Mechanisms for Repeated Bargaining[*]

Andrzej Skrzypacz[†]   Juuso Toikka[‡]

May 3, 2013

Abstract

How does the feasibility of efficient dynamic contracting in repeated trade depend on the features of the bargaining environment such as persistence of values, the agent's private information about the evolution of uncertainty, or trading frequency? To answer this question, we derive a necessary and sufficient condition for efficient, unsubsidized, and individually rational trade in environments satisfying a payoff-equivalence property. The characterization takes the form of a joint restriction on the sensitivity of the expected gains from trade to the agents' initial private information, implying that efficient contracting requires sufficient congruence of the agents' private expectations. This restriction translates to bounds on the persistence of values, or on the amount of asymmetric information about their evolution. It also allows us to distinguish between the effects of increasing patience and of more frequent interaction; if values are autocorrelated, the latter need not facilitate efficiency even when the former does. Finally, we discuss second-best mechanisms, and explain how our results extend to general dynamic Bayesian collective choice problems.

1 Introduction

In this paper, we study the feasibility of efficient contracting in dynamic environments. For concreteness, we cast the analysis in the context of a (finitely or infinitely) repeated bilateral trade problem such as the one faced by the supplier and the buyer of a service which can be

provided in multiple periods. A natural benchmark for such problems is given by the impossibility theorem of Myerson and Satterthwaite (1983): In a one-shot problem with two-sided private information, there do not exist satisfactory trading mechanisms, that is, mechanisms that achieve efficient trade while being incentive compatible, individually rational, and budget balanced. We characterize informational conditions under which this negative result is overturned when the agents bargain over dynamic contracts.

A dynamic setting brings about two novel features: First, the agents' privately-known values may change over time, with the resulting process exhibiting some form of serial dependence. Second, the agents may have private information about this process beyond just knowing their values for the current transaction. For instance, the seller of a service may have superior information not only about his current cost, but also about his long-run average cost, or about the likelihood of shocks to his cost structure. In order to accommodate such multi-dimensional asymmetric information, we model the agents' values as evolving over time according to a pair of Markov processes, the parameters of which may be part of the agents' initial private information along with the starting values of the processes.

We start by showing that satisfactory mechanisms exist if and only if the expected present value of (first-best) gains from trade is no higher than the sum of worst-case expectations of gains from trade from the perspective of each agent. Specifically, an agent's worst-case expectation is computed as the infimum of the expected present value of gains from trade conditional on the agent's initial information, where the infimum is over both the starting value and the parameters of his Markov process. Thus the condition imposes a restriction on the sensitivity of expected gains from trade to the agents' initial information, implying that efficient dynamic contracting requires sufficient congruence of the agents' private expectations.[1]

The above characterization provides a precise, yet simple, answer to the question: How does the feasibility of efficient dynamic contracting depend on the features of the bargaining environment? Namely, anything that makes an agent's worst-case conditional expectation of gains from trade more similar to the unconditional expectation is (weakly) good news for efficiency. Conversely, any change that leads to the unconditional and conditional expectations being more different makes efficiency harder to achieve. For example, extending the duration of the relationship is beneficial only to the extend that the agents have less information about their future values (say, because of shocks to their costs) and that they are sufficiently patient for this to be reflected in the expected present value calculation.

---

[1]In a static problem, or in a dynamic problem with perfectly persistent values and costs, the worst-case expectations never add up to the unconditional expected gains from trade, yielding an alternative proof of the Myerson-Satterthwaite theorem due to Williams (1999) and Krishna and Perry (2000). This is immediate to see when the supports of values and costs coincide as then the worst-case expectation is zero for each agent.

We illustrate how this result can be put to work beyond general level intuitions by considering three applications:

First, for processes with positive serial correlation in the sense of first-order stochastic dominance, we show that the condition for efficiency is harder to satisfy—and hence satisfactory mechanisms are less likely to exist—the more persistent the process. This formalizes the intuition that persistence is detrimental to efficiency as persistent information is difficult to elicit due to it affecting payoffs in many periods.

Second, if the agents have private information about the parameters of their type processes, we show that satisfactory mechanisms may exist despite this information being fully persistent provided that the range of possible processes is "not too large." We also discuss which features of the distributions of the parameters are relevant for the conclusion. Moreover, we show that given enough uncertainty over possible processes, satisfactory mechanisms may fail to exist even if the agents are arbitrarily patient and values are purely transitory.

