii) the dynamic insurance literature. We discuss these in order below.

There has been a very extensive and old literature on collective decision making. We are most closely related to the models in which players' values are independent of each other. Thus each player's preferred action does not directly depend on the private information of the other players. Only recently has this literature started to consider the case where rather than only one decision there are potentially multiple ones to be taken. The question then is how the availability of many decisions is to be taken into account to improve efficiency.

The first paper in this regard is Casella (2005) who studies the problem of a committee that each period must take one of two possible decisions and players have a continuum of preference intensities for these two actions. She proposes the "storable vote" mechanism which endows players with one vote per period. Players can then either cast their vote or store it for the future. This effectively allows them to have more influence on the future decisions since they are taken by majority vote. In a two period setting, the storable vote mechanism can indeed enhance efficiency.⁴ Although appealing from a practical standpoint, the storable vote mechanism is not optimal.

Hortala-Vallve (2012) allows a continuum of preference intensities over issues in a static environment in which players know their preferences over all issues when making their decision. He looks at the "qualitative voting" mechanism, which endows players with a large number of votes and allows them to cast their votes over the issues according to their preferences. He shows this mechanism is optimal in the case of two issues and two or three players.

Drexler and Kleiner (2014) study committee decision making in which two privately informed players repeatedly make a binary decision without monetary transfers. They characterize conditions on prior distributions and classes of mechanisms among which a voting rule is optimal.

Jackson and Sonnenschein (2007) study a dynamic Bayesian mechanism without side-payments, which they call the "linking mechanism." Agents must budget their representations of preferences so that the frequency of reported preferences across issues mirrors the underlying distribution of preferences. They show that agents' incentives are to satisfy their budget by being as truthful as possible and that when the number of problems being linked goes to infinity this mechanism is arbitrary close to being efficient. Although efficient in the limit, their mechanism, unlike ours, is not optimal. Also, although both their and our mechanisms are approximately efficient, the small inefficiencies are caused by very different aspects of the mechanisms. In theirs, the inefficiency arises from the fact that for any finite number of problems the realized distribution does not exactly match the budget which is based on the underlying preference distribution. The inefficiency in our mechanism arises instead from variation in continuation values to provide incentives. None of these previous papers have a dictatorship result as we do.

A special case of our model can capture the problem of two players with income shocks trying to self-insure, where the common action is regarded as a transfer from one player to another. The problem between one agent and one principal has been looked at before by Thomas and Worrall (1990). Within the macro literature, this problem has been examined with a continuum of agents by Atkeson and Lucas (1992), Green (1987), and Phelan (1995). Phelan (1998) and Wang (1995)

⁴ Skrzypacz and Hopenhayn (2004, Section 3.1) use a similar "chips mechanism" to sustain collusion in a repeated auction environment. They numerically demonstrate that the chips mechanism converges to an optimal scheme.

study a finite number of players in a dynamic insurance setting.⁵ In a related environment, Friedman (1998) and Zhao (2007) study a repeated moral hazard problem in which a finite number of risk-averse players produce a perishable consumption good each period.

One important implication of private information on dynamic insurance problems is the immiseration result: an agent's continuation value diverges to negative infinity in the long run. Interestingly, unlike the other papers, Wang (1995) does not have the immiseration result but rather has a non-degenerate long-run distribution of consumption.⁶ Phelan (1998, p. 184) discusses Wang's result and argues that it is not due to a finite number of players but rather due to the fact that Wang imposes a lower bound on consumption and utility: "This makes clear that the results of Wang (1995) depend crucially on his assumption that $c \geq 0$." In Wang's model, any transfer by a player must be at most his endowment. We instead constrain transfers to be in some fixed compact set. The set could be within the set of transfers allowed in Wang or include the set of transfers allowed by Wang. In either case, it implies bounds both on consumption and utilities. The important difference is that, in our model, the action space is independent of the realized shocks. Thus, we show that the non-immiseration result in Wang (1995) does not depend on the bound of utilities itself but rather on the action space being determined by the realized types. This restriction makes it impossible to always transfer all of one player's endowment to the other in an incentive compatible manner, and thus the immiseration result is not supported as a truthful equilibrium.⁷

Dynamic insurance problems have the incentive structure in which players faced with the first-best insurance contract would have an incentive to claim to be of the lowest possible type to obtain the highest insurance payment. Instead, in our general setting, when faced by the first-best mechanism, although the players still have an incentive to exaggerate their types, they would not in general want to claim to be the lowest (or highest) possible type. The reason is that the stage game, at the interim stage, is no longer a zero-sum game. There is a possibility that both players want a similar action to be taken even after they learn their private information. Suppose for example that two players have their preferred actions drawn from the i.i.d. uniform distribution on $[-1, 1]$ and have a quadratic loss utility in the distance between their type and the joint action. A player with realized type 0.01 expects the other player to have a lower type and thus wants to exaggerate his type in order to bring the joint decision closer to his preferred action.⁸ However, he does not want to claim to be of type 1 because if the other player reports a high type this could lead to a joint decision close to 1. The fact that the level of conflict or congruence of preferences is not known allows for some cooperation to take place even in a one-shot game. In contrast, there is no scope for insurance in a one-shot version of the dynamic insurance problem.

The paper is organized as follows. Section 2 introduces the model. Section 3 characterizes the optimal renegotiation-proof dynamic Bayesian mechanism. Concluding remarks are provided in Section 4. All proofs are relegated to Appendix A. Appendix B, available online, provides the results of numerical simulations of the repeated collective decision and insurance models, respectively.

⁵ Phelan (1998, Section 3) assumes that there is a price-taking representative firm which can separately contract with each agent.

⁶ Phelan (1995) also has a non-degenerate long-run consumption distribution in a principal-agent environment with (one-sided) limited commitment.

⁷ In Phelan (1995), the limited commitment excludes the immiseration as an equilibrium value.

⁸ The first-best symmetric decision rule is simply the average of both types in this case.

2. Model

In each period $t \in \mathbb{N}$, a group of two players (denoted by $N := \{1, 2\}$) must take a joint action a from a compact convex set $\mathcal{A} := [\underline{a}, \bar{a}]$, where its interior \mathcal{A}° is not empty.⁹ Each period, each player i privately receives a preference shock (his private type) from a finite set Θ_i . Preference shocks are drawn from a probability density function $f_i(\cdot) > 0$ on Θ_i . They are assumed to be independent across players and i.i.d. over time for each player. We denote the set of type profiles by $\Theta := \Theta_1 \times \Theta_2$ and the joint density f on Θ by $f(\theta) := f_1(\theta_1) \times f_2(\theta_2)$.

Each player’s instantaneous utility $u_i : \mathcal{A} \times \Theta_i \rightarrow \mathbb{R}$ depends on a common action and his own private type, i.e., we deal with the case of private values, and side-payments are not allowed. We assume that u_i is continuously differentiable and strictly concave in common action $a \in \mathcal{A}$ for any given type. Each i has his unique best action $a_i(\theta_i) (\in \mathcal{A})$ which maximizes his utility $u_i(\cdot, \theta_i)$ for each type θ_i .¹⁰

Our model captures the following three standard environments as special cases: (i) a collective decision problem with single-peaked preferences, (ii) a dynamic insurance partnership model, and (iii) a multiplicative preference shock model. In the first environment, the players have to take a joint action repeatedly, where they have single-peaked preferences with their types being their most preferred actions, i.e., $a_i(\theta_i) = \theta_i$. A typical example is that each player has a quadratic utility function $u_i(a, \theta_i) = -(a - \theta_i)^2$ with $\Theta_1 = \Theta_2 \subseteq \mathcal{A}$.

Second, two players insure themselves against private endowment shocks. We can reinterpret the common action as the “transfer” from player 1 to 2 (this is possible because the mechanism itself does not allow side-payments). Unlike Atkeson and Lucas (1992), Green (1987), and Thomas and Worrall (1990), transfers are restricted in the compact set $\mathcal{A} = [\underline{a}, \bar{a}]$.¹¹ Player 1’s (resp. 2’s) transfer is restricted to at most \bar{a} (resp. $-\underline{a}$) each time. Note that Wang (1995) has a more stringent feasibility condition that each player can only transfer his own endowments, that is, $-\theta_2 \leq a(\theta_1, \theta_2) \leq \theta_1$ for each $(\theta_1, \theta_2) \in \Theta$. In our setting, the set \mathcal{A} of feasible transfers from player 1 to 2 can either include or be included in a set of feasible transfers allowed in Wang (1995).

Players’ utility functions are written as $u_1(a, \theta_1) = \tilde{u}_1(\theta_1 - a)$ and $u_2(a, \theta_2) = \tilde{u}_2(\theta_2 + a)$, where each \tilde{u}_i is a strictly increasing and strictly concave function defined on an interval. By construction, both players consume the aggregate endowments each time, i.e., the date-by-date resource constraint is satisfied.

The third environment is a multiplicative preference shock model. Each player’s utility is $u_i(a, \theta_i) := \theta_i \tilde{u}_i(a)$, where \tilde{u}_i is strictly concave and $\Theta_i \subseteq \mathbb{R}_{++}$.¹² Each player has his (fixed) favorite action $a_i(\theta_i) = a_i \in \mathcal{A}$ with $a_1 \neq a_2$. The multiplicative preference shocks affect the (marginal) utility of a common action. Thus preference intensities are private information. Also,

⁹ One can extend the action space to a (compact convex) multi-dimensional space.

¹⁰ First, $a_i(\theta_i)$ may lie on the boundary of \mathcal{A} if $u_i(\cdot, \theta_i)$ is monotone. Second, we further assume that $a_1(\theta_1)$ and $a_2(\theta_2)$ are not identical with positive probability (i.e., $a_1(\theta_1) \neq a_2(\theta_2)$ for some (θ_1, θ_2)). As a result, the Pareto frontier of the feasible payoff set is not a singleton because the players cannot simultaneously achieve their maximum payoffs $\mathbb{E}_\theta [u_i(a_i(\theta_i), \theta_i)]$.

¹¹ Posing a restriction on possible transfers in an insurance model makes the problem technically rather challenging. It limits the degree to which each player can be punished in a single period. Also, an optimal transfer may lie on the boundary of feasible transfers.

¹² Whenever the utility function is written as $u_i(a, \theta_i) = r_i(\theta_i) \tilde{u}_i(a)$, where $r_i : \Theta_i \rightarrow \mathbb{R}_{++}$ and $\Theta_i \subseteq \mathbb{R}_{++}$, we can redefine it as $\theta_i k_i \tilde{u}_i(a)$, where the density function is $\tilde{f}(\theta_i) := (r_i(\theta_i) / (k_i \theta_i)) f_i(\theta_i)$ with $k_i := \sum_{\theta_i \in \Theta_i} (r_i(\theta_i) \cdot f_i(\theta_i) / \theta_i) (> 0)$.

if each player has a CARA utility in the insurance environment, then endowment shocks are isomorphic to multiplicative preference shocks.

It is worth pointing out that the collective decision environment has a distinct feature from the insurance partnership in that the players can possibly make a better decision by sharing more information. Although the players have an incentive to exaggerate their private information, they do not know how aligned their preferences are. Hence it is possible even in the static model to have an efficient allocation dependent on the players’ types in an incentive compatible way. On the other hand, in the static insurance model, it is impossible to have an efficient allocation depend on the players’ types in an incentive compatible manner.

Following Fudenberg et al. (1994), we restrict attention to Perfect Public Equilibria (PPE). That is, each player conditions only on the “public” histories of the game. Letting the initial history be h^0 , a public history at time $t \in \mathbb{N}$, h^t , is a sequence of (i) past announcements of the players (they make reports $(\hat{\theta}_i)_{i \in N} \in \Theta$ after they observe their preference shocks), (ii) past realized actions, and (iii) possibly realizations of a public randomization device. Let H^t be the set of all possible public histories h^t , and let $H^0 = \{h^0\}$.

Given the current reports and the past public history of the game, history-dependent allocations are determined according to an enforceable contract, which is a sequence of functions $a = (a^t)_{t \in \mathbb{N}}$, where $a^t : H^{t-1} \times \Theta \rightarrow \mathcal{A}$. The contract is enforceable in the sense that the players commit themselves to it a priori (at time “0”) before they learn their preference shocks.

A public strategy for player i is a sequence of functions $\hat{\theta}_i = (\hat{\theta}_i^t)_{t \in \mathbb{N}}$, where $\hat{\theta}_i^t : H^{t-1} \times \Theta_i \rightarrow \Theta_i$. Each strategy profile $\hat{\theta} = (\hat{\theta}_i)_{i \in N}$ induces a probability distribution over public histories. Letting $\delta \in (0, 1)$ be the common discount factor, player i ’s ex-ante expected discounted average payoff is given by

$$\mathbb{E} \left[(1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} u_i \left(a^t(h^{t-1}, \hat{\theta}^t(h^{t-1}, \theta^t)), \theta_i^t \right) \right].$$

3. Optimal mechanisms

We analyze the repeated game using the recursive methods developed by Abreu et al. (1990). Specifically, letting $\mathcal{W} \subseteq \mathbb{R}^2$ be the set of pure strategy PPE payoff profiles, we can decompose the players’ payoffs into (i) a profile of current utilities $(u_i(a(\theta), \theta_i))_{i \in N}$ and (ii) a profile of continuation values $(w_i(\theta))_{i \in N} \in \mathcal{W}$. In other words, any PPE payoff profile is written as $(\mathbb{E}_{\theta}[(1 - \delta)u_i(a(\theta), \theta_i) + \delta w_i(\theta)])_{i \in N}$ and it can be summarized by a current decision rule $a(\hat{\theta})$ and equilibrium continuation values $w(\hat{\theta}) \in \mathcal{W}$ as a function of a profile of announcements $\hat{\theta}$. A mechanism (a, w) is a pair of current allocation $a : \Theta \rightarrow \mathcal{A}$ and continuation value functions $w : \Theta \rightarrow \mathbb{R}^2$.