Third, we distinguish between increasing patience and more frequent interaction, and show that if values are auto-correlated, the latter need not facilitate efficiency even if the former does. This is because an increase in the trading frequency has two effects: it leads to a smaller per-period discount factor, but implies higher correlation of values in adjacent periods as there is less time for values to evolve between trades. The first effect is simply an increase in patience, which makes unconditional and conditional expectations of gains from trade more similar, whereas the second amounts to an increase in persistence, which goes in the opposite direction. As a result, the continuous-time limit may exhibit inefficiency even if efficiency can be achieved in the same model by sending the discount rate to zero. This gives an adverse-selection analog of the result of Abreu, Milgrom, and Pearce (1991), who point out the difference of the two limits for two-sided moral hazard.

For completeness, we also discuss second-best analysis, and illustrate that the dynamics of distortions may be qualitatively different depending on whether the agents' initial information is about their initial value and cost, or about a process parameter.[2]

As we explain in the last section, our characterization and its consequences readily extend to general dynamic Bayesian collective choice problems such as double auctions, public good provision, or allocation of resources within a team. The extension presented there provides a dynamic generalization of Williams' (1999) characterization of the existence of satisfactory mechanisms in static private value environments (see also Krishna and Perry, 2000).[3]

---

[2]This result is closely related to the contrasting findings by Battaglini (2005) and Boleslavsky and Said (2012) who study monopolistic price discrimination under analogous assumptions. However, here the distortions arise in a social-surplus-maximizing mechanism rather than as solutions to a profit-maximization problem.

[3]The existence of satisfactory mechanisms depends also on the specification of the agents' outside options,

In terms of the model, this paper is most closely related to the work of Athey and Miller (2007), and Athey and Segal (2007, 2012). Athey and Miller study repeated bilateral trade when types are iid across periods, thus abstracting from the two features central to our analysis. Athey and Segal (2012) establish a general "folk theorem" for Markov games with transfers, which implies that in a discrete-type version of our problem, if the horizon is infinite, initial information is only about starting values, and processes are ergodic, then satisfactory mechanisms exist given sufficiently little discounting. Athey and Segal (2007) specialize this result to two examples of bilateral trade. In contrast, our characterization allows us to compute the critical discount factor, establish comparative statics (e.g., with respect to persistence and trading frequency), and study the case where initial private information is about process parameters which is ruled out by their ergodicity assumption.

In terms of the methods, we draw on recent advances in dynamic mechanism design, which in turn build on earlier contributions by Baron and Besanko (1984), Courty and Li (2000), Battaglini (2005), and Eső and Szentes (2007) among others. In particular, our analysis applies to environments that satisfy a dynamic version of the payoff-equivalence property familiar from static quasilinear environments. We provide sufficient conditions for it to hold in the Appendix where we establish a payoff-equivalence theorem for environments with multi-dimensional initial information by extending the necessary condition for incentive compatibility by Pavan, Segal, and Toikka (2012, PST henceforth) to our setting. Our discussion of second best also builds on their results. Athey and Segal (2012) and Bergemann and Välimäki (2010) construct efficient dynamic mechanisms, which extend VCG mechanisms to dynamic private-value environments. Our results on ex post budget balance rely on Athey and Segal's method of balancing transfers, which yields a dynamic version of the AGV mechanism of d'Aspremont and Gerard-Varet (1979).

2   The Environment

Consider the following dynamic bargaining environment with two-sided private information: There are two agents, a buyer (B) and a seller (S), who may trade a non-storable good and a numeraire in each of countably many periods indexed by $t = 0, 1, \ldots, T$, with $T \leq \infty$. (For example, the seller can provide a service for which the buyer pays in cash.) If the period-t allocation of the good is $x_t \in \{0, 1\}$ and agent $i \in \{B, S\}$ receives $p_{i,t}$ units of the numeraire, then the resulting flow payoffs are $x_t v_t + p_{B,t}$ for the buyer and $p_{S,t} - x_t c_t$ for the seller. The *value* v_t and *cost* c_t are private information of the respective agents, and evolve

---

or the "status quo." See Segal and Whinston (2011) and the references therein for the static case. However, as this issue is tangential to our analysis, we omit a discussion of it in the interest of space.

4

as described below. The agents evaluate streams of flow payoffs according to the discounted average criterion with a common discount factor $\delta \in (0,1]$ with $\delta < 1$ if $T = \infty$.

The buyer's values are given by a privately-observed stochastic process $V = (V_t)_{t=0}^{T}$ on the (possibly unbounded) interval $\mathcal{V} \subset \mathbb{R}$. Let $\mathcal{V}_0 \subset \mathcal{V}$ and $\Omega_B \subset \mathbb{R}^k$ be convex sets. The law of V is given by a mixture of Markov chains defined by an initial distribution F_0 on $\Omega_B \times \mathcal{V}_0$ and a parameterized family $\{F(\cdot \mid \cdot; \theta_B)\}_{\theta_B \in \Omega_B}$ of Markov kernels $F(\cdot \mid \cdot; \theta_B) : \mathcal{V} \times \mathcal{V} \to [0,1]$.[4] Specifically, given any $(\theta_B, v_0) \in \Omega_B \times \mathcal{V}_0$, V follows the Markov chain $\langle v_0, F(\cdot \mid \cdot; \theta_B) \rangle$ with transitions $F(\cdot \mid \cdot; \theta_B)$ and initial value v_0. The initial condition $(\theta_B, v_0) \in \Omega_B \times \mathcal{V}_0$ is determined as the realization of the random vector (Θ_B, V_0) with distribution F_0.

The buyer's private information in period 0 consists of the vector (θ_B, v_0), and hence from his perspective V is simply the Markov chain $\langle v_0, F(\cdot \mid \cdot; \theta_B) \rangle$. The seller is only informed of the distribution F_0 and the family $\{F(\cdot \mid \cdot; \theta_B)\}_{\theta_B \in \Omega_B}$, which are assumed common knowledge as usual. Thus the buyer's payoff-relevant private information in any period t consists of the transitory component v_t as well as the permanent component θ_B.

Analogously, the seller's privately-observed cost evolves on the interval $\mathcal{C} \subset \mathbb{R}$ according to a process $C = (C_t)_{t=0}^{T}$ generated by the kernels $\{G(\cdot \mid \cdot; \theta_S)\}_{\theta_S \in \Omega_S}$ and the joint distribution G_0 of (Θ_S, C_0) on $\Omega_S \times \mathcal{C}_0$ for some convex sets $\Omega_S \subset \mathbb{R}^k$ and $\mathcal{C}_0 \subset \mathcal{C}$. From the seller's perspective, C is the Markov chain $\langle c_0, G(\cdot \mid \cdot; \theta_S) \rangle$ determined by his period-0 private information (θ_S, c_0), whereas the buyer only knows the pair $(G_0, \{G(\cdot \mid \cdot; \theta_S)\}_{\theta_S \in \Omega_S})$. In every period t, the seller's payoff-relevant private information consists of the vector (θ_S, c_t).

We assume that the processes V and C are independent, which is the dynamic extension of the independent-types assumption familiar from static models. For simplicity, we also assume that the initial distributions F_0 and G_0 have full support on their respective domains. Finally, to ensure that expected allocation utilities are well-behaved, assume $\mathbb{E}[\sum_{t=0}^{T} \delta^t |V_t| \mid \theta_B, v_0]$ and $\mathbb{E}[\sum_{t=0}^{T} \delta^t |C_t| \mid \theta_S, c_0]$ are finite for all $(\theta, v_0, c_0) \in \Omega \times \mathcal{V}_0 \times \mathcal{C}_0$, where $\theta := (\theta_B, \theta_S) \in \Omega_B \times \Omega_S =: \Omega$.[5]

The following three examples illustrate the environment. We use versions of them (along with other examples) throughout the paper.

Example 1 (Conditionally iid types) *Given any $\theta \in \Omega$, values V and costs C are iid draws from distributions $F(\cdot \mid \theta_B)$ and $G(\cdot \mid \theta_S)$, respectively, whose parameters are private information of the agents. For example, θ_i may be the mean (or, more generally, the vector*

---

[4]That is, $F(\cdot \mid v; \theta_B)$ is a cumulative distribution function (cdf) for all $(v, \theta_B) \in \mathcal{V} \times \Omega_B$, and $F(v \mid \cdot; \cdot) : \mathcal{V} \times \Omega_B \to [0,1]$ is (Borel) measurable for all $v \in \mathcal{V}$.

[5]All conditional expectations where the conditioning event is measure zero are to be interpreted as the (unique) version obtained by using the kernels $F(\cdot \mid \cdot)$ and $G(\cdot \mid \cdot)$ while respecting the independence of the agents' type processes.

of the first k moments) of the distribution.

The next example generalizes the previous one by introducing persistence.