The set of PPE payoffs has the recursive structure. For any candidate equilibrium payoff value set $W \subseteq \mathbb{R}^2$, a mechanism (a, w) is *admissible* with respect to W if it satisfies the following (Bayesian) incentive constraints and self-enforceability condition:

$$\mathbb{E}_{\theta_{-i}} [(1 - \delta)u_i(a(\theta), \theta_i) + \delta w_i(\theta)] \geq \mathbb{E}_{\theta_{-i}} [(1 - \delta)u_i(a(\hat{\theta}_i, \theta_{-i}), \theta_i) + \delta w_i(\hat{\theta}_i, \theta_{-i})] \text{ for all } i \in N \text{ and } \theta_i, \hat{\theta}_i \in \Theta_i; \text{ and} \tag{1}$$

$$w(\theta) \in W \text{ for all } \theta \in \Theta. \tag{2}$$

Expression (1) requires each player to prefer to tell the truth given that the other player is expected to tell the truth. Expression (2) is the self-enforceability condition that continuation values

must be in the candidate equilibrium value set. Defining the mapping B from the collection of candidate equilibrium value sets (namely, the power set $\mathcal{P}(\mathbb{R}^2)$ of \mathbb{R}^2) into itself by the set of payoffs attainable by admissible mechanisms, the set of PPE payoffs is given by the largest bounded fixed point of the operator B , where we assume that the players can utilize some public randomization device.¹³

Formally, $B : \mathcal{P}(\mathbb{R}^2) \rightarrow \mathcal{P}(\mathbb{R}^2)$ maps any subset W of \mathbb{R}^2 to

$$B(W) := \text{co} \left\{ \left(\mathbb{E}_\theta [(1 - \delta)u_i(a(\theta), \theta_i) + \delta w_i(\theta)] \right)_{i \in N} \mid (a, w) \text{ is admissible w.r.t. } W \right\}.$$

With slight modifications to Abreu et al. (1990), it can be shown that the monotone mapping B preserves compactness and that the PPE value set \mathcal{W} is a compact convex set. Hence, we define \bar{v}_i (resp. \underline{v}_i) as player i 's maximum (resp. minimum) equilibrium payoff. Indeed, \bar{v}_i is player i 's expected payoff when his preferred action is always taken: $\bar{v}_i = \mathbb{E}_\theta [u_i(a_i(\theta_i), \theta_i)]$.¹⁴

To study the dynamics of efficient mechanisms, we formalize the Bellman equation that characterizes the frontier of PPE values. Denote each player's possible PPE value set by $\mathcal{W}_i := \{v_i \in \mathbb{R} \mid \text{there exists } v_{-i} \in \mathbb{R} \text{ such that } (v_i, v_{-i}) \in \mathcal{W}\} = [\underline{v}_i, \bar{v}_i]$. For a given expected utility $v \in \mathcal{W}_2$ promised to player 2, define $V(v)$ as the highest value delivered to player 1. The value function $V : \mathcal{W}_2 \rightarrow \mathcal{W}_1$ satisfies the following functional equation.

$$V(v) = \max_{a(\cdot) \in \mathcal{A}, w(\cdot)} \mathbb{E}_\theta [(1 - \delta)u_1(a(\theta), \theta_1) + \delta V(w(\theta))]$$

subject to

$$\mathbb{E}_\theta [(1 - \delta)u_2(a(\theta), \theta_2) + \delta w(\theta)] = v; \tag{PK}$$

$$\begin{aligned} &\mathbb{E}_{\theta_2} [(1 - \delta)u_1(a(\theta), \theta_1) + \delta V(w(\theta))] \\ &\geq \mathbb{E}_{\theta_2} [(1 - \delta)u_1(a(\hat{\theta}_1), \theta_1) + \delta V(w(\hat{\theta}_1, \theta_2))] \text{ for all } \theta_1, \hat{\theta}_1 \in \Theta_1; \end{aligned} \tag{IC1}$$

$$\begin{aligned} &\mathbb{E}_{\theta_1} [(1 - \delta)u_2(a(\theta), \theta_2) + \delta w(\theta)] \\ &\geq \mathbb{E}_{\theta_1} [(1 - \delta)u_2(a(\theta_1, \hat{\theta}_2), \theta_2) + \delta w(\theta_1, \hat{\theta}_2)] \text{ for all } \theta_2, \hat{\theta}_2 \in \Theta_2; \text{ and} \end{aligned} \tag{IC2}$$

$$w(\theta) \in \mathcal{W}_2 \text{ for all } \theta \in \Theta. \tag{SE}$$

The promise-keeping constraint (PK) requires that player 2's expected discounted utility delivered by a mechanism must be the promised value v . Expressions (IC i) are the players' (Bayesian) incentive compatibility constraints: each player i prefers to tell the truth, taking into account of the future consequences of his announcement, given that the other player is expected to be truthful. The self-enforceability condition (SE) requires the continuation values to be in the equilibrium value set.

A mechanism (a, w) is *incentive compatible* (IC) if it satisfies (IC i) for all $i \in N$. A mechanism (a, w) is *incentive feasible* at $v \in \mathcal{W}_2$ if it satisfies (PK), (IC i) for all $i \in N$, and (SE) when the promised value is v . We sometimes denote a mechanism (a, w) by, for example,

¹³ This assumption is introduced to convexify the set of PPE payoffs. While we implicitly introduce a public randomization device, we will explicitly formulate, in Appendix A, a stochastic mechanism as a collection of probability distributions on the current actions for each profile of announcements of preference shocks.

¹⁴ When player i 's preferred action is always taken, he cannot do any better and the opponent is ignored and hence has no incentive to lie. In other words, a dictatorship is incentive compatible. Also, we assume that, in a one-shot game, the players can achieve a (static) equilibrium payoff profile v that Pareto-dominates the payoff profile associated with some random dictatorial allocation (i.e., a randomization of $a_1(\cdot)$ and $a_2(\cdot)$). This ensures that the set of PPE values has a non-empty interior.

$(a(\cdot|v, \delta), w(\cdot|v, \delta))$ to make explicit the dependence of the mechanism on the promised value v or the discount factor δ .

While our general model accommodates the dynamic insurance model as a special case, the repeated collective decision problem has distinct features from a pure dynamic insurance model in terms of incentives. In the collective decision environment, while the players have an incentive to claim that their type is more extreme than it actually is, they do not always have an incentive to exaggerate their preferences all the way to the extremes. This is in contrast to the dynamic insurance environment, where the players have an incentive to claim to be poorer than they actually are and those facing the first-best contract have an incentive to misreport all the way to the lowest possible endowment to obtain the maximal amount of transfer from the opponent. Moreover, in the collective decision problem, the relevant direction in which the IC constraints bind depends on whether one player's favorite action is above or below the other player's average type. The relevant constraints for a player whose favorite action is above (resp. below) his opponent's average are those that ensure he does not want to lie upwards (resp. downwards).

Since player i 's best (and $-i$'s worst) Pareto efficient payoff is attained in PPE, the PPE value function passes through the payoff profile $(\bar{v}_i, \underline{v}_{-i}^P)$, where $\underline{v}_{-i}^P := \mathbb{E}_\theta[u_{-i}(a_i(\theta_i), \theta_{-i})]$, for each $i \in N$. Thus, the Pareto frontier of PPE values is always characterized by $(v, V(v))$ with $v \in \mathcal{W}_2^* := [\underline{v}_2^P, \bar{v}_2]$ (i.e., the frontier of PPE values is well-defined and downward-sloping on \mathcal{W}_2^*). Whenever $\underline{v}_2 < \underline{v}_2^P$, on the other hand, any frontier of PPE value on $[\underline{v}_2, \underline{v}_2^P]$ is upward-sloping. This can occur, for example, when the players can take extreme actions in the quadratic preference collective decision problem (i.e., \mathcal{A} is "large" relative to Θ_i).

3.1. Renegotiation-proof PPE

We have so far characterized the frontier of the full-commitment PPE values and we have seen that some PPE values may be inefficient. In the collective decision environment, for example, the players may be able to take an efficient joint action today by using an inefficient continuation contract (which may be associated with taking extreme actions from tomorrow on). An efficient full-commitment contract may be susceptible to renegotiation.

Hence, we study renegotiation-proof (RP, for short) PPE. In infinite horizon games like ours, however, there is no clear-cut definition of RP equilibrium, in contrast to the finite horizon games where renegotiation-proofness is well defined by backward induction. Among various concepts of renegotiation-proofness in infinite horizon games, we focus on (*internally*) renegotiation-proof PPE, following the concept of "internally renegotiation-proof" equilibrium sets by Ray (1994).¹⁵ The idea that a mechanism is internally renegotiation-proof is captured by the condition that it is (weakly) Pareto efficient such that its continuation contract is renegotiation-proof. It extends to the backward-induction definition of renegotiation-proofness in finite horizon games.

The key observation is that the payoff vector $(\bar{v}_i, \underline{v}_{-i}^P)$, the one associated with player i 's best and $-i$'s worst Pareto efficient feasible payoff, is supported by an equilibrium of the one-shot game. Any PPE payoff vector v (where $v_i < \underline{v}_i^P$ for some $i \in N$) is Pareto-dominated. With this in mind, consider the frontier of the largest self-generating set W (i.e., $W \subseteq B(W)$) included in $[\underline{v}_2^P, \bar{v}_2] \times [\underline{v}_1^P, \bar{v}_1]$. The frontier of W is downward sloping and defined entirely on $[\underline{v}_2^P, \bar{v}_2]$.

¹⁵ Zhao (2006) formulates internal renegotiation-proofness in terms of a value function in a repeated principal-agent problem. Bergin and MacLeod (1993) introduce a similar concept, "weak full recursive efficiency," and compare various notions of renegotiation-proofness. In our analysis, as is standard, renegotiation can occur at the beginning of any given period.

Moreover, the set W can be obtained as the largest fixed point of an auxiliary operator $B^* : \mathcal{P}(\mathbb{R}^2) \rightarrow \mathcal{P}(\mathbb{R}^2)$ that maps any $W \in \mathcal{P}(\mathbb{R}^2)$ to¹⁶

$$B^*(W) := B(W \cap ([\underline{v}_2^P, \bar{v}_2] \times [\underline{v}_1^P, \bar{v}_1])) \cap ([\underline{v}_2^P, \bar{v}_2] \times [\underline{v}_1^P, \bar{v}_1]).$$

We now characterize the frontier of W by the graph of the value function V . The graph of V has the property that the Pareto-frontier of the generated value function coincides with the original value function V . Thus, the RP-PPE value function $V : \mathcal{W}_2^* \rightarrow \mathcal{W}_1^*$ (where $\mathcal{W}_1^* := [\underline{v}_1^P, \bar{v}_1]$) satisfies the following Bellman equation. For each $v \in \mathcal{W}_2^*$,

$$V(v) = \max_{(a,w)(\cdot) \in \mathcal{A} \times \mathcal{W}_2^*} \mathbb{E}_\theta [(1 - \delta)u_1(a(\theta), \theta_1) + \delta V(w(\theta))] \tag{OPT}$$

subject to (PK) and (IC $_i$) for each $i \in N$.

We abuse notation to denote V for (full-commitment) PPE and RP-PPE value functions. Since V is strictly decreasing on \mathcal{W}_2^* , the promise-keeping constraint in the Bellman equation can be replaced with the weak inequality.

In terms of the Bellman equation (OPT), renegotiation-proofness imposes a lower limit on player 2’s continuation values as in the limited-commitment problem. As stressed by Zhao (2006), however, the difference between the renegotiation-proof and limited-commitment problems is that while an agent can unilaterally walk away from an ongoing contract in limited-commitment environments, renegotiation-proofness requires that both agents mutually agree on any change in an ongoing contract. Hence, whenever $\underline{v}_2 = \underline{v}_2^P$ as in the insurance model, no restriction is imposed.

The Bellman equations characterizing RP (as well as full-commitment) PPE value functions have a special feature that the value function itself enters into the player 1’s incentive constraint, as in Wang (1995) and Zhao (2007). Thus, for any two given candidate value functions, the sets of incentive compatible mechanisms under these value functions could be distinct. This implies that the one-step operator (defined by the right-hand side of the Bellman equation) may not necessarily be monotone, and hence the standard technique (i.e., Blackwell condition (see, for example, Aliprantis and Border (2006, Theorem 3.53))) to verify that the one-step operator is a contraction mapping cannot readily be invoked. Despite this, we have established that the RP-PPE value function is well defined. We also have an algorithm to compute the RP-PPE value function of our model by iterating the auxiliary operator B^* .¹⁷

Moreover, by using this observation, we show in Appendix A that the largest fixed point $W = W(\delta)$ of B^* is monotonically non-decreasing in the discount factor δ . Thus, the RP (as well as the full-commitment) PPE value functions are monotonically non-decreasing in the discount factor.

Note also that internal renegotiation-proofness allows a current action to be ex post inefficient. While an optimal mechanism is required to be (weakly) Pareto efficient subject to the continuation mechanism being “internally renegotiation-proof,” ex post inefficiency may still occur due to the cost of keeping incentive constraints. For concreteness, consider the quadratic collective decision problem in which the ex-ante symmetric players with the lowest type being -1 . Internal renegotiation-proofness is weaker than the requirement that the current action be between players’ reported types (i.e., $a(\theta_1, \theta_2) \in [\min(\theta_1, \theta_2), \max(\theta_1, \theta_2)]$). Indeed, the second-best current

¹⁶ We can characterize W in slightly different yet equivalent ways. For example, it is the largest set satisfying $W = B(W) \cap ([\underline{v}_2^P, \bar{v}_2] \times [\underline{v}_1^P, \bar{v}_1])$.

¹⁷ In contrast, an internally renegotiation-proof equilibrium does not necessarily exist. See Ray (1994) and Zhao (2006).

action $a(-1, -1)$ can be lower than -1 . While making $a(-1, -1)$ less than -1 is costly to both players, it relaxes the incentive constraint of the second lowest type imitating the lowest type. If player i 's true type is the second lowest, then he is less inclined to exaggerate his announcement because his announcement of the lowest type induces the joint action further away from his true type, which occurs with the probability that his opponent's shock is the lowest. Thus, whenever the second lowest type has an incentive to mimic the lowest type, the ex-ante loss from such "money burning" turns out to be second order compared to the first-order gain from relaxing the incentive costs (see Online Appendix B for more details).

3.2. Approximate efficiency

The first-best allocation maximizes a weighted sum of players' instantaneous utilities. Formally, the first-best value function on \mathcal{W}_2^* , $V^{FB} : \mathcal{W}_2^* \rightarrow \mathcal{W}_1^*$, is characterized by the highest feasible payoff delivered to player 1 given that player 2' expected value is at least $v \in \mathcal{W}_2^*$:

$$V^{FB}(v) = \max_{a(\cdot) \in \mathcal{A}} \mathbb{E}_\theta [u_1(a(\theta), \theta_1)] \text{ subject to } \mathbb{E}_\theta [u_2(a(\theta), \theta_2)] \geq v. \tag{FB}$$

Denote by $a^{FB}(\cdot|v) : \Theta \rightarrow \mathcal{A}$ the first-best allocation, i.e., the solution to the first-best Pareto problem (FB).

In the quadratic collective decision model, the first-best action is a weighted average of the players' types: $a^{FB}(\theta|v) = \sqrt{\frac{v}{v_2^P}}\theta_1 + \left(1 - \sqrt{\frac{v}{v_2^P}}\right)\theta_2$, where $v \in [v_2^P, 0]$ and $v_2^P = -\mathbb{E}_\theta[(\theta_1 - \theta_2)^2]$. The difficulty of implementing the first-best decision rule is that, whenever a player's preference shock is different from his opponent's average type, he would have an incentive to exaggerate his report towards the extremes. If player 2's average shock is zero, then player 1 with type θ_1 would announce $\sqrt{\frac{v_2^P}{v}}\theta_1$ (except at $v = \bar{v}_2$, in which case player 1's incentive is trivially satisfied).

When $\delta > 0$, continuation values can be used as an additional instrument to get the players to report their types truthfully. The key is to present the players with a trade-off between the benefit of a larger influence on the current decision and the loss they will incur in future continuation values. This allows the mechanism to take into account the intensity of the players' preferences, which, in turn, leads to efficiency gains when compared to a static problem.

Continuation values play a similar role to the one side-payments play in standard static incentive problems. The difference between side-payments and continuation values is that the latter can only imperfectly transfer utilities across the players. Continuation values must be drawn from the equilibrium value set. To transfer continuation utility from one player to another in any period, the decisions for subsequent periods must be altered in an incentive compatible way. Together with the lack of observability, generally the players cannot attain *exact* efficiency as an equilibrium outcome. Indeed, exact efficiency would call for an equilibrium in which for all histories future decisions would not respond to current announcements. Hence, truth-telling would have to be a static best response for the players, and this cannot be attained in an incentive compatible manner except for rather trivial cases in which the first-best allocation is achieved in a one-shot game.¹⁸

¹⁸ Consider a quadratic collective decision problem among two ex-ante symmetric players with two preference shocks. A less trivial case is where they have three shocks and where the middle shock is the average type. Then, the middle type does not have any incentive to lie upward or downward.

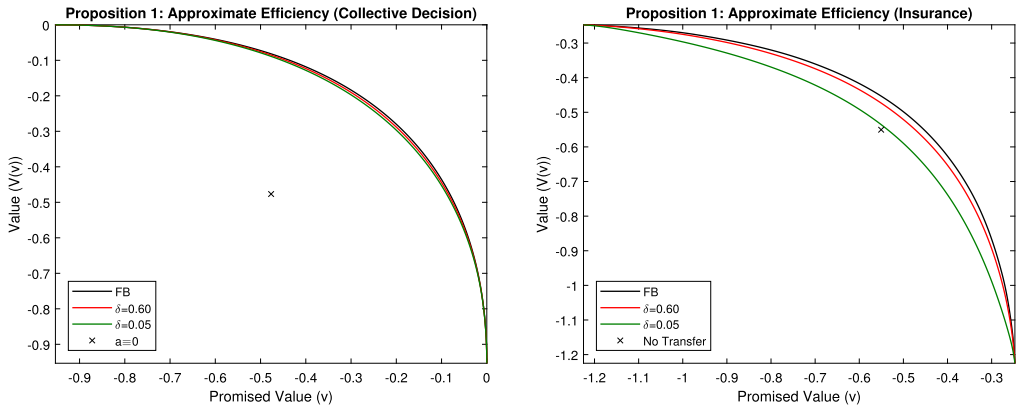


Fig. 1. The Value Functions V_{δ}^{SB} and V^{FB} in the Quadratic Collective Decision Problem (Left) and in the CARA Insurance Problem (Right).

Although full efficiency cannot generally be attained, one can arbitrarily approximate it as the players become patient.

Proposition 1 (Approximate Efficiency). *For any $\varepsilon > 0$, there exists $\bar{\delta} \in (0, 1)$ such that, if $\delta \in (\bar{\delta}, 1)$ then the RP-PPE value $V_{\delta}^{SB}(v) := V(v)$ is within ε of the first-best value $V^{FB}(v)$ for all $v \in \mathcal{W}_2^*$. Put differently, V_{δ}^{SB} converges uniformly to V^{FB} as δ approaches 1. Moreover, the convergence of V_{δ}^{SB} is monotone in the discount factor δ .*

In Fig. 1, we plot the RP-PPE value functions V_{δ}^{SB} for different values of δ and compare them with the first-best value function V^{FB} . We do this for the quadratic collective decision problem on the left and for the CARA insurance problem on the right.¹⁹ Also, the “x” in the left panel denotes the payoff profile associated with a constant action $a \equiv 0$ (a weighted sum of ex-ante average types $\mathbb{E}_{\theta}[\gamma\theta_1 + (1 - \gamma)\theta_2] = 0$).²⁰ The “x” in the right panel denotes the symmetric static second-best values (i.e., the values with no transfer).

In both cases, as stated in Proposition 1, V_{δ}^{SB} uniformly converges to V^{FB} as δ approaches 1. If we were to zoom in we could also see that they actually never quite attain the first-best values for interior v (see also Online Appendix B, which reports the social welfare loss in %). In addition, it is worth noting that the behavior of both problems is quite different for low values of δ . In particular, the efficiency losses due to the private information are much larger in the insurance problem than in the collective decision problem. In the insurance environment, when $\delta = 0$, there is no way the players can mutually insure themselves. In the collective decision

¹⁹ For the quadratic collective decision problem, we set $\Theta_1 = \Theta_2 = \{-1.1 + 0.2k \mid k \in \{0, 1, \dots, 11\}\}$ with the uniform density. The action space is $\mathcal{A} = [-1.3, 1.3]$. For the CARA insurance problem, we let $u_i(a, \theta_i) = -\exp(-(\theta_i + (-1)^i a))$. We set $\Theta_i = \{0, 0.8, 1.6\}$ with the uniform density. The set of feasible transfers is $\mathcal{A} = [-0.8, 0.8]$. For each problem we discretize the set $\mathcal{W}_2^* = [v_2^P, \bar{v}_2]$ to compute the value functions. See Online Appendix B for more details.

²⁰ The “x” in the left panel of Fig. 1 (the payoff profile $(\mathbb{E}_{\theta_2}[-(\theta_2)^2], \mathbb{E}_{\theta_1}[-(\theta_1)^2])$) is exactly the midpoint of (v_2^P, \bar{v}_1) and (\bar{v}_2, v_1^P) . This is generally the case in any symmetric quadratic preference model with $\mathbb{E}_{\theta_i}[\theta_i] = 0$. Also, if the players, who would face the first-best allocation, would always be able to exactly announce $\frac{\theta_i}{\gamma_i}$ and if the resulting action $\theta_1 + \theta_2$ would be feasible, then the resulting (hypothetical) expected utility profile would be $(\mathbb{E}_{\theta_2}[-(\theta_2)^2], \mathbb{E}_{\theta_1}[-(\theta_1)^2])$.

problem, however, even in the static setting $\delta = 0$ the players can credibly convey some of their private information since they are not in a zero-sum game as is the case in the static insurance problem. It is also worth noting that the approximate efficiency of the RP-PPE value function immediately implies that of the (full-commitment) PPE value function (on \mathcal{W}_2^*).

The intuition behind the approximate efficiency result is as follows. As the discount factor δ approaches unity, the utility value of the current period, which is weighted by $(1 - \delta)$, becomes insignificant relative to the continuation values. Hence, in order to guarantee truth-telling in the current period, continuation values have to vary only minimally, and the associated losses from variation in continuation values become negligible. It is worth noting that the trade-off between current allocation and continuation values is indeed less steep for high values of δ . When δ is close to unity, it is easier to provide incentives and there will be less variability in future allocations.

Similarly to Athey and Bagwell (2001), we construct continuation values that replicate the expected payments of the expected externality mechanism proposed by Arrow (1979) and d'Aspremont and Gérard-Varet (1979), that guarantee efficiency in standard static Bayesian mechanism design problems with transfers.²¹ This approach differs from Fudenberg et al. (1994).²²

In the context of collective decision problems, our approximate efficiency result can be contrasted with Jackson and Sonnenschein (2007). In their linking mechanism, the players are allowed to report each possible type a fixed number of times according to the frequency with which that type should statistically be realized. Jackson and Sonnenschein (2007, Corollary 2) demonstrate that, for any $\varepsilon > 0$, their linking mechanism is less than ε inefficient relative to the first best if players are patient and face sufficiently many identical problems. The sources of the efficiency losses in their mechanism and in our scheme are thus quite different. In the linking mechanism, when the last periods get closer, players may not be able to report their types truthfully, as they might have run out of their budgeted reports for a particular type. Then, the linking mechanism forces them to lie. Instead, in our setting, the inefficiency arises from variation in continuation values over time.

3.3. Dynamics of optimal mechanisms

In this section, we find the optimal mechanism by solving the second-best problem (OPT) for each promised value, and establish the long-run properties of the optimal mechanism. As we have to follow a lot of technical steps, it is useful to first preview what we do. We start by taking the *first-order approach* to solving the second-best problem (OPT), i.e., we consider a relaxed problem where each player's relevant incentive constraints are local ones.²³ Thus, we first represent the local incentive constraints (ICi-UP) and (ICi-DW). Next, we form the Lagrangian (LAG*) associated with the relaxed problem, and formalize the optimality conditions.

²¹ Athey and Bagwell (2001) analyze an infinitely repeated Bertrand duopoly game, where each firm privately receives discrete cost shocks. In contrast to our approximate efficiency result, they establish that, for a discount factor strictly less than one, monopoly profits can be *exactly* attained by firms making use of asymmetric continuation values. The difference from our setting is that, for some states that occur with positive probability, firms in their model can transfer profits without any costs.

²² Furthermore, their conditions for the Folk theorem are not satisfied in our model.

²³ We do this only for ease of exposition. Appendix A establishes our main result regardless of the first-order approach. It also provides a condition under which the first-order approach is valid for the following standard environments: quadratic, CARA, and multiplicative preferences.

The Lagrangian turns out to be the weighted sum of the players’ virtual utilities (VU) that incorporate the costs of incentives, and the weights attached on the players vary over time in order to keep incentives (Lemma 2). Moreover, the players’ relative Pareto weights (i.e., a measure of “inequality”) increase over time in expectation (Lemma 1). These observations lead us to our main theorem (Theorem 1): the optimal mechanism converges to the adoption of one player’s preferred allocation. Even though we can have “regime shifts” where the identity of the “despot” might switch over time, on average the dynamics of inequality measured by the share of Pareto weights accrued to the better-off agent converges to one (Corollary 1).

To formulate the first-order approach, we denote by $U_i(\theta_i, \hat{\theta}_i|a, w)$ player i ’s expected utility of a mechanism (a, w) when his type is θ_i and he announces $\hat{\theta}_i$:

$$U_i(\theta_i, \hat{\theta}_i|a, w) := \begin{cases} \mathbb{E}_{\theta_{-1}} \left[(1 - \delta)u_1(a(\hat{\theta}_1, \theta_2), \theta_1) + \delta V(w(\hat{\theta}_1, \theta_2)) \right] & \text{if } i = 1 \\ \mathbb{E}_{\theta_{-2}} \left[(1 - \delta)u_2(a(\theta_1, \hat{\theta}_2), \theta_2) + \delta w(\theta_1, \hat{\theta}_2) \right] & \text{if } i = 2 \end{cases}$$

We denote by $U_i(\theta_i|a, w) = U_i(\theta_i, \theta_i|a, w)$ or simply by $U_i(\theta_i)$ whenever it is clear from the context. Second, we assume that each $\Theta_i = \{\theta_i^{(k_i)}\}_{k_i=1}^{m_i}$ is in \mathbb{R} such that $\theta_i^{(1)} < \theta_i^{(2)} < \dots < \theta_i^{(m_i)}$.²⁴ Now, each player’s local upward and downward incentive constraints (ICi-UP) and (ICi-DW) are written as follows.

$$U_i \left(\theta_i^{(k_i)} \middle| a, w \right) \geq U_i \left(\theta_i^{(k_i)}, \theta_i^{(k_i+1)} \middle| a, w \right) \text{ for all } k_i \in \{1, \dots, m_i - 1\}; \text{ and} \quad (\text{ICi-UP})$$

$$U_i \left(\theta_i^{(k_i)} \middle| a, w \right) \geq U_i \left(\theta_i^{(k_i)}, \theta_i^{(k_i-1)} \middle| a, w \right) \text{ for all } k_i \in \{2, \dots, m_i\}. \quad (\text{ICi-DW})$$

We find the optimal mechanism by solving the following relaxed second-best problem (OPT*) for each promised value $v \in \mathcal{W}_2^*$:

$$V(v) = \max_{\substack{a(\cdot) \in \mathcal{A} \\ w(\cdot) \in \mathcal{W}_2^*}} \mathbb{E}_\theta [(1 - \delta)u_1(a(\theta), \theta_1) + \delta V(w(\theta))] \quad (\text{OPT}^*)$$

subject to (PK) and (ICi-UP) and (ICi-DW) for each $i \in N$.

The optimal dynamic mechanism design problem is reduced to maximizing the Lagrangian of (OPT*) at each promised value $v \in \mathcal{W}_2^*$. Define the Lagrange multipliers $\lambda = (\lambda_1, \lambda_2) = ((\lambda_{i,k_i}^{\text{UP}})_{k_i=1}^{m_i-1}, (\lambda_{i,k_i}^{\text{DW}})_{k_i=2}^{m_i})_{i \in N}$ and $\gamma = (\gamma_1, \gamma_2)$ as follows. Let $(\lambda_{i,k_i}^{\text{UP}})_{k_i}$ and $(\lambda_{i,k_i}^{\text{DW}})_{k_i}$ be those on the player i ’s local upward and downward incentive compatibility constraints (ICi-UP) and (ICi-DW), respectively. The multiplier γ_2 is the one on the player 2’s promise-keeping constraint (PK), and γ_1 is that on the objective function (player 1’s utility function). Since $(\gamma_1, \gamma_2) \neq 0$, we normalize it by $\gamma_1 + \gamma_2 = 1$ (alternatively, we could normalize $\gamma_1 = 1$ whenever $\gamma_1 > 0$). The multipliers γ_1 and γ_2 play the role of Pareto weights.

Formally, there exists a vector of non-negative Lagrange multipliers (λ, γ) with $\gamma_1 + \gamma_2 = 1$ such that the following two conditions characterize an optimal mechanism. First, the Lagrangian $\mathcal{L}(\cdot, \cdot|\lambda, \gamma, v)$ is maximized at an optimal mechanism (a, w) , i.e., $\mathcal{L}(a, w|\lambda, \gamma, v) \geq \mathcal{L}(\tilde{a}, \tilde{w}|\lambda, \gamma, v)$ for all functions $(\tilde{a}, \tilde{w}) : \Theta \rightarrow \mathcal{A} \times \mathcal{W}_2^*$. Second, the complementary slackness conditions are satisfied.

Following Myerson (1981, 1984, 1985), we view the Lagrangian as the sum of the players’ virtual utilities. To that end, we introduce the dummy multipliers $\lambda_{i,0}^{\text{UP}} = \lambda_{i,m_i}^{\text{UP}} = \lambda_{i,1}^{\text{DW}} = \lambda_{i,m_i+1}^{\text{DW}} = 0$

²⁴ This single-dimensional assumption can be dropped when we take care of the entire IC constraints. See Appendix A.

and dummy types $\theta_i^{(0)}$ and $\theta_i^{(m_i+1)}$. Let $\gamma_i \tilde{u}_i(a, \theta_i)$ be player i 's instantaneous virtual utility at θ_i with respect to γ_i and λ_i :

$$\begin{aligned} \gamma_i \tilde{u}_i(a, \theta_i^{(k_i)}) &:= \gamma_i \left\{ u_i(a, \theta_i^{(k_i)}) \left(1 + \frac{\tilde{\lambda}_{i,k_i}^{UP} - \tilde{\lambda}_{i,k_i-1}^{UP} + \tilde{\lambda}_{i,k_i}^{DW} - \tilde{\lambda}_{i,k_i+1}^{DW}}{f_i(\theta_i^{(k_i)})} \right) \right. \\ &+ \left\{ u_i(a, \theta_i^{(k_i)}) - u_i(a, \theta_i^{(k_i-1)}) \right\} \frac{\tilde{\lambda}_{i,k_i-1}^{UP}}{f_i(\theta_i^{(k_i)})} \\ &\left. + \left\{ u_i(a, \theta_i^{(k_i)}) - u_i(a, \theta_i^{(k_i+1)}) \right\} \frac{\tilde{\lambda}_{i,k_i+1}^{DW}}{f_i(\theta_i^{(k_i)})} \right\}, \end{aligned} \tag{VU}$$

where we re-define the Lagrange multipliers $\tilde{\lambda}_{i,k_i}^{UP} = \frac{\lambda_{i,k_i}^{UP}}{\gamma_i}$ and $\tilde{\lambda}_{i,k_i}^{DW} = \frac{\lambda_{i,k_i}^{DW}}{\gamma_i}$ whenever $\gamma_i > 0$.