Example 2 (Renewal model) *Given $\theta_B \in \Omega_B$, the buyer's value evolves as follows: The initial value V_0 is distributed according to $F_0(\cdot \mid \theta_B)$. For every period $t > 0$, given $V_{t-1} = v$, the distribution of V_t is given by the kernel $F(\cdot \mid v; \theta_B) = \gamma_B \mathbf{1}_{[v,\infty)}(\cdot) + (1 - \gamma_B)F_0(\cdot \mid \theta_B)$ for some $\gamma_B \in [0,1]$. That is, the buyer's type stays constant with probability γ_B, and it is drawn anew from the privately known distribution $F_0(\cdot \mid \theta_B)$ (independently of past types) with the complementary probability. Similarly, given $\theta_S \in \Omega_S$, the seller's initial cost C_0 is distributed according to $G_0(\cdot \mid \theta_S)$, and for every $t > 0$, given $C_{t-1} = c$, the distribution of C_t is $G(\cdot \mid c; \theta_S) = \gamma_S \mathbf{1}_{[c,\infty)}(\cdot) + (1 - \gamma_S)G_0(\cdot \mid \theta_S)$. Note that taking $\gamma_i = 0$ yields the above iid case, whereas $\gamma_i = 1$ corresponds to perfectly persistent types.*

Finally, we consider linear autoregressive processes with Gaussian shocks:

Example 3 (Linear AR(1)) *The buyer's value evolves on $\mathbb{R}$ according to the linear first-order autoregressive, or AR(1), process $v_t = \gamma_B v_{t-1} + (1 - \gamma_B)m_B + \varepsilon_{B,t}$, where m_B is the long-term mean, and $(\varepsilon_{B,t})$ are a sequence of iid draws from a Normal distribution. Similarly, the seller's value evolves on $\mathbb{R}$ according to $c_t = \gamma_S c_{t-1} + (1 - \gamma_S)m_S + \varepsilon_{S,t}$, where $(\varepsilon_{S,t})$ are iid draws from a Normal distribution. The constants γ_i, m_i, and the parameters of the Normal distributions may be private information.*

3 Trading Games and Mechanisms

By the revelation principle, in order to characterize incentive compatible outcomes, it is without loss to focus on truthful equilibria of direct revelation mechanisms where in each period the agents simply report their new private information, and the mechanism determines the allocation and transfers as a function of the history of reports. However, in a dynamic setting the set of implementable decision rules in general depends on the degree of transparency in the mechanism. The most permissive results are achieved with the least amount of information disclosure as hiding information from an agent amounts to pooling his incentive constraints (see, e.g., Myerson, 1986). As will be clear from our results, for the purposes of the current paper this issue can be addressed by considering only the two extreme cases: a fully transparent mechanism, where all reports, allocations, and transfers are public, and the fully opaque mechanism, where reports are confidential and transfers are never observed by either agent. (The latter is best viewed as a purely theoretical construct.)

Formally, a *decision rule* is a (measurable) map

$$\mu = (x, p) : \Omega \times (\mathcal{V} \times \mathcal{C})^{T+1} \to \{0, 1\}^{T+1} \times \mathbb{R}^{2(T+1)},$$

where x is the *allocation rule* and p is the *transfer rule*, whose period-t components x_t and p_t are functions only of the (reported) parameters, values, and costs in periods $0, \ldots, t$.[6] A decision rule μ induces a multi-stage game form where in every period $t = 0, 1, \ldots, T$, given history of reports $(\hat{\theta}_B, \hat{\theta}_S, \hat{v}_0, \hat{c}_0, \ldots, \hat{v}_{t-1}, \hat{c}_{t-1}) \in \Omega \times (\mathcal{V} \times \mathcal{C})^t$, timing is as follows:

t.1 The agents privately observe their own current types $v_t \in \mathcal{V}$ and $c_t \in \mathcal{C}$ (or $(\theta_B, v_0) \in \Omega_B \times \mathcal{V}_0$ and $(\theta_S, c_0) \in \Omega_S \times \mathcal{C}_0$ if $t = 0$).

t.2 The buyer reports $\hat{v}_t \in supp\, F(\cdot \mid \hat{v}_{t-1}; \hat{\theta}_B)$ and the seller reports $\hat{c}_t \in supp\, G(\cdot \mid \hat{c}_{t-1}; \hat{\theta}_S)$ (or, respectively, $(\hat{\theta}_B, \hat{v}_0) \in \Omega_B \times \mathcal{V}_0$ and $(\hat{\theta}_S, \hat{c}_0) \in \Omega_S \times \mathcal{C}_0$ if $t = 0$).

t.3 The decision rule μ determines the allocation $x_t \in \{0, 1\}$ and transfers $(p_{B,t}, p_{S,t}) \in \mathbb{R}^2$ as a function of the reports $(\hat{\theta}_B, \hat{\theta}_S, \hat{v}_0, \hat{c}_0, \ldots, \hat{v}_t, \hat{c}_t)$.

Note that the agents are restricted to reporting types that are consistent with the supports of the type process.