The first term of $\tilde{u}_i(a, \theta_i^{(k_i)})$ is the utility $u_i(a, \theta_i^{(k_i)})$ adjusted by the cost of incentives associated with type $\theta_i^{(k_i)}$ lying upward $\theta_i^{(k_i+1)}$ and downward $\theta_i^{(k_i-1)}$ and types $\theta_i^{(k_i+1)}$ and $\theta_i^{(k_i-1)}$ imitating $\theta_i^{(k_i)}$. The second term is the upward information rent due to type $\theta_i^{(k_i-1)}$ imitating $\theta_i^{(k_i)}$. The third term is the downward information rent due to type $\theta_i^{(k_i+1)}$ imitating $\theta_i^{(k_i)}$.²⁵

Now, the Lagrangian amounts to the expected sum of the players' virtual utilities:

$$\begin{aligned} \mathcal{L}(a, w|\lambda, \gamma, v) &= \mathbb{E}_\theta \left[\left\{ \sum_{i=1}^2 (1 - \delta) \underbrace{\gamma_i \tilde{u}_i(a(\theta_i^{(k_i)}), \theta_{-i}, \theta_i^{(k_i)})}_{\text{player } i \text{'s virtual instantaneous utility}} \right. \right. \\ &+ \underbrace{\delta \gamma_1 V(w(\theta_1^{(k_1)}, \theta_2^{(k_2)})) \left(1 + \frac{\tilde{\lambda}_{1,k_1}^{UP} - \tilde{\lambda}_{1,k_1-1}^{UP} + \tilde{\lambda}_{1,k_1}^{DW} - \tilde{\lambda}_{1,k_1+1}^{DW}}{f_1(\theta_1^{(k_1)})} \right)}_{\text{player 1's virtual continuation utility}} \\ &\left. \left. + \underbrace{\delta \gamma_2 w(\theta_1^{(k_1)}, \theta_2^{(k_2)}) \left(1 + \frac{\tilde{\lambda}_{2,k_2}^{UP} - \tilde{\lambda}_{2,k_2-1}^{UP} + \tilde{\lambda}_{2,k_2}^{DW} - \tilde{\lambda}_{2,k_2+1}^{DW}}{f_2(\theta_2^{(k_2)})} \right)}_{\text{player 2's virtual continuation utility}} \right\} - \gamma_2 v. \end{aligned} \tag{LAG*}$$

The optimal current allocation maximizes the weighted sum of the players' *virtual* instantaneous utilities \tilde{u}_i with the weight given to player i being γ_i .²⁶ The optimal continuation value function $w(\cdot)$, on the other hand, maximizes the weighted sum of players' virtual continuation utilities with the weight given to player i being γ_i . Again, player i 's continuation utility is adjusted by the costs of incentives associated with type $\theta_i^{(k_i)}$ imitating $\theta_i^{(k_i+1)}$ and $\theta_i^{(k_i-1)}$ and types

²⁵ If we do not take the first-order approach, then we incorporate the costs of incentives associated with type $\theta_i^{(k_i)}$ imitating all the other types and all the other types imitating $\theta_i^{(k_i)}$. See Appendix A.

²⁶ Precisely, the player i 's virtual utility itself (the left-hand side of (VU)) depends on his weight γ_i . That is, the optimal current allocation maximizes the sum of the players' instantaneous virtual utilities $\gamma_i \tilde{u}_i$.

$\theta_i^{(k_i+1)}$ and $\theta_i^{(k_i-1)}$ imitating $\theta_i^{(k_i)}$.²⁷ We remark that while we can solve the optimal current allocation and the optimal continuation value function separately, the Pareto weights attached on the players’ virtual instantaneous utilities and virtual continuation utilities are the same: γ_1 and γ_2 .

Hence, the dynamics of repeated decision making are fully determined by (i) a decision rule that, at each period, maximizes the weighted sum of the players’ instantaneous *virtual* utilities, and (ii) the process that governs the evolution of the weights attached on the players’ virtual utilities. While the Pareto weights (Lagrange multipliers) of the first-best problem are time-invariant, the Pareto weights in the second-best problem (the ratio of which corresponds to the “slope” of the value function) have to vary in order to keep incentives.

Thus, we consider the stochastic process of players’ relative Pareto weights. With any promised value $v \in \mathcal{W}_2^*$, we can associate players’ Pareto weights $(\gamma_1, \gamma_2) = (\gamma_1(v), \gamma_2(v))$ with $\gamma_1 + \gamma_2 = 1$. Let $(\bar{\gamma}_1, \underline{\gamma}_2)$ and $(\underline{\gamma}_1, \bar{\gamma}_2)$ be the Pareto weights associated with $(\bar{v}_1, \underline{v}_2^P)$ and $(\underline{v}_1^P, \bar{v}_2)$, respectively. Then, for each $\theta \in \Theta$, the next-period’s weights (the weights associated with $w(\theta|v)$) are denoted by $(\gamma'_1(\theta), \gamma'_2(\theta))$. Now, we can find each player’s next-period relative Pareto weight $\frac{\gamma'_i(\theta)}{\gamma'_{-i}(\theta)}$ explicitly as a function of the relative Pareto weight $\frac{\gamma_i}{\gamma_{-i}}$ and the Lagrange multipliers by the first-order conditions of the Lagrangian (LAG*). When the continuation values are in the interior, it turns out that

$$\underbrace{\frac{\gamma'_{-i}(\theta_1^{(k_1)}, \theta_2^{(k_2)})}{\gamma'_i(\theta_1^{(k_1)}, \theta_2^{(k_2)})}}_{\text{player } -i \text{'s next-period relative Pareto weight at } (\theta_1^{(k_1)}, \theta_2^{(k_2)})} = \underbrace{\frac{\gamma_{-i}}{\gamma_i}}_{\text{player } -i \text{'s relative Pareto weight}} \cdot \underbrace{\frac{1 + \frac{\tilde{\lambda}_{-i,k_{-i}}^{\text{UP}} - \tilde{\lambda}_{-i,k_{-i}-1}^{\text{UP}} + \tilde{\lambda}_{-i,k_{-i}}^{\text{DW}} - \tilde{\lambda}_{-i,k_{-i}+1}^{\text{DW}}}{f_{-i}(\theta_{-i}^{(k_{-i})})}}{1 + \frac{\tilde{\lambda}_{i,k_i}^{\text{UP}} - \tilde{\lambda}_{i,k_i-1}^{\text{UP}} + \tilde{\lambda}_{i,k_i}^{\text{DW}} - \tilde{\lambda}_{i,k_i+1}^{\text{DW}}}{f_i(\theta_i^{(k_i)})}}}_{\text{ratio of coefficients put on virtual continuation utilities}}. \tag{3}$$

The numerator of the right-hand side of Expression (3) consists of player $-i$ ’s current Pareto weight and the coefficient put on his virtual continuation utility. Likewise, the denominator consists of player i ’s current Pareto weight and the coefficient put on his virtual continuation utility. By taking expectation of relative Pareto weight given by Expression (3), the sum of the Lagrange multipliers associated with type $\theta_i^{(k_i)}$ imitating all the other types and the ones θ_i associated with all the other types imitating $\theta_i^{(k_i)}$ cancel out with each other, and we obtain the “submartingale” property, irrespective of the first-order approach. Each player’s relative Pareto weight $\frac{\gamma_i}{\gamma_{-i}}$ always increases in expectation over time.

Lemma 1 (Submartingale). *Let $v \in (\underline{v}_2^P, \bar{v}_2)$ be such that $w(\cdot|v) \in (\underline{v}_2^P, \bar{v}_2)$. Each player i ’s relative Pareto weight satisfies*

$$\frac{\gamma_i}{\gamma_{-i}} = \mathbb{E}_{\theta_i} \left[\frac{1}{\mathbb{E}_{\theta_{-i}} \left[\frac{\gamma'_{-i}(\theta)}{\gamma'_i(\theta)} \right]} \right] < \mathbb{E}_{\theta} \left[\frac{\gamma'_i(\theta)}{\gamma'_{-i}(\theta)} \right].$$

Lemma 1 contrasts with Thomas and Worrall (1990), which shows that the marginal value (i.e., the negative of relative Pareto weight) follows a martingale. The difference hinges on the

²⁷ If we do not take the first-order approach, each player’s continuation utility is adjusted by the costs of incentives associated with type $\theta_i^{(k_i)}$ imitating all the other types and all the other types imitating $\theta_i^{(k_i)}$.

fact that both players have private information in our model. In fact, if player 1 (= $-i$) had no preference shock, then his marginal value with respect to player 2's preference shocks would follow a “martingale” process with respect to player 2's type distribution. In the above expression, the strict inequality becomes the equality as $\mathbb{E}_{\theta_{-i}}[\cdot]$ is degenerate.

Yet, the intuition behind Lemma 1 is similar to Thomas and Worrall (1990). In the optimal dynamic mechanism it will always be efficient to have continuation values vary over time for any given promised value $v \in (\underline{v}_2^P, \bar{v}_2)$. Continuation values allow the players with a more extreme type in the current period (poorer agents in the insurance models) to get more weight in the current allocation choice (higher current consumption) in exchange for forgoing future decision rights (future consumption).

Lemma 2 (Spreading of Values). *Let $v \in (\underline{v}_2^P, \bar{v}_2)$. For each $i \in N$, there is positive probability of $\frac{\gamma'_i}{\gamma_{-i}} \neq \frac{\gamma_i}{\gamma_{-i}}$. If $w(\cdot|v) \in (\underline{v}_2^P, \bar{v}_2)$, then for each $i \in N$, there is positive probability of both $\frac{\gamma'_i}{\gamma_{-i}} > \frac{\gamma_i}{\gamma_{-i}}$ and $\frac{\gamma'_i}{\gamma_{-i}} < \frac{\gamma_i}{\gamma_{-i}}$.*

While continuation values have to vary over time, our main result (Theorem 1) asserts that the decision process becomes dictatorial in the long run. To see the long-run dictatorial property, let $v^0 \in \mathcal{W}_2^*$ be an initial promised value, and let $(v^t)_{t \in \mathbb{N}}$ be a stochastic process of player 2's promised values induced by an optimal mechanism and a path of shocks $(\theta^t)_t$. Let (γ^t_1, γ^t_2) be players' Pareto weights associated with v^t for each $t \geq 0$.

Let W be the countable set of promised values that can be reached in a finite time from v^0 by the optimal mechanism. We take v^0 to be in the interior $(\underline{v}_2^P, \bar{v}_2)$, because if $v^0 \in \{\underline{v}_2^P, \bar{v}_2\}$ then $v^t = v^0$ for all $t \in \mathbb{N}$. Consider the two cases: (i) $w(\theta|v) \in \{\underline{v}_2^P, \bar{v}_2\}$ for some $(\theta, v) \in \Theta \times W$; and (ii) $w(\theta|v) \in (\underline{v}_2^P, \bar{v}_2)$ for all $(\theta, v) \in \Theta \times W$.

Case (i): Let t be the first time such that $w(\theta|v^{t-1}) \in \{\underline{v}_2^P, \bar{v}_2\}$ for some $\theta \in \Theta$ and $v^{t-1} \in W \cap (\underline{v}_2^P, \bar{v}_2)$. For all $\tau \geq t$, the promised values $v^\tau = v^t$ get absorbed into a point in $\{\underline{v}_2^P, \bar{v}_2\}$. At time t , the set of promised values that can be reached from v^0 is finite. Thus, in Case (i), with positive probability, a promised value can reach exactly \underline{v}_2^P or \bar{v}_2 in a finite time. For any other promised value v that can be reached from v^0 at time t (there are only finitely many such promised values), we can consider the two cases with the new initial promised value being v . Hence, from now on, we study the case in which an initial promised value v^0 satisfies $w(\theta|v) \in (\underline{v}_2^P, \bar{v}_2)$ for all $(\theta, v) \in \Theta \times W$.

Case (ii): Consider the case in which a promised value can never reach exactly \underline{v}_2^P or \bar{v}_2 in a finite time. With this assumption, the stochastic process $(\frac{\gamma^t_2}{\gamma^t_1})_t$ follows a sub-martingale process by Lemma 1, and we consider next the dynamics of the relative Pareto weights in the limit as $t \rightarrow \infty$.

Theorem 1 (Dictatorship in the Limit). *Let $v^0 \in (\underline{v}_2^P, \bar{v}_2)$ belong to Case (ii).*

1. *Suppose that the marginal value function is bounded as in the insurance problem: $\sup_{v \in [\underline{v}_2^P, \bar{v}_2]} |V'(v)| < \infty$. Almost surely along any path, $\frac{\gamma^t_2}{\gamma^t_1} \rightarrow \frac{1-\bar{\gamma}_1}{\bar{\gamma}_1} \in [0, \infty)$ or $\frac{\gamma^t_2}{\gamma^t_1} \rightarrow \frac{\bar{\gamma}_2}{1-\bar{\gamma}_2} \in (0, \infty)$.*
2. *Suppose that the marginal value function is unbounded. Almost surely along any path:*

- (a) $\frac{\gamma_2^t}{\gamma_1^t} \rightarrow \frac{1-\bar{\gamma}_1}{\bar{\gamma}_1} \in [0, \infty)$ or $\frac{\gamma_2^t}{\gamma_1^t} \rightarrow \frac{\bar{\gamma}_2}{1-\bar{\gamma}_2} \in (0, \infty]$, if the limit distribution of $(\frac{\gamma_2^t}{\gamma_1^t})_t$ exists.
- (b) Either $\frac{1-\bar{\gamma}_1}{\bar{\gamma}_1} \in [0, \infty)$ or $\frac{\bar{\gamma}_2}{1-\bar{\gamma}_2} \in (0, \infty]$ (or both) are accumulation points.

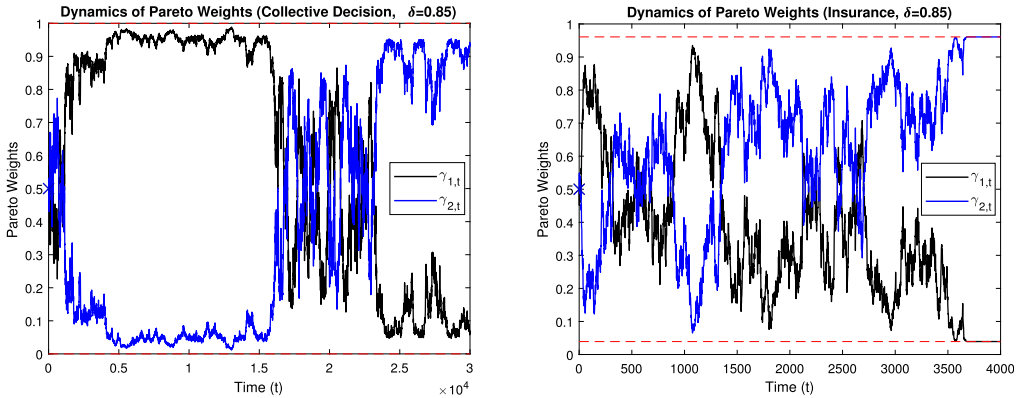


Fig. 2. Sample Paths of $(\gamma_1^t, \gamma_2^t)_t$: Collective Decision (Left) and Insurance (Right).

Theorem 1 is related to the previous immiseration results, such as Thomas and Worrall (1990), which show that no interior point can be in the support of the long-run distribution. The right panel of Fig. 2 illustrates that the decision making process can eventually be absorbed with one of the players becoming a dictator for all future decisions.²⁸ As illustrated in the left panel of Fig. 2, however, it is also possible to have “regime shifts” where each of the two players is arbitrarily close to being a dictator for a long time until a long string of negative shocks eventually shifts almost all of the decision rights to the other player. Importantly though, even when there are regime shifts, the dynamics of inequality as measured by the share of Pareto weights accrued to the better-off player can converge only to 1 over time.

Corollary 1 (Long-run Inequality).

- 1. For any $v^0 \in [v_2^P, \bar{v}_2]$, if $\frac{\gamma_2^t}{\gamma_1^t} \rightarrow \frac{1-\bar{\gamma}_1}{\bar{\gamma}_1}$ or $\frac{\gamma_2^t}{\gamma_1^t} \rightarrow \frac{\bar{\gamma}_2}{1-\bar{\gamma}_2}$, then

$$\lim_{t \rightarrow \infty} \max \left(\frac{\gamma_1^t - \underline{\gamma}_1}{\bar{\gamma}_1 - \underline{\gamma}_1}, \frac{\gamma_2^t - \underline{\gamma}_2}{\bar{\gamma}_2 - \underline{\gamma}_2} \right) = 1.$$

- 2. For any $v^0 \in (v_2^P, \bar{v}_2)$ belonging to Case (ii), almost surely along any path,

$$\limsup_{t \rightarrow \infty} \max \left(\frac{\gamma_1^t - \underline{\gamma}_1}{\bar{\gamma}_1 - \underline{\gamma}_1}, \frac{\gamma_2^t - \underline{\gamma}_2}{\bar{\gamma}_2 - \underline{\gamma}_2} \right) = 1.$$

Thus, the only way that $\left(\max \left(\frac{\gamma_1^t - \underline{\gamma}_1}{\bar{\gamma}_1 - \underline{\gamma}_1}, \frac{\gamma_2^t - \underline{\gamma}_2}{\bar{\gamma}_2 - \underline{\gamma}_2} \right) \right)_t$ can converge is to 1.

²⁸ As we discretized \mathcal{W}_2^* , we also discretize possible pairs of Pareto weights (γ_1, γ_2) by those that are associated with corresponding discretized promised values. Thus, the next-period Pareto weights are chosen from the discretized grid of Pareto weights.

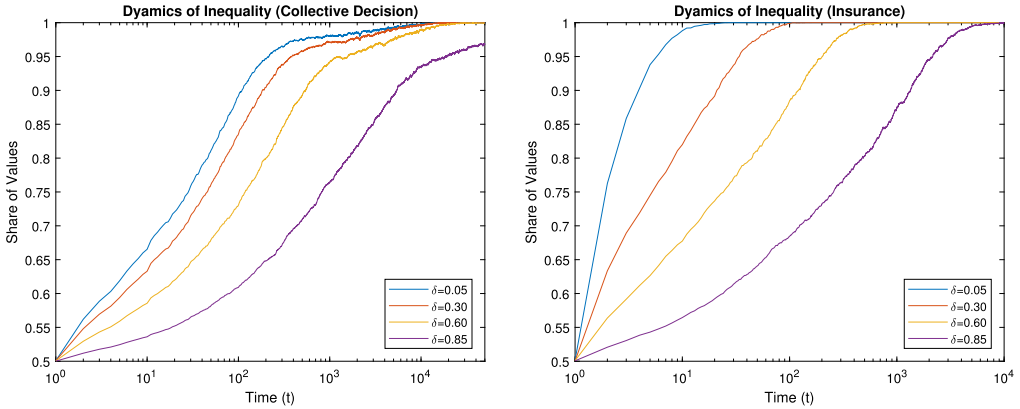


Fig. 3. Dynamics of Inequality: Collective Decision Problem (Left) and Dynamic Insurance Problem (Right).

Our result contrasts with Wang (1995) on existence of non-degenerate limiting wealth distribution.²⁹ Fig. 3 illustrates this point both for the insurance and collective decision models in the left and right panels, respectively. For each discount factor, the dynamics of the share of values captured by the better-off player are calculated as the average of 1500 simulations. It is also worth noting that inequality grows faster when more variation in continuation values is needed to provide incentives. This is the case when δ is closer to zero. Also, since the conflict of interest is stronger in the insurance model than in the collective decision problem, we get faster growth of inequality in the right panel than in the left panel.

In Case (ii), regime shifts are not possible when the marginal value function is bounded. For example, the quadratic collective decision problem violates this premise because $V'(v_2^P) = 0$ and $V'(\bar{v}_2) = -\infty$.³⁰ Intuitively, when one player has almost all the decision rights, moving the decision epsilon away from this player leads to a second-order loss for him while in expectation it leads to a first-order gain to the other player. For this reason, even after a favorable shock θ the optimal contract would make the continuation value of the dominant player increase by ever smaller amounts as it approximates full decision rights. The process thus gets sticky around the dictatorial points.³¹ Instead, when the players are insatiable, as in the insurance model, the gains and losses are both of the first order. This implies that the change in continuation values can remain bounded away from zero. Thus, the process for decision rights can eventually be absorbed into one of the limits.

The possibility of having regime shifts is what drives the use of accumulation points in the statement of Theorem 1. Recall that a relative Pareto weight $\frac{\gamma_2}{\gamma_1}$ is an accumulation point if there is a sub-sequence of the path $(\frac{\gamma_2^t}{\gamma_1^t})_{t \geq 0}$ converging to $\frac{\gamma_2}{\gamma_1}$. Proving this does not rely on the (bounded) sub-martingale convergence theorem. Indeed, the sub-martingale process should be

²⁹ The crucial difference between his and our models is existence of endogenous feasibility constraint. In Wang (1995), the players can transfer only their endowments so that any promised value has to vary each time.

³⁰ More generally, if $v_i < v_i^P$ for each $i \in N$ (and consequently the issue of renegotiation-proofness would matter), then $V'(v_2^P) = 0$ and $V'(\bar{v}_2) = -\infty$.

³¹ Formally, almost surely along any path, for any open neighborhood of v_2^P or \bar{v}_2 , and for any arbitrarily large $N \in \mathbb{N}$, there is $t \in \mathbb{N}$ such that v^τ is in the neighborhood for all $\tau \in \{t, \dots, t + N\}$.

unbounded.³² Thus, we only use the continuity of the continuation value function and Lemma 2 in the proof. This contrasts with Thomas and Worrall (1990), in which one of the players does not have private information and consequently the marginal value function follows a non-positive martingale (i.e., each player's relative Pareto weight follows a non-negative martingale). Then, Thomas and Worrall (1990) invoke the martingale convergence theorem. Unfortunately, once we generalize our set-up as explained above, we can no longer rely on this very powerful tool because, unlike the (non-positive) martingale convergence theorem, the sub-martingale convergence theorem requires the relative Pareto weights to be bounded.³³

We conclude that a dictatorship is a prevalent ex-post consequence of an optimal dynamic mechanism in repeated decision making problems with private information. It is worth pointing out that although the linking mechanism of Jackson and Sonnenschein (2007) does not have this long-run implication, it can lead to even lower values in the long run.

4. Final remarks

We have brought together under one model two strands of literature, collective decision and dynamic insurance problems. We have shown that, in an optimal renegotiation-proof dynamic Bayesian mechanism, the provision of intertemporal incentives necessarily leads to a dictatorial mechanism: in the long run, the optimal mechanism converges to the adoption of one player's favorite action.

Our model leaves several interesting avenues for future research. The first is to incorporate endogenous, type dependent, limited commitment. In environments with endogenous participation constraints, such as Fuchs and Lippi (2006), the threat of abandoning the partnership puts a limit on the extent to which one of the players can dominate the decision process. The second is persistence of private information and correlation of shocks among the players. Although more realistic, such generalizations will pose significant analytical challenges since when considering a deviation in the current report players must also consider the effects such a deviation has on the distribution of beliefs for future periods.³⁴ An additional element to consider in such a generalized environment is the way information is communicated, because it will affect the amount of learning that will take place on equilibrium path. For example, one could have agents confidentially report to a central mechanism that simply reports back the recommended action or a public announcement where all players learn the other players' reports. Despite the difficulties we believe that it is worth trying to incorporate these considerations in future work.

Appendix A

In this Appendix, we explicitly formulate stochastic mechanisms, which are history-dependent stochastic allocations (i.e., probability distributions on the joint actions \mathcal{A}).³⁵ To

³² If this is the case, $w(\cdot|v) \in (\underline{v}_2^P, \bar{v}_2)$ would be obtained as an interior solution for an interior promised value $v \in (\underline{v}_2^P, \bar{v}_2)$. Thus, any interior $v \in (\underline{v}_2^P, \bar{v}_2)$ belongs to Case (ii).

³³ See, for example, Doob (1953, p. 324).

³⁴ Pavan et al. (2014) might prove useful in implementing such an extension. Guo and Hörner (2015) study a dynamic principal-agent problem without monetary transfers in which the valuations of a good follow a two-state Markov chain.

³⁵ Myerson (1979) studies stochastic mechanisms in a static Bayesian collective decision problem. Goltsman et al. (2009) and Kováč and Mylovánov (2009) consider stochastic mechanisms in static mechanism design problems without side-payments in the contexts of cheap talk communication and delegation problems, respectively.

obtain the recursive representation of the problem, we first define the following notations. Let $\Delta(\mathcal{A})$ be the set of all probability measures on the Borel sets on \mathcal{A} . For a given subset \mathcal{W} of \mathbb{R}^2 , a *stochastic (recursive) mechanism* is defined as a pair of mappings $(P, w) : \Theta \ni \theta \mapsto (P(\cdot|\theta), w(\theta)) \in \Delta(\mathcal{A}) \times \mathcal{W}$. We denote by $\mathcal{P} := (\Delta(\mathcal{A}))^\Theta$ the set of stochastic current allocations.

The set $\mathcal{W} \subseteq \mathbb{R}^2$ of PPE payoffs associated with stochastic mechanisms has the same recursive structure as the deterministic mechanisms do. Defining the mapping $B : \mathcal{P}(\mathbb{R}^2) \rightarrow \mathcal{P}(\mathbb{R}^2)$ by the set of payoffs generated by an admissible stochastic mechanism (where admissibility is defined as in the main text), the set of PPE payoffs is given by the largest bounded fixed point of B . Formally, for any $W \in \mathcal{P}(\mathbb{R}^2)$, define $B(W) \in \mathcal{P}(\mathbb{R}^2)$ by

$$B(W) := \left\{ \left(\mathbb{E}_\theta \left[\int (1 - \delta) u_i(a, \theta_i) dP(a|\theta) + \delta w_i(\theta) \right] \right)_{i \in N} \mid (P, w) \text{ is admissible w.r.t. } W \right\}.$$

As in Abreu et al. (1990), it can be shown that the monotone mapping B preserves compactness and that the PPE value set $\mathcal{W} = B(\mathcal{W})$ is a compact convex set.

The frontier of the PPE value set (associated with stochastic mechanisms) has the same recursive representation with player 2's promised value being a state variable as in the deterministic case. Define $V(v)$ as the highest value delivered to player 1 given that player 2's promised value is $v \in \mathcal{W}_2$. The value function $V : \mathcal{W}_2 \rightarrow \mathcal{W}_1$ is characterized as follows.

$$V(v) = \max_{(P, w) \in \mathcal{P} \times \mathcal{W}_2} \mathbb{E}_\theta \left[\int (1 - \delta) u_1(a, \theta_1) dP(a|\theta) + \delta V(w(\theta)) \right] \tag{OPT-S}$$

subject to

$$\mathbb{E}_\theta \left[\int (1 - \delta) u_2(a, \theta_2) dP(a|\theta) + \delta w(\theta) \right] = v; \tag{PK-S}$$

$$\begin{aligned} & \mathbb{E}_{\theta_2} \left[\int (1 - \delta) u_1(a, \theta_1) dP(a|\theta) + \delta V(w(\theta)) \right] \\ & \geq \mathbb{E}_{\theta_2} \left[\int (1 - \delta) u_1(a, \theta_1) dP(a|\hat{\theta}_1, \theta_2) + \delta V(w(\hat{\theta}_1, \theta_2)) \right] \text{ for all } \theta_1, \hat{\theta}_1 \in \Theta_1; \text{ and} \end{aligned} \tag{IC1-S}$$

$$\begin{aligned} & \mathbb{E}_{\theta_1} \left[\int (1 - \delta) u_2(a, \theta_2) dP(a|\theta) + \delta w(\theta) \right] \\ & \geq \mathbb{E}_{\theta_1} \left[\int (1 - \delta) u_2(a, \theta_2) dP(a|\theta_1, \hat{\theta}_2) + \delta w(\theta_1, \hat{\theta}_2) \right] \text{ for all } \theta_2, \hat{\theta}_2 \in \Theta_2. \end{aligned} \tag{IC2-S}$$

The renegotiation-proof value function V is characterized by the Bellman equation (OPT-S) where we replace \mathcal{W}_2 with $\mathcal{W}_2^* = [\underline{v}_2^*, \bar{v}_2^*]$. Note that, by abusing the notation, we denote by V both the (full-commitment) PPE and renegotiation-proof PPE value functions as in the main text.

A.1. Renegotiation-proof PPE

We use the similar argument to Abreu et al. (1990, Theorem 6) to show that the largest bounded fixed point of B^* is non-decreasing in the discount factor.³⁶ Fix a discount factor

³⁶ The (full-commitment) PPE value set is non-decreasing in δ because it is convex (Abreu et al., 1990, Theorem 6).

$\delta \in (0, 1)$ and the largest bounded fixed point $\mathcal{W}^*(\delta)$ of $B_\delta^* := B^*$. Fix $v \in \mathcal{W}^*(\delta)$. Choose an admissible pair (a^δ, w^δ) that delivers v (for ease of notation, we use a deterministic current allocation a^δ). Then, for any $\delta' \in [\delta, 1)$, we consider the following continuation value function of player $i \in N$:

$$w'_i(\cdot) = \frac{(1 - \delta')\delta}{(1 - \delta)\delta'} w_i^\delta(\cdot) + \frac{\delta' - \delta}{(1 - \delta)\delta'} v_i \in [\underline{v}_i^P, \bar{v}_i].$$

Now, observe that

$$\begin{aligned} \mathbb{E}_\theta [(1 - \delta')u_i(a^\delta(\theta), \theta_i) + \delta' w'_i(\theta)] &= \frac{1 - \delta'}{1 - \delta} \mathbb{E}_\theta [(1 - \delta)u_i(a^\delta(\theta), \theta_i) + \delta w^\delta(\theta)] \\ &\quad + \frac{\delta' - \delta}{1 - \delta} v_i = v_i. \end{aligned}$$

It can be seen that the mechanism (a^δ, w') is IC and delivers the expected utility v when the discount factor is δ' . Hence, $\mathcal{W}^*(\delta) \subseteq B_{\delta'}^*(\mathcal{W}^*(\delta))$. This implies that $v \in \mathcal{W}^*(\delta')$, and hence $\mathcal{W}^*(\delta) \subseteq \mathcal{W}^*(\delta')$. Then, we obtain $V_\delta^{\text{SB}}(v_2) \leq V_{\delta'}^{\text{SB}}(v_2)$. Since $v_2 \in [\underline{v}_2^P, \bar{v}_2]$ is arbitrary, we obtain the desired result.

A.2. Approximate efficiency

Proof of Proposition 1 (Approximate Efficiency). It is sufficient to prove that the RP-PPE value function V_δ^{SB} approaches point-wise the first-best value function V^{FB} as δ approaches unity. Indeed, since V_δ^{SB} monotonically converges pointwise to V^{FB} on the compact set \mathcal{W}_2^* , Dini’s theorem (see, for example, Aliprantis and Border (2006, Theorem 2.66)) implies that the convergence is uniform.

We also make the following three remarks. First, since $V_\delta^{\text{SB}}(\cdot) (\leq V^{\text{FB}}(\cdot))$ is non-decreasing in δ , it follows that $V_\delta^{\text{SB}}(\cdot)$ converges uniformly to a concave limit function V_1 , which we show is V^{FB} . Second, since the limit function V_1 and the first-best value function V^{FB} are continuous, it is sufficient to show that the point-wise convergence is almost everywhere. This means that we can assume without loss of generality that the limit concave function is differentiable. Third, the first-best value $V^{\text{FB}}(v)$ cannot exactly be attained for any given $\delta \in (0, 1)$ when $v \in (\underline{v}_2^P, \bar{v}_2)$, provided that the players cannot attain the first-best in the one-shot game. This follows because that would call for the decision $a^{\text{FB}}(\cdot|v)$ to be taken in every period and after every history, so that future decisions would not respond to current announcements. Hence, truth-telling would have to be a static best response for the players. This cannot be attained in an incentive compatible manner.