If the agents observe each other's reports at stage $t.2$ as well as the allocation and transfers at stage $t.3$, then this game form is the *public mechanism with decision rule μ*, or simply the *public mechanism μ*. In contrast, if each agent observes neither transfers nor the other agent's reports, then the above game form is the *blind mechanism with decision rule μ*, or the *blind mechanism μ* for short. In what follows, we use 'mechanism μ' as the general term to refer to both the public and the blind mechanism with decision rule μ.

The set of feasible period-t histories of agent i in a mechanism μ is denoted $H_{i,t}^\mu$, or simply $H_{i,t}$, if μ is clear from the context.[7] A strategy for the buyer is then a sequence of (measurable) functions $\sigma_B = (\sigma_{B,t})_{t=0}^T$ where $\sigma_{B,0} : H_{B,0} \to \Omega_B \times \mathcal{V}_0$ and $\sigma_{B,t} : H_{B,t} \to \mathcal{V}$ for $t \geq 1$. The seller's strategies $\sigma_S = (\sigma_{S,t})_{t=0}^T$ are defined analogously. We say that agent i's history $h_{i,t} \in H_{i,t}$ is *truthful* if his own reports have been truthful in all periods $0, \ldots, t-1$.

---

[6]Restricting attention to deterministic mechanisms is without loss for our results so we do so throughout to simplify notation.

[7]Formally, the sets of buyer's histories in a public mechanism are defined recursively by setting $H_{B,0} := \Omega_B \times \mathcal{V}_0$, $H_{B,1} := H_{B,0} \times \Omega \times \mathcal{V}_0 \times \mathcal{C}_0 \times \{0, 1\} \times \mathbb{R}^2 \times \mathcal{V}$, and $H_{B,t} := H_{B,t-1} \times \mathcal{V} \times \mathcal{C} \times \{0, 1\} \times \mathbb{R}^2 \times \mathcal{V}$ for $t \geq 2$. (I.e., in each period, the history is augmented with the previous-period reports, allocation, and transfers, and the buyer's new value.) In a blind mechanism these become $H_{B,0} := \Omega_B \times \mathcal{V}_0$, $H_{B,1} := H_{B,0} \times \Omega_B \times \mathcal{V}_0 \times \mathcal{V} \times \{0, 1\}$, and $H_{B,t} := H_{B,t-1} \times \mathcal{V} \times \mathcal{V} \times \{0, 1\}$ for $t \geq 2$. (I.e., in each period, the history is augmented with the buyer's previous-period report and allocation as well as his new value.) The sets $H_{S,t}$ for the seller are defined analogously.

We say that strategy σ_i for agent i is *truthful*, denoted σ_i^*, if it reports truthfully at all truthful histories.

A mechanism μ and a strategy profile $\sigma := (\sigma_B, \sigma_S)$ induce an allocation process X on $\{0,1\}$ and a payment process $P := (P_B, P_S)$ on $\mathbb{R}^2$ in the obvious way. For any period t and truthful histories $h_{B,t} \in H_{B,t}$ and $h_{S,t} \in H_{S,t}$, we denote the expected continuation utilities for the buyer and the seller, respectively, by

$$U_t^{\mu,\sigma}(h_{B,t}) := \mathbb{E}^{\mu,\sigma}\left[\frac{1-\delta}{1-\delta^{T+1}}\sum_{\tau=t}^{T}\delta^{\tau}(X_{\tau}V_{\tau} + P_{B,\tau}) \mid h_{B,t}\right],$$

and

$$\Pi_t^{\mu,\sigma}(h_{S,t}) := \mathbb{E}^{\mu,\sigma}\left[\frac{1-\delta}{1-\delta^{T+1}}\sum_{\tau=t}^{T}\delta^{\tau}(P_{S,\tau} - X_{\tau}C_{\tau}) \mid h_{S,t}\right].$$

Our convention is to omit σ if the strategies are truthful (e.g., $U_t^{\mu} := U_t^{\mu,\sigma^*}$).

The following definitions are standard (see, e.g., Athey and Segal, 2012, Bergemann and Välimäki, 2010, or PST):

Definition 1 *A mechanism $\mu = (x,p)$ is* efficient *(E) if for all t, $x_t = \mathbf{1}_{\{v_t \geq c_t\}}$.*

Definition 2 *A mechanism μ is* Bayesian incentive compatible *(IC) if for all strategies σ_B, σ_S, and all $(\theta, v_0, c_0) \in \Omega \times \mathcal{V}_0 \times \mathcal{C}_0$,*

$$U_0^{\mu}(\theta_B, v_0) \geq U_0^{\mu,(\sigma_B,\sigma_S^*)}(\theta_B, v_0) \quad and \quad \Pi_0^{\mu}(\theta_S, c_0) \geq \Pi_0^{\mu,(\sigma_B^*,\sigma_S)}(\theta_S, c_0).$$