Now, we show that the RP-PPE value function V_δ^{SB} approaches point-wise the first-best Pareto frontier V^{FB} as δ tends to 1. If $v = \underline{v}_2^P$ ($v = \bar{v}_2$), then player 1’s (2’s) best action $a_1(\theta_1)$ ($a_2(\theta_2)$) is always implemented, and hence $V^{\text{FB}}(v) = V_\delta^{\text{SB}}(v)$ at such v irrespective of δ . Thus, fix $v \in (\underline{v}_2^P, \bar{v}_2)$. Suppose to the contrary that $\lim_{\delta \rightarrow 1} V_\delta^{\text{SB}}(v) = V_1(v) < V^{\text{FB}}(v)$. Thus, $V_\delta^{\text{SB}}(v) \leq V_1(v) < V^{\text{FB}}(v)$ for all $\delta \in (0, 1)$.

Recall that the Lagrangian of the first-best problem (FB) when player 2’s promised value is v is $\mathcal{L} = \mathbb{E}_\theta \left[\sum_{i=1}^2 \gamma_i u_i(a(\theta), \theta_i) \right] - \gamma_2 v$. Note that (γ_1, γ_2) depends on v . Now, we consider the following continuation value functions which replicate the expected payments of the expected externality mechanism proposed by Arrow (1979) and d’Aspremont and Gérard-Varet (1979). We start with player 2’s continuation value function $w_2(\cdot) = w_2(\cdot|\delta, v)$:

$$w_2(\theta_1, \hat{\theta}_2|\delta, v) := \left(\frac{1-\delta}{\delta}\right) \left[\xi_2(\hat{\theta}_2|v) - \mathbb{E}_{\theta_2} [\xi_2(\theta_2|v)]\right] + \mathbb{E}_{\theta} [u_2(a^{FB}(\theta|v), \theta_2)], \quad (A.1)$$

where $\xi_2(\theta_2|v) := \mathbb{E}_{\theta_1} \left[\frac{\gamma_1}{\gamma_2} u_1(a^{FB}(\theta_1, \theta_2|v), \theta_1)\right]$. Since $w_2(\theta_1, \theta_2)$ does not depend on player 1’s announcements, we sometimes suppress θ_1 .

Player 2’s IC constraint, if a^{FB} is implemented and he faces w_2 as a continuation value function, is

$$\theta_2 \in \operatorname{argmax}_{\hat{\theta}_2 \in \Theta_2} \mathbb{E}_{\theta_1} \left[(1-\delta)u_2(a^{FB}(\hat{\theta}_2, \theta_1|v), \theta_2) + \delta w_2(\hat{\theta}_2, \theta_1) \right] \text{ for each } \theta_2 \in \Theta_2.$$

Since it is only the first term $\xi_2(\hat{\theta}_2)$ that depends on player 2’s announcement $\hat{\theta}_2$ in Expression (A.1), we obtain, for each $\theta_2 \in \Theta_2$,

$$\begin{aligned} & \operatorname{argmax}_{\hat{\theta}_2 \in \Theta_2} \mathbb{E}_{\theta_1} \left[(1-\delta)u_2(a^{FB}(\hat{\theta}_2, \theta_1|v), \theta_2) + \delta w_2(\hat{\theta}_2, \theta_1) \right] \\ &= \operatorname{argmax}_{\hat{\theta}_2 \in \Theta_2} (1-\delta)\mathbb{E}_{\theta_1} \left[\sum_{j=1}^2 \gamma_j u_j(a^{FB}(\theta_1, \hat{\theta}_2), \theta_j) \right]. \end{aligned}$$

Thus, effectively, the player 2 faces the planner’s problem.³⁷ Hence, the announcement $\hat{\theta}_2 = \theta_2$ is optimal for player 2. Player 2’s expected payoff is exactly v .

Next, we define player 1’s continuation value function $w_1(\cdot) = w_1(\cdot|\delta, v)$ as

$$w_1(\theta_1, \theta_2) := V_{\delta}^{SB}(w_2(\theta_2)) + \frac{1-\delta}{\delta} \frac{\gamma_2}{\gamma_1} \left(u_2(a^{FB}(\theta_1, \theta_2), \theta_2) - \max_{\theta \in \Theta} u_2(a^{FB}(\theta), \theta_2) \right).$$

Recall again that w_2 does not depend on θ_1 . It can be seen that (a^{FB}, w_1) respects the player 1’s incentive constraint. Also, when δ is close to unity, (w_1, w_2) is self-enforceable and $w_i(\cdot) \geq \underline{v}_i^P$ for each $i \in N$.

Now, player 1’s expected payoff associated with (a^{FB}, w_1) is at most $V_{\delta}^{SB}(v)$. Thus, we obtain

$$\begin{aligned} V_{\delta}^{SB}(v) &\geq V^{FB}(v) - \frac{\delta}{1-\delta} \left(V_{\delta}^{SB}(v) - \mathbb{E}_{\theta} [V_{\delta}^{SB}(w_2(\theta))] \right) \\ &\quad - (1-\delta) \frac{\gamma_2}{\gamma_1} \left(\max_{\theta \in \Theta} u_2(a^{FB}(\theta), \theta_2) - v \right). \end{aligned}$$

Hence, it suffices to show that the second and third terms of the right-hand side of the above expression go to 0 as δ tends to one. Noting that the third term clearly tends to 0, we follow the arguments by Lockwood and Thomas (1989, Theorem) and Thomas and Worrall (1990, Proposition 4) by showing that the efficiency loss $(V_{\delta}^{SB}(v) - \mathbb{E}_{\theta} [V_{\delta}^{SB}(w_2(\theta))])$ goes to 0 faster than δ goes to 1:

$$\lim_{\delta \rightarrow 1} \frac{\delta}{1-\delta} \left(V_{\delta}^{SB}(v) - \mathbb{E}_{\theta} [V_{\delta}^{SB}(w_2(\theta|\delta, v))] \right) = 0.$$

Recall that V_1 is concave and continuous and is differentiable almost everywhere on $[\underline{v}_2^P, \bar{v}_2]$. Also, the convergence is monotone and uniform. As discussed, we can assume without loss of generality that $\frac{d}{dv} V_1(v)$ exists. Now, we have

³⁷ Recall that a^{FB} and (γ_1, γ_2) depend on v .

$$\begin{aligned} & \frac{\delta}{1-\delta} \mathbb{E}_\theta \left[V_\delta^{\text{SB}}(v) - V_\delta^{\text{SB}}(w_2(\theta)) \right] \\ & \leq \frac{\delta}{1-\delta} \mathbb{E}_\theta \left[\frac{d}{dw} V_\delta^{\text{SB}}(w_2(\theta))(v - w_2(\theta)) \right] \\ & = \mathbb{E}_\theta \left[\frac{d}{dw} V_\delta^{\text{SB}}(w_2(\theta)) (\xi_2(\theta_2) - \mathbb{E}_{\theta_2} [\xi_2(\theta_2)]) \right] \\ & = \mathbb{E}_\theta \left[\left(\frac{d}{dw} V_\delta^{\text{SB}}(w_2(\theta)) - \frac{d}{dw} V_1(v) \right) (\xi_2(\theta_2) - \mathbb{E}_{\theta_2} [\xi_2(\theta_2)]) \right], \end{aligned}$$

where $\frac{d}{dw} V_\delta^{\text{SB}}(w(\theta|\delta, v))$ is an appropriate one-sided derivative (i.e., the right derivative if $w(\theta|\delta, v) > v$ and the left derivative if $w(\theta|\delta, v) < v$). Thus, we obtain

$$\frac{\delta}{1-\delta} \mathbb{E}_\theta \left[V_\delta^{\text{SB}}(v) - V_\delta^{\text{SB}}(w_2(\theta)) \right] \leq \kappa \mathbb{E}_\theta \left[\left| \frac{d}{dw} V_\delta^{\text{SB}}(w_2(\theta|\delta, v)) - \frac{d}{dw} V_1(v) \right| \right], \tag{A.2}$$

where $\kappa := \max_{\hat{\theta}_2 \in \Theta_2} \left| \xi_2(\hat{\theta}_2) - \mathbb{E}_{\theta_2} [\xi_2(\theta_2)] \right| < +\infty$. Finally, the right-hand side of Expression (A.2) goes to 0 as $\delta \rightarrow 1$, as V_δ^{SB} uniformly converges to a uniformly continuous concave function V_1 . See Online Appendix B. \square

A.3. Dynamics of optimal mechanisms

In order to characterize incentive constraints, we define the expected utility of player i with type $\theta_i^{(k_i)}$ reporting $\hat{\theta}_i^{(k_i)}$ under a given stochastic mechanism (P, w) as follows:

$$\tilde{U}_i \left(\theta_i, \hat{\theta}_i \mid P, w \right) := \begin{cases} \mathbb{E}_{\theta_2} \left[\int (1-\delta)u_1(a, \theta_1) dP(a|\hat{\theta}_1, \theta_2) + \delta V(w(\hat{\theta}_1, \theta_2)) \right] & \text{if } i = 1 \\ \mathbb{E}_{\theta_1} \left[\int (1-\delta)u_2(a, \theta_2) dP(a|\theta_1, \hat{\theta}_2) + \delta w(\theta_1, \hat{\theta}_2) \right] & \text{if } i = 2 \end{cases}$$

Player i 's expected utility of the mechanism at θ_i is given by $\tilde{U}_i(\theta_i|P, w) := \tilde{U}_i(\theta_i, \theta_i|P, w)$. We often omit (P, w) .

A.3.1. Validity of the first-order approach

Next, we provide a condition which validates the first-order approach. To that end, as in the main text, let each $\Theta_i = \{\theta_i^{(k_i)}\}_{k_i=1}^{m_i}$ be in \mathbb{R} such that $\theta_i^{(1)} < \theta_i^{(2)} < \dots < \theta_i^{(m_i)}$. We also assume that u_i is continuously differentiable with respect to θ_i on an open interval including Θ_i .

We show that the following expected monotonicity condition on marginal utility (**MONi-MU-S**) is sufficient to guarantee that the IC constraints are replaced by the local (upward and downward) ones:

$$\mathbb{E}_{\theta_{-i}} \left[\int \frac{\partial}{\partial \theta_i} u_i(a, \tau) dP(a|\theta_i, \theta_{-i}) \right] \text{ is non-decreasing in } \theta_i \text{ at every } \tau. \tag{MONi-MU-S}$$

Next, we show that, in several standard environments (quadratic, CARA, and multiplicative preferences), the incentive constraint (**ICi-S**) can be replaced with the local ones (**ICi-UP-S**) and (**ICi-DW-S**) and the expected monotonicity condition (**MONi-MU-S**). Moreover, if each player's utility function is quadratic, then the monotonicity condition (**MONi-MU-S**) can be replaced by the expected monotonicity condition on the current allocation.

Proposition A.1 (Incentive Compatibility). *A stochastic mechanism (P, w) satisfies (ICi-S) if it satisfies (MONi-MU-S) and the following local IC constraints:*

$$\tilde{U}_i \left(\theta_i^{(k_i)} \mid P, w \right) \geq \tilde{U}_i \left(\theta_i^{(k_i)}, \theta_i^{(k_i+1)} \mid P, w \right) \text{ for each } k_i \in \{1, \dots, m_i - 1\}; \text{ and} \tag{ICi-UP-S}$$

$$\tilde{U}_i \left(\theta_i^{(k_i)} \mid P, w \right) \geq \tilde{U}_i \left(\theta_i^{(k_i)}, \theta_i^{(k_i-1)} \mid P, w \right) \text{ for each } k_i \in \{2, \dots, m_i\}. \tag{ICi-DW-S}$$

The converse is also true in either of the following standard environments: (i) multiplicative preferences (i.e., $u_i(a, \theta_i) = \theta_i \tilde{u}_i(a)$); (ii) CARA preferences in the insurance setting (i.e., $u_i(a, \theta_i) = -\exp(-r_i(\theta_i + (-1)^i a))$ with $r_i > 0$); or (iii) quadratic preferences (i.e., $u_i(a, \theta_i) = -(a - \theta_i)^2$).

Moreover, in the quadratic case, (MONi-MU-S) can be replaced with the following monotonicity condition:

$$\mathbb{E}_{\theta_{-i}} \left[\int ad P(a \mid \theta_i, \theta_{-i}) \right] \text{ is non-decreasing in } \theta_i \in \Theta_i. \tag{MONi-AL-S}$$

Proof of Proposition A.1. The proof consists of two steps.

Step 1. Fix $(k_i, k'_i) \in \{1, 2, \dots, m_i\}^2$ with $k_i > k'_i$. First, (ICi-UP-S) implies that, for all ℓ_i with $k'_i \leq \ell_i < k_i$,

$$\begin{aligned} \tilde{U}_i(\theta_i^{(\ell_i)}) &\geq \tilde{U}_i(\theta_i^{(\ell_i+1)}) + \tilde{U}_i(\theta_i^{(\ell_i)}, \theta_i^{(\ell_i+1)}) - \tilde{U}_i(\theta_i^{(\ell_i+1)}, \theta_i^{(\ell_i+1)}) \\ &= \tilde{U}_i(\theta_i^{(\ell_i+1)}) - (1 - \delta) \int_{\theta_i^{(\ell_i)}}^{\theta_i^{(\ell_i+1)}} \mathbb{E}_{\theta_{-i}} \left[\int \frac{\partial}{\partial \theta_i} u_i(a, \tau) dP(a \mid \theta_i^{(\ell_i+1)}, \theta_{-i}) \right] d\tau. \end{aligned}$$

By summing the above expression up for each ℓ_i with $k'_i \leq \ell_i \leq k_i - 1$, we have

$$(1 - \delta) \sum_{\ell_i=k'_i}^{k_i-1} \int_{\theta_i^{(\ell_i)}}^{\theta_i^{(\ell_i+1)}} \mathbb{E}_{\theta_{-i}} \left[\int \frac{\partial}{\partial \theta_i} u_i(a, \tau) dP(a \mid \theta_i^{(\ell_i+1)}, \theta_{-i}) \right] d\tau \geq \tilde{U}_i(\theta_i^{(k_i)}) - \tilde{U}_i(\theta_i^{(k'_i)}).$$

Now, it follows from the expected monotonicity condition (MONi-MU-S) that

$$\begin{aligned} \tilde{U}_i(\theta_i^{(k'_i)}) &\geq \tilde{U}_i(\theta_i^{(k_i)}) - (1 - \delta) \mathbb{E}_{\theta_{-i}} \left[\int \left\{ u_i(a, \theta_i^{(k_i)}) - u_i(a, \theta_i^{(k'_i)}) \right\} dP(a \mid \theta_i^{(k_i)}, \theta_{-i}) \right] \\ &= \tilde{U}_i(\theta_i^{(k'_i)}, \theta_i^{(k_i)}). \end{aligned}$$

Second, (ICi-DW-S) implies that, for all ℓ_i with $k'_i \leq \ell_i < k_i$,

$$\begin{aligned} \tilde{U}_i(\theta_i^{(\ell_i+1)}) &\geq \tilde{U}_i(\theta_i^{(\ell_i)}) + \tilde{U}_i(\theta_i^{(\ell_i+1)}, \theta_i^{(\ell_i)}) - \tilde{U}_i(\theta_i^{(\ell_i)}, \theta_i^{(\ell_i)}) \\ &= \tilde{U}_i(\theta_i^{(\ell_i)}) + (1 - \delta) \int_{\theta_i^{(\ell_i)}}^{\theta_i^{(\ell_i+1)}} \mathbb{E}_{\theta_{-i}} \left[\int \frac{\partial}{\partial \theta_i} u_i(a, \tau) dP(a \mid \theta_i^{(\ell_i)}, \theta_{-i}) \right] d\tau. \end{aligned}$$