The mechanism μ is perfect Bayesian incentive compatible *(PIC) if it is IC and the game induced by the mechanism has a perfect Bayesian equilibrium in truthful strategies. The mechanism is* within-period ex post incentive compatible *(EPIC) if it is PIC and for all t, all pairs of truthful histories $(h_{B,t}, h_{S,t}) \in H_{B,t} \times H_{S,t}$, and each agent i, a truthful strategy is a best response for agent i even if he knows the history $h_{j,t}$ of agent $j \neq i$.*

An agent can never observe that the other agent has deviated from truthful reporting due to the restriction on reports at stage $t.2$, and hence the difference between IC and PIC is essentially just the fact that the former requires sequential rationality only almost surely, whereas the latter imposes it everywhere.

Definition 3 *A mechanism μ is* individually rational in period 0 *(IR$_0$) if for all $(\theta, v_0, c_0) \in \Omega \times \mathcal{V}_0 \times \mathcal{C}_0$,*

$$U_0^{\mu}(\theta_B, v_0) \geq 0 \quad and \quad \Pi_0^{\mu}(\theta_S, c_0) \geq 0.$$

The mechanism μ is individually rational *(IR) if for all t, and all truthful histories $(h_{B,t}, h_{S,t}) \in H_{B,t} \times H_{S,t}$,*

$$U_t^\mu(h_{B,t}) \geq 0 \quad and \quad \Pi_t^\mu(h_{S,t}) \geq 0.$$

Individual rationality in period 0 corresponds to a situation where the agents, having observed their initial private information, decide whether to commit to a long-term contract or to take their outside option, which yields a payoff of zero. For individual rationality we require, in addition, that the agents expected continuation payoffs under truthful reporting remain nonnegative in all future periods.

Definition 4 *A mechanism μ is* ex ante budget balanced *(BB$_0$) if*

$$\mathbb{E}^\mu \left[\sum_{t=0}^{T} \delta^t \left(P_{B,t} + P_{S,t} \right) \right] \leq 0.$$

The mechanism $\mu = (x, p)$ is budget balanced *(BB) if $p_B + p_S \equiv 0$.*

Ex ante budget balance is the relevant notion in situations where the operation of the mechanism can be financed by a third party. Then it corresponds to the requirement that, in terms of the expected present value, the profit to the third party be nonnegative. In contrast, in a budget balanced mechanism the sum of transfers is identically zero at all possible reporting histories, and hence such a mechanism requires no outside financing.

Remark 1 *Fix a decision rule (x, p). Let μ and η denote the public and the blind mechanism with decision rule (x, p). Observe that in every period t and for each agent i, the set of histories $H_{i,t}^\eta$ in the blind mechanism η corresponds to a partition of the set of histories $H_{i,t}^\mu$ in the public mechanism μ. Therefore, if μ has any of the properties listed in Definitions 1–4, then η has the same property, but the converse is clearly not true in general. More generally, fix any mechanism λ with decision rule (x, p) and with an arbitrary information disclosure policy (e.g., λ may involve sending noisy, private signals to the players about the history of the other agent's reports and transfers). The information structure in the blind mechanism η amounts to a coarsening of the agents' information in the mechanism λ, and hence if λ has any of the properties in Definitions 1–4, then η has the same property. (Indeed, this is simply the dynamic revelation principle.) On the other hand, the information structure in the public mechanism μ refines the information in λ, which makes the properties harder to satisfy.*

4   A Characterization

We provide a necessary and sufficient condition for the existence of mechanisms that are efficient, individually rational, and budget-balanced. The tightness of our condition relies on the following property of the bargaining environment:

Definition 5 *The environment has the* payoff-equivalence property *if for all IC mechanisms* $\mu = (x,p)$ *and* $\eta = (x',p')$ *such that* $x = x'$, *there exist constants* $a, b \in \mathbb{R}$ *such that* $U_0^\mu = U_0^\eta + a$ *and* $\Pi_0^\mu = \Pi_0^\eta + b$.

The above notion is a dynamic analog of the familiar static payoff-equivalence (or revenue-equivalence) property, which here obtains as a special case by taking $T = 0$. Similarly to static settings, the property is satisfied in sufficiently well-behaved environments. We provide an informal discussion of sufficient conditions for payoff equivalence at the end of this section, and relegate a formal account—including a novel payoff-equivalence theorem for dynamic models with multi-dimensional initial information (Proposition A.1)—to Appendix A.1.