By summing the above expression up for each ℓ_i with $k'_i \leq \ell_i \leq k_i - 1$, we have

$$\tilde{U}_i(\theta_i^{(k_i)}) - \tilde{U}_i(\theta_i^{(k'_i)}) \geq (1 - \delta) \sum_{\ell_i=k'_i}^{k_i-1} \int_{\theta_i^{(\ell_i)}}^{\theta_i^{(\ell_i+1)}} \mathbb{E}_{\theta_{-i}} \left[\int \frac{\partial}{\partial \theta_i} u_i(a, \tau) dP(a|\theta_i^{(\ell_i)}, \theta_{-i}) \right] d\tau.$$

Thus, it follows from the expected monotonicity condition (**MONi-MU-S**) that

$$\begin{aligned} \tilde{U}_i(\theta_i^{(k_i)}) &\geq \tilde{U}_i(\theta_i^{(k'_i)}) + (1 - \delta) \mathbb{E}_{\theta_{-i}} \left[\int \left\{ u_i(a, \theta_i^{(k_i)}) - u_i(a, \theta_i^{(k'_i)}) \right\} dP(a|\theta_i^{(k'_i)}, \theta_{-i}) \right] \\ &= \tilde{U}_i(\theta_i^{(k_i)}, \theta_i^{(k'_i)}). \end{aligned}$$

Step 2. First, suppose that each player has multiplicative preference shocks: $u_i(a, \theta_i) = \theta_i \tilde{u}_i(a)$. Observing that $\frac{\partial}{\partial \theta_i} u_i(a, \theta_i) = \tilde{u}_i(a)$, if a mechanism (P, w) is IC then for any $\theta_i, \hat{\theta}_i \in \Theta_i$,

$$\begin{aligned} 0 \leq & \left(\theta_i - \hat{\theta}_i \right) \left\{ \mathbb{E}_{\theta_{-i}} \left[\int \frac{\partial}{\partial \theta_i} u_i(a, \tau) dP(a|\theta_i, \theta_{-i}) \right] \right. \\ & \left. - \mathbb{E}_{\theta_{-i}} \left[\int \frac{\partial}{\partial \theta_i} u_i(a, \tau) dP(a|\hat{\theta}_i, \theta_{-i}) \right] \right\}. \end{aligned}$$

Hence, we obtain the expected monotonicity condition (**MONi-MU-S**).

Second, suppose that each player’s utility function is of the CARA form in the insurance setting: $u_i(a, \theta_i) = -\exp(-r_i(\theta_i + (-1)^i a))$ with $r_i > 0$. The argument for the multiplicative case can be applied since we can recast this environment as one of the multiplicative preference shocks. Let his utility function be given by $\theta_i \tilde{u}_i(a)$ with $\tilde{u}_i(a) = k_i(-1) \exp(-r_i(-1)^i a)$ and let his type distribution be given by $\tilde{f}_i(\theta_i) = \frac{\exp(-r_i \theta_i)}{k_i \theta_i} f_i(\theta_i) (> 0)$, where $k_i = \int_{\underline{\theta}_i}^{\bar{\theta}_i} \frac{\exp(-r_i \theta_i)}{\theta_i} f_i(\theta_i) d\theta_i (> 0)$.

Third, suppose that each player’s utility is quadratic: $u_i(a, \theta_i) = -(a - \theta_i)^2$. If a mechanism (P, w) is IC, then for any $\theta_i, \hat{\theta}_i \in \Theta_i$,

$$0 \leq 2 \left(\theta_i - \hat{\theta}_i \right) \left\{ \mathbb{E}_{\theta_{-i}} \left[\int adP(a|\theta_i, \theta_{-i}) \right] - \mathbb{E}_{\theta_{-i}} \left[\int adP(a|\hat{\theta}_i, \theta_{-i}) \right] \right\}.$$

Hence, we obtain the expected monotonicity condition on the current allocation (**MONi-AL-S**). Now, noting that $\frac{\partial}{\partial \theta_i} u_i(a, \theta_i) = 2(a - \theta_i)$, (**MONi-MU-S**) and (**MONi-AL-S**) are equivalent. □

We make two remarks. First, for deterministic mechanisms, the expected monotonicity condition on marginal utilities is:

$$\mathbb{E}_{\theta_{-i}} \left[\frac{\partial}{\partial \theta_i} u_i(a(\theta_i, \theta_{-i}), \tau) \right] \text{ is non-decreasing in } \theta_i \in \Theta_i \text{ at every } \tau \in \Theta_i.$$

Likewise, the expected monotonicity condition on the current allocation is:

$$\mathbb{E}_{\theta_{-i}} [a(\theta_i, \theta_{-i})] \text{ is non-decreasing in } \theta_i \in \Theta_i.$$

Second, the expected monotonicity condition (**MONi-MU-S**) also establishes that if the local upward (resp. downward) constraint at type $\theta_i^{(k_i)}$ is binding then the local downward (resp. upward) constraint at type $\theta_i^{(k_i+1)}$ (resp. $\theta_i^{(k_i-1)}$) is satisfied. Suppose that the IC constraint of player i with type θ_i imitating $\hat{\theta}_i$ is binding. Then, the IC constraint of player i with type $\hat{\theta}_i$ imitating θ_i is rewritten as follows:

$$\begin{aligned} & \mathbb{E}_{\theta_{-i}} \left[\int u_i(a, \hat{\theta}_i) dP(a|\hat{\theta}_i, \theta_{-i}) - \int u_i(a, \hat{\theta}_i) dP(a|\theta_i, \theta_{-i}) \right] \\ & \geq \mathbb{E}_{\theta_{-i}} \left[\int u_i(a, \theta_i) dP(a|\hat{\theta}_i, \theta_{-i}) - \int u_i(a, \theta_i) dP(a|\theta_i, \theta_{-i}) \right]. \end{aligned}$$

Now, this inequality is equivalent to:

$$\int_{\hat{\theta}_i}^{\theta_i} \left\{ \mathbb{E}_{\theta_{-i}} \left[\int \frac{\partial}{\partial \theta_i} u_i(a, \tau) dP(a|\hat{\theta}_i, \theta_{-i}) \right] - \mathbb{E}_{\theta_{-i}} \left[\int \frac{\partial}{\partial \theta_i} u_i(a, \tau) dP(a|\theta_i, \theta_{-i}) \right] \right\} d\tau \geq 0,$$

which follows from the expected monotonicity condition.

A.3.2. The Lagrangian associated with the full problem

We have taken the first-order approach throughout the main text. Here we sketch how our results hold irrespective of the first-order approach. For ease of notation, we consider deterministic mechanisms as in the main text. Each $\Theta_i = \{\theta_i^{(k_i)}\}_{k_i=1}^{m_i}$ is a finite set. For every $\theta_i^{(k_i)}, \theta_i^{(k'_i)} \in \Theta_i$, the incentive constraint $U_i(\theta_i^{(k_i)}|a, w) \geq U_i(\theta_i^{(k'_i)}, \theta_i^{(k'_i)}|a, w)$ can be written as follows:

$$U_i(\theta_i^{(k_i)}) - U_i(\theta_i^{(k'_i)}) - (1 - \delta) \mathbb{E}_{\theta_{-i}} \left[u_i(a(\theta_i^{(k'_i)}, \theta_{-i}), \theta_i^{(k_i)}) - u_i(a(\theta_i^{(k'_i)}, \theta_{-i}), \theta_i^{(k'_i)}) \right] \geq 0.$$

We formulate the Lagrangian $\hat{\mathcal{L}}(a, w|\lambda, \gamma, v)$ of the problem (OPT) as:

$$\begin{aligned} \hat{\mathcal{L}}(a, w|\lambda, \gamma, v) &= \sum_{i=1}^2 \gamma_i \sum_{k_i=1}^{m_i} f_i(\theta_i^{(k_i)}) U_i(\theta_i^{(k_i)}) \\ &+ \sum_{i=1}^2 \gamma_i \sum_{k_i=1}^{m_i} \sum_{k'_i=1}^{m_i} \tilde{\lambda}_{k_i, k'_i}^i \left\{ U_i(\theta_i^{(k_i)}) - U_i(\theta_i^{(k'_i)}) \right\} \\ &+ (1 - \delta) \sum_{i=1}^2 \gamma_i \sum_{k_i=1}^{m_i} \sum_{k'_i=1}^{m_i} \tilde{\lambda}_{k_i, k'_i}^i \mathbb{E}_{\theta_{-i}} \left[u_i(a(\theta_i^{(k'_i)}, \theta_{-i}), \theta_i^{(k_i)}) - u_i(a(\theta_i^{(k'_i)}, \theta_{-i}), \theta_i^{(k_i)}) \right]. \end{aligned}$$

If the Lagrange multiplier of the non-local incentive constraints are zero, then the above Lagrangian reduces to the one associated with the first-order approach.

It can be verified that the Lagrangian is written as follows.

$$\begin{aligned} \hat{\mathcal{L}}(a, w|\lambda, \gamma, v) &= \sum_{i=1}^2 \gamma_i \sum_{k_i=1}^{m_i} f_i(\theta_i^{(k_i)}) U_i(\theta_i^{(k_i)}) \left(1 + \frac{\sum_{k'_i=1}^{m_i} (\tilde{\lambda}_{k_i, k'_i}^i - \tilde{\lambda}_{k'_i, k_i}^i)}{f_i(\theta_i^{(k_i)})} \right) \\ &+ (1 - \delta) \sum_{i=1}^2 \gamma_i \sum_{k_i=1}^{m_i} f_i(\theta_i^{(k_i)}) \sum_{k'_i=1}^{m_i} \frac{\tilde{\lambda}_{k_i, k'_i}^i}{f_i(\theta_i^{(k_i)})} \\ &\times \mathbb{E}_{\theta_{-i}} \left[u_i(a(\theta_i^{(k'_i)}, \theta_{-i}), \theta_i^{(k'_i)}) - u_i(a(\theta_i^{(k'_i)}, \theta_{-i}), \theta_i^{(k_i)}) \right]. \end{aligned}$$

Equivalently, we have:

$$\begin{aligned} \hat{\mathcal{L}}(a, w|\lambda, \gamma, v) &= \sum_{i=1}^2 \gamma_i \sum_{k_i=1}^{m_i} f_i(\theta_i^{(k_i)}) U_i(\theta_i^{(k_i)}) \left(1 + \frac{\sum_{k'_i=1}^{m_i} (\tilde{\lambda}_{k_i, k'_i}^i - \tilde{\lambda}_{k'_i, k_i}^i)}{f_i(\theta_i^{(k_i)})} \right) \\ &+ (1 - \delta) \sum_{i=1}^2 \gamma_i \sum_{k_i=1}^{m_i} f_i(\theta_i^{(k_i)}) \sum_{k'_i=1}^{m_i} \frac{(\tilde{\lambda}_{k_i, k'_i}^i - \tilde{\lambda}_{k'_i, k_i}^i)}{f_i(\theta_i^{(k_i)})} \\ &\times \mathbb{E}_{\theta_{-i}} \left[u_i(a(\theta_i^{(k'_i)}), \theta_{-i}, \theta_i^{(k'_i)}) - u_i(a(\theta_i^{(k_i)}), \theta_{-i}, \theta_i^{(k_i)}) \right] \\ &+ (1 - \delta) \sum_{i=1}^2 \gamma_i \sum_{k_i=1}^{m_i} f_i(\theta_i^{(k_i)}) \sum_{k'_i=1}^{m_i} \frac{\tilde{\lambda}_{k'_i, k_i}^i}{f_i(\theta_i^{(k_i)})} \\ &\times \mathbb{E}_{\theta_{-i}} \left[u_i(a(\theta_i^{(k'_i)}), \theta_{-i}, \theta_i^{(k'_i)}) - u_i(a(\theta_i^{(k_i)}), \theta_{-i}, \theta_i^{(k_i)}) \right]. \end{aligned}$$

Note that if the Lagrange multiplier of the non-local incentive constraints are zero, then we obtain the Lagrangian (LAG*) associated with the first-order approach.

Moreover, these expressions yield the player’s (general) virtual utility (which incorporates all the non-local IC constraints). We denote by $\gamma_i \tilde{u}_i(a, \theta_i^{(k_i)})$ the player i ’s instantaneous virtual utility at $\theta_i^{(k_i)}$ with respect to γ_i and λ_i as follows:

$$\begin{aligned} \gamma_i \tilde{u}_i(a, \theta_i^{(k_i)}) &:= \gamma_i \left\{ u_i(a, \theta_i^{(k_i)}) \left(1 + \frac{\sum_{k'_i=1}^{m_i} \tilde{\lambda}_{k_i, k'_i}^i}{f_i(\theta_i^{(k_i)})} \right) - \frac{\sum_{k'_i=1}^{m_i} u_i(a, \theta_i^{(k'_i)}) \tilde{\lambda}_{k'_i, k_i}^i}{f_i(\theta_i^{(k_i)})} \right\} \\ &= \gamma_i \left\{ u_i(a, \theta_i^{(k_i)}) \left(1 + \frac{\sum_{k'_i=1}^{m_i} (\tilde{\lambda}_{k_i, k'_i}^i - \tilde{\lambda}_{k'_i, k_i}^i)}{f_i(\theta_i^{(k_i)})} \right) \right. \\ &\quad \left. + \sum_{k'_i=1}^{m_i} \frac{\tilde{\lambda}_{k'_i, k_i}^i}{f_i(\theta_i^{(k'_i)})} \mathbb{E}_{\theta_{-i}} \left[u_i(a, \theta_i^{(k_i)}) - u_i(a, \theta_i^{(k'_i)}) \right] \right\}. \end{aligned}$$

Hence, the Lagrangian $\hat{\mathcal{L}} = \hat{\mathcal{L}}(a, w|\lambda, \gamma, v)$ can be expressed as

$$\begin{aligned} \hat{\mathcal{L}} &= \mathbb{E}_{\theta} \left[\sum_{i=1}^2 (1 - \delta) \gamma_i \tilde{u}_i(a(\theta), \theta_i) \right. \\ &+ \delta \gamma_1 V(w(\theta)) \left(1 + \frac{\sum_{k'_1=1}^{m_1} (\tilde{\lambda}_{k_1, k'_1}^1 - \tilde{\lambda}_{k'_1, k_1}^1)}{f_1(\theta_1^{(k_1)})} \right) \\ &\left. + \delta \gamma_2 w(\theta) \left(1 + \frac{\sum_{k'_2=1}^{m_2} (\tilde{\lambda}_{k_2, k'_2}^2 - \tilde{\lambda}_{k'_2, k_2}^2)}{f_2(\theta_2^{(k_2)})} \right) \right]. \end{aligned}$$

A.3.3. The dictatorship result

This part of the appendix demonstrates our main theorem. We start by proving Lemmas 1 and 2. The lemmas are proven without relying on the first-order approach. If we knew it was valid then it could be simplified by replacing $\sum_{k'_i=1}^{m_i} (\tilde{\lambda}_{k_i, k'_i}^i - \tilde{\lambda}_{k'_i, k_i}^i)$ with $\tilde{\lambda}_{i, k_i}^{UP} - \tilde{\lambda}_{i, k_i-1}^{UP} + \tilde{\lambda}_{i, k_i}^{DW} - \tilde{\lambda}_{i, k_i+1}^{DW}$ (or assigning zero to the Lagrange multipliers of all the non-local IC constraints).