Denote the *first-best gains from trade* by

$$Y := \frac{1-\delta}{1-\delta^{T+1}} \sum_{t=0}^{T} \delta^t (V_t - C_t)^+,$$

where for any $a \in \mathbb{R}$, we write $a^+ := \max\{0, a\}$. Our main tool for the analysis of repeated bargaining is given by the following result, the proof of which can be found in Appendix A.2 (along with all other omitted proofs).

Proposition 1 *Suppose that the environment has the payoff-equivalence property. Then the following are equivalent:*

1. *The first-best gains from trade, Y, satisfy*

$$\inf_{\theta_B, v_0} \mathbb{E}[Y \mid \theta_B, v_0] + \inf_{\theta_S, c_0} \mathbb{E}[Y \mid \theta_S, c_0] \geq \mathbb{E}[Y]. \tag{1}$$

2. *There exists a blind mechanism that is E, IC, IR_0, and BB_0.*

3. *There exists a public mechanism that is E, PIC, IR_0, and BB.*

4. *There exists a public mechanism that is E, EPIC, IR, and BB_0.*

Remark 2 *If (1) is not satisfied, then for each of the statements 2–4, the minimum (expected) subsidy required for the existence of a mechanism having the listed properties is given by* $\mathbb{E}[Y] - \inf_{\theta_B, v_0} \mathbb{E}[Y \mid \theta_B, v_0] - \inf_{\theta_S, c_0} \mathbb{E}[Y \mid \theta_S, c_0]$.

10

Proposition 1 shows that inequality (1) is a necessary and sufficient condition for the existence of a bargaining mechanism that delivers efficient, unsubsidized, and voluntary trade. Note that the properties invoked in the second statement are arguably the weakest possible requirements for a satisfactory mechanism as incentive compatibility is imposed without perfection, balancing the budget may rely on an unbounded credit line, the agents are able to commit to a long-term contract, and minimal feedback is provided to the agents as the mechanism is blind—see Remark 1. Thus there is no scope for relaxing (1).

In the other direction, the third and the fourth statement serve to show, respectively, that (1) is in fact sufficient for efficient trade to be perfect Bayesian incentive compatible in a *public* mechanism that is either budget balanced period by period and individual rational in period 0, or individually rational in every period and budget balanced ex ante. (In the latter case perfect Bayesian incentive compatibility can be strengthened to within-period ex post incentive compatibility.) However, (1) is not in general sufficient to simultaneously guarantee ex post budget balance and individual rationality in every period, i.e., it is not enough for the existence of a mechanism that is E, PIC, IR, and BB.[8]

For the special case of a static model (i.e., for $T = 0$), Proposition 1 follows from a characterization by Williams (1999) (see also Krishna and Perry, 2000). However, in the static model, inequality (1) is satisfied only in the trivial case where there is common knowledge of positive gains from trade (see Myerson and Satterthwaite, 1983). In contrast, in dynamic settings the existence of future surplus allows (1) to be satisfied in many cases where the gains from trade are not certain.

In order to interpret inequality (1), we note that it imposes a joint restriction on the sensitivity of the expected (first-best) gains from trade, $\mathbb{E}[Y]$, on each agent's initial private information.[9] In our dynamic setting, there are two new channels through which this information matters: First, fixing the parameters $\theta = (\theta_B, \theta_S)$, V and C are Markov processes, and hence varying v_0 or c_0 will in general vary the distribution of the period-t gains from trade, $(V_t - C_t)^+$, in every period t (rather than just in period 0), with the effect being more pronounced the more persistent the process. Second, for fixed initial values (v_0, c_0), the distri-

---

[8]Athey and Segal (2007, 2012) show that such a mechanism exists in a broad class of ergodic Markov environments with known parameters and finitely many types if the agents are sufficiently patient. We conjecture that the same is true in our model provided that (1) is satisfied for some $\delta < 1$. See also Athey and Miller (2007) for an exploration of different combinations of properties in the case of iid types drawn from known distributions.

[9]Trivially, if at most one agent has private information in period 0 (i.e., if $\Omega_B \times \mathcal{V}_0$ or $\Omega_S \times \mathcal{C}_0$ is a singleton), then at least one of the infimum terms equals $\mathbb{E}[Y]$, implying that the inequality is satisfied. Thus satisfactory mechanisms exist in this case despite there being bilateral private information in the future. On the other hand, if initial private information is two-sided, and it is possible that one of the agents knows based on his initial information that there are no gains from trade (i.e., if $\inf \mathbb{E}[Y \mid \theta_B, v_0] \wedge \inf \mathbb{E}[Y \mid \theta_S, c_0] = 0$), then the inequality is never satisfied.