Proof of Lemma 1. The first-order condition of the Lagrangian with respect to $w(\theta_1^{(k_1)}, \theta_2^{(k_2)}) = w(\theta_1^{(k_1)}, \theta_2^{(k_2)})$ is:

$$V'(w(\theta_1^{(k_1)}, \theta_2^{(k_2)})) \left(1 + \frac{\sum_{k'_1=1}^{m_1} (\tilde{\lambda}_{k_1, k'_1}^1 - \tilde{\lambda}_{k'_1, k_1}^1)}{f_1(\theta_1^{(k_1)})} \right) = -\frac{\gamma_2}{\gamma_1} \left(1 + \frac{\sum_{k'_2=1}^{m_2} (\tilde{\lambda}_{k_2, k'_2}^2 - \tilde{\lambda}_{k'_2, k_2}^2)}{f_2(\theta_2^{(k_2)})} \right).$$

In other words, the next-period relative Pareto weight is given as:

$$\frac{\gamma'_{-i}(\theta_1^{(k_1)}, \theta_2^{(k_2)})}{\gamma'_i(\theta_1^{(k_1)}, \theta_2^{(k_2)})} = \gamma_{-i} \frac{1 + \frac{\sum_{k'_{-i}=1}^{m_{-i}} (\tilde{\lambda}_{k_{-i}, k'_{-i}}^{-i} - \tilde{\lambda}_{k'_{-i}, k_{-i}}^{-i})}{f_{-i}(\theta_{-i}^{(k_{-i})})}}{1 + \frac{\sum_{k'_i=1}^{m_i} (\tilde{\lambda}_{k_i, k'_i}^i - \tilde{\lambda}_{k'_i, k_i}^i)}{f_i(\theta_i^{(k_i)})}}. \tag{A.3}$$

We remark again that, under the first-order approach, Expression (A.3) reduces to

$$\frac{\gamma'_{-i}(\theta_1^{(k_1)}, \theta_2^{(k_2)})}{\gamma'_i(\theta_1^{(k_1)}, \theta_2^{(k_2)})} = \gamma_{-i} \frac{1 + \frac{\tilde{\lambda}_{-i, k_{-i}}^{UP} - \tilde{\lambda}_{-i, k_{-i}-1}^{UP} + \tilde{\lambda}_{-i, k_{-i}}^{DW} - \tilde{\lambda}_{-i, k_{-i}+1}^{DW}}{f_{-i}(\theta_{-i}^{(k_{-i})})}}{1 + \frac{\tilde{\lambda}_{i, k_i}^{UP} - \tilde{\lambda}_{i, k_i-1}^{UP} + \tilde{\lambda}_{i, k_i}^{DW} - \tilde{\lambda}_{i, k_i+1}^{DW}}{f_i(\theta_i^{(k_i)})}},$$

which is Expression (3) in the main text.

Taking expected values of both sides of Expression (A.3) and applying Jensen’s inequality, we obtain

$$\frac{\gamma_i}{\gamma_{-i}} = \mathbb{E}_{\theta_i} \left[\frac{1}{\mathbb{E}_{\theta_{-i}} \left[\frac{\gamma'_{-i}(\theta)}{\gamma'_i(\theta)} \right]} \right] \leq \mathbb{E}_{\theta} \left[\frac{\gamma'_i(\theta)}{\gamma'_{-i}(\theta)} \right].$$

Finally, Lemma 2 implies that the inequality is strict (note that the proof of Lemma 2 does not hinge on the strict inequality). □

Proof of Lemma 2. First, each player i ’s relative Pareto weight cannot be constant because, as we asserted in Proposition 1, full efficiency is not attainable at $v \in (v_2^P, \bar{v}_2)$.

Now, assume that v is such that $w(\theta|v) \in (v_2^P, \bar{v}_2)$ for all $\theta \in \Theta$. Suppose further that $\frac{\gamma'_i(\theta)}{\gamma'_{-i}(\theta)} \leq \frac{\gamma_i}{\gamma_{-i}}$ for all $\theta \in \Theta$. Since Lemma 1 implies $\mathbb{E}_{\theta} \left[\frac{\gamma'_i(\theta)}{\gamma'_{-i}(\theta)} \right] \geq \frac{\gamma_i}{\gamma_{-i}}$, it follows that $\frac{\gamma'_i(\theta)}{\gamma'_{-i}(\theta)} = \frac{\gamma_i}{\gamma_{-i}}$ for all $\theta \in \Theta$. Player i ’s relative Pareto weight is constant, a contradiction.

Next, assume that $\frac{\gamma'_i(\theta)}{\gamma'_{-i}(\theta)} \geq \frac{\gamma_i}{\gamma_{-i}}$ for all $\theta \in \Theta$. Then, $\frac{\gamma'_{-i}(\theta)}{\gamma'_i(\theta)} \leq \frac{\gamma_{-i}}{\gamma_i}$ for all $\theta \in \Theta$. As in the previous argument, this implies that player $-i$ ’s relative Pareto weight is constant, a contradiction. □

Proof of Theorem 1. Part 1. Suppose first that V' is bounded. It follows from Lemma 1 that the process $(\frac{\gamma'_i}{\gamma'_i})_t$ is a bounded sub-martingale, or, $(V'(v_t))_t$ is a bounded super-martingale.

We prove this part in terms of the promised values by invoking the theorem of the maximum to the continuation value function (see, for example, Aliprantis and Border (2006, Section 17.5)).

Thus, we establish the preconditions for the theorem of the maximum. Fix an optimal current allocation. The set of continuation value functions $w : \Theta \rightarrow \mathcal{W}_2^*$ is identified as a compact set $(\mathcal{W}_2^*)^\Theta$. Since V is weakly concave on \mathcal{W}_2^* , the value function is continuous on the interior $(\underline{v}_2^P, \bar{v}_2)$. Indeed, since the largest fixed point \mathcal{W}^* of B^* is a compact set, the value function is continuous on \mathcal{W}_2^* . Then, the set of continuation value functions w which satisfy the promise-keeping and incentive compatibility constraints forms a closed subset of the compact set $(\mathcal{W}_2^*)^\Theta$, which is compact. The correspondence that associates, with each promised value v , the set of continuation value functions that respect the promise-keeping and incentive compatibility constraints is also continuous. The objective function is also continuous. Thus, applying the theorem of the maximum, the set of optimal continuation value functions is a compact-valued and upper hemi-continuous correspondence.

Now, we invoke Doob’s bounded super-martingale convergence theorem as in Thomas and Worrall (1990): $V'(v_t)$ has to converge to some random variable. We show that, almost surely along any path, $V'(v_t)$ converges to a constant in $\{\underline{v}_2^P, \bar{v}_2\}$. If not, consider $v^t \rightarrow v \in (\underline{v}_2^P, \bar{v}_2)$. By Lemma 2, there is $\theta \in \Theta$ such that $w(\theta|v) \neq v$. Since Θ is finite, with probability one, there is a subsequence $(v^{t(n)})_n$ such that the shock θ occurs at each $t(n)$. Since $v^t \rightarrow v$, there is a subsequence $(v^{\tau(m)})_m$ of $(v^{t(n)})_n$ such that $w(\theta|v) = \lim_{m \rightarrow \infty} w(\theta|v^{\tau(m)}) = \lim_{m \rightarrow \infty} v^{\tau(m)+1} = \lim_{m \rightarrow \infty} v^{\tau(m)} = v$, a contradiction.

Part 2. Suppose that V' is unbounded. For (2a), if the limit distribution exists, then the previous argument applies. For (2b), we show that, almost surely along any path, either \underline{v}_2^P or \bar{v}_2 (or both) are accumulation points. To that end, take $\varepsilon > 0$. For each $v \in [\underline{v}_2^P + \varepsilon, \bar{v}_2 - \varepsilon]$, there is $\theta_v \in \Theta$ with $w(\theta_v|v) - v > 0$. Then, there are $\rho_v > 0$ and $\delta_v > 0$ such that $w(\theta_v|v') - v' > \delta_v$ for all $v' \in (v - \rho_v, v + \rho_v)$. Since $\{(v - \rho_v, v + \rho_v)\}_{v \in [\underline{v}_2^P + \varepsilon, \bar{v}_2 - \varepsilon]}$ is an open cover of a compact set $[\underline{v}_2^P + \varepsilon, \bar{v}_2 - \varepsilon]$, there is a finite set $\{v_n\}_{n=1}^N$ such that $[\underline{v}_2^P + \varepsilon, \bar{v}_2 - \varepsilon] \subseteq \bigcup_{n=1}^N (v_n - \rho_{v_n}, v_n + \rho_{v_n})$. Let $\delta = \min_{n \in \{1, \dots, N\}} \delta_{v_n} > 0$. Thus, at each $v \in [\underline{v}_2^P + \varepsilon, \bar{v}_2 - \varepsilon]$, if θ_v occurs, $w(\theta_v|v) - v > \delta$. Thus, after sufficiently many occurrences of θ_v ’s, a promised value v^t will be in $(\bar{v}_2 - \varepsilon, \bar{v}_2)$ from $v \in [\underline{v}_2^P + \varepsilon, \bar{v}_2 - \varepsilon]$. In the analogous manner, for each $v \in [\underline{v}_2^P + \varepsilon, \bar{v}_2 - \varepsilon]$, after sufficiently many occurrences of some θ_v ’s, a promised value v^t will be in $[\underline{v}_2^P, \underline{v}_2^P + \varepsilon)$ from $v \in [\underline{v}_2^P + \varepsilon, \bar{v}_2 - \varepsilon]$. Hence, almost surely along any path, either \underline{v}_2^P or \bar{v}_2 (or both) are accumulation points. □

Proof of Corollary 1. It follows from Theorem 1 that the largest accumulation point of $\left(\left(\frac{\gamma_1^t - \underline{\gamma}_1}{\bar{\gamma}_1 - \underline{\gamma}_1}, \frac{\gamma_2^t - \underline{\gamma}_2}{\bar{\gamma}_2 - \underline{\gamma}_2} \right) \right)_t$ is 1. If $\frac{\gamma_2^t}{\gamma_1^t} \rightarrow \frac{1 - \bar{\gamma}_1}{\bar{\gamma}_1}$ or $\frac{\gamma_2^t}{\gamma_1^t} \rightarrow \frac{\bar{\gamma}_2}{1 - \bar{\gamma}_2}$, then $\left(\left(\frac{\gamma_1^t - \underline{\gamma}_1}{\bar{\gamma}_1 - \underline{\gamma}_1}, \frac{\gamma_2^t - \underline{\gamma}_2}{\bar{\gamma}_2 - \underline{\gamma}_2} \right) \right)_t$ has a unique accumulation point (which is the limit) 1. □

Appendix B. Supplementary material

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.jet.2019.03.007>.

References

- Abreu, D., Pearce, D., Stacchetti, E., 1990. Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica* 58, 1041–1063.
- Aliprantis, C.D., Border, K., 2006. *Infinite Dimensional Analysis: A Hitchhiker's Guide*, third edition. Springer-Verlag Berlin Heidelberg.
- Arrow, K.J., 1979. The property rights doctrine and demand revelation under incomplete information. In: Boskin, M.J. (Ed.), *Economics and Human Welfare*. Academic Press, pp. 23–39.
- Athey, S., Bagwell, K., 2001. Optimal collusion with private information. *Rand J. Econ.* 32, 428–465.
- Atkeson, A., Lucas Jr., R.E., 1992. On efficient distribution with private information. *Rev. Econ. Stud.* 59, 427–453.
- Bergin, J., MacLeod, W., 1993. Efficiency and renegotiation in repeated games. *J. Econ. Theory* 61, 42–73.
- Carrasco, V., Fuchs, W., 2009. Dividing and Discarding: A Procedure for Taking Decisions with Non-transferable Utility. Casella, A., 2005. Storable votes. *Games Econ. Behav.* 51, 391–419.
- d'Aspremont, C., Gérard-Varet, L.-A., 1979. Incentives and incomplete information. *J. Public Econ.* 11, 25–45.
- Doob, J.L., 1953. *Stochastic Processes*. John Wiley and Sons, Inc., New York.
- Drexler, M., Kleiner, A., 2014. Preference Intensities in Repeated Collective Decision-Making.
- Friedman, E., 1998. Risk Sharing and the Dynamics of Inequality.
- Fuchs, W., Lippi, F., 2006. Monetary union with voluntary participation. *Rev. Econ. Stud.* 73, 437–457.
- Fudenberg, D., Levine, D., Maskin, E., 1994. The folk theorem with imperfect public information. *Econometrica* 62, 997–1039.
- Goltsman, M., et al., 2009. Mediation, arbitration and negotiation. *J. Econ. Theory* 144, 1397–1420.
- Green, E.J., 1987. Lending and the smoothing of uninsurable income. In: Prescott, E.C., Wallace, N. (Eds.), *Contractual Arrangements for Intertemporal Trade*. The University of Minnesota Press, Minneapolis, pp. 3–25.
- Guo, Y., Hörner, J., 2015. *Dynamic Mechanism Design without Money*.
- Hortala-Vallve, R., 2012. Qualitative voting. *J. Theor. Polit.* 24, 526–554.
- Jackson, M.O., Sonnenschein, H.F., 2007. Overcoming incentive constraints by linking decisions. *Econometrica* 75, 241–257.
- Kováč, E., Mylovanov, T., 2009. Stochastic mechanisms in settings without monetary transfers: the regular case. *J. Econ. Theory* 144, 1373–1395.
- Lockwood, B., Thomas, J., 1989. Asymptotic efficiency in principal-agent models with hidden information. *Econ. Lett.* 30, 297–301.
- Myerson, R.B., 1979. Incentive compatibility and the bargaining problem. *Econometrica* 47, 61–73.
- Myerson, R.B., 1981. Optimal auction design. *Math. Oper. Res.* 6, 58–73.
- Myerson, R.B., 1984. Cooperative games with incomplete information. *Int. J. Game Theory* 13, 69–96.
- Myerson, R.B., 1985. Bayesian equilibrium and incentive-compatibility: an introduction. In: Hurwicz, L., Schmeidler, D., Sonnenschein, H. (Eds.), *Social Goals and Social Organization: Essays in Memory of Elisha Pazner*. Cambridge University Press, pp. 229–259.
- Pavan, A., Segal, I., Toikka, J., 2014. Dynamic mechanism design: a Myersonian approach. *Econometrica* 82, 601–653.
- Phelan, C., 1995. Repeated moral hazard and one-sided commitment. *J. Econ. Theory* 66, 488–506.
- Phelan, C., 1998. On the long run implications of repeated moral hazard. *J. Econ. Theory* 79, 174–191.
- Ray, D., 1994. Internally renegotiation-proof equilibrium sets: limit behavior with low discounting. *Games Econ. Behav.* 6, 162–177.
- Skrzypacz, A., Hopenhayn, H., 2004. Tacit collusion in repeated auctions. *J. Econ. Theory* 114, 153–169.
- Thomas, J., Worrall, T., 1990. Income fluctuation and asymmetric information: an example of a repeated principal-agent problem. *J. Econ. Theory* 51, 367–390.
- Wang, C., 1995. Dynamic insurance with private information and balanced budgets. *Rev. Econ. Stud.* 62, 577–595.
- Zhao, R.R., 2006. Renegotiation-proof contract in repeated agency. *J. Econ. Theory* 131, 263–281.
- Zhao, R.R., 2007. Dynamic risk-sharing with two-sided moral hazard. *J. Econ. Theory* 136, 601–640.