11

bution of Y may depend on the privately known parameters θ. Inequality (1) imposes a joint lower bound on the most pessimistic period-0 expectations about Y that may be held by each agent. Hence we may interpret Proposition 1 as saying that the agents' expectations have to be sufficiently congruent for bilateral bargaining to be efficient. For example, it immediately follows that repeated interaction (i.e., $T > 0$) is beneficial only if there is less asymmetric information about the future than the present.

Since the payoff-equivalence property is taken as given, the proof of Proposition 1 parallels that of the static case by Williams (1999), with a dynamic VCG scheme and the balancing trick of Athey and Segal (2012) replacing, respectively, the static VCG and AGV mechanisms in the argument. Analogously to the static setting, the proof provides another interpretation of inequality (1): the terms on the left are the utilities of the worst initial types of the buyer and the seller under the repetition of the static Pivot mechanism, whereas the term on the right is the expected budget deficit under that scheme. Hence, (1) is the condition under which this simple mechanism could be financed (in expectation) by charging type-independent participation fees in period 0. If (1) is not satisfied, a mechanism designer wanting to achieve efficient, individually rational, and budget balanced trade (in any mechanism) would need to subsidize it by exactly the shortfall.

In the remainder of this section, we discuss sufficient conditions for the environment to have the payoff-equivalence property, relegating most of the technical details to Appendix A.1. The Reader more interested in the implications of Proposition 1 may proceed directly to the next section.

Recall that in static models where each agent's type is parameterized by a real vector, the payoff-equivalence property can be established as an implication of an appropriately formulated envelope theorem (see Milgrom and Segal, 2002). The regularity assumption required for this approach is that each agent's utility function vary smoothly in the type with a bounded derivative. In a dynamic setting, this translates, roughly, to the requirement that a change in an agent's initial private information have a smooth effect on his expected utility over dynamic allocation and transfer rules. As the expected utility depends on both the (Bernoulli) utility function as well as the stochastic process for the agent's future types, a condition stated on the primitives must require that both the utility function and the stochastic process vary "sufficiently smoothly" in the agent's initial private information. PST provide sufficient conditions for this in a general class of dynamic models where each agent's new private information is one dimensional in every period, and establish an envelope theorem for such problems (PST, Theorem 1).

Our environment is not covered by the results of PST as here the agents' initial information is multi-dimensional. However, by combining the standard payoff-equivalence argument

for static multi-dimensional models (see, e.g., Holmström, 1979) with the dynamic envelope formula of PST, it is possible to establish a payoff-equivalence result for dynamic models with multi-dimensional initial information. We provide general sufficient conditions for this in Appendix A.1 by introducing the notion of a *regular* environment (Definition A.2) and proving that such environments have the payoff-equivalence property (Proposition A.1). These conditions are somewhat cumbersome to state, and hence it is convenient to introduce the following easy-to-check smoothness condition, which is a special case of regular environments (see Lemma A.1), and hence (grossly) sufficient for payoff equivalence:

Definition 6 *Let $Z = (Z_t)_{t=0}^T$ be a stochastic process on the interval $\mathcal{Z} \subset \mathbb{R}$ generated by the kernels $\{H(\cdot \mid \cdot; \theta_i)\}_{\theta_i \in \Omega_i}$ and the initial distribution H_0 on the convex set $\Omega_i \times \mathcal{Z}_0 \subset \mathbb{R}^k \times \mathcal{Z}$. The process Z is* smooth *if the following conditions hold:*

1. *Every $H(\cdot \mid z; \theta_i)$, $(z, \theta_i) \in \mathcal{Z} \times \Omega_i$, is absolutely continuous with density $h(\cdot \mid z; \theta_i)$ strictly positive on $\mathcal{Z}$.*

2. *The kernel $H(\cdot \mid \cdot; \cdot) : \mathbb{R}^{k+2} \to \mathbb{R}$ is continuously differentiable and there exist constants $b < \frac{1}{\delta}$ and $d < \infty$ such that for all $z, z' \in \text{int } \mathcal{Z}$ and all $\theta_i \in \text{int } \Omega_i$,*

$$\frac{|\partial_z H(z' \mid z; \theta_i)|}{h(z' \mid z; \theta_i)} < b \quad and \quad \frac{\|\nabla_{\theta_i} H(z' \mid z; \theta_i)\|}{h(z' \mid z; \theta_i)} < d.$$

The environment is smooth *if the processes V and C are smooth.*

The first condition is an assumption of "continuous types." (It also imposes full support and rules out atoms, which is not essential, but simplifies exposition.) The second condition ensures that the process Z is a sufficiently smooth Lipschitz function of the initial information (z_0, θ_i), which is comparable to the differentiability and bounded-derivative assumptions invoked in static models. For example, the conditionally iid types in Example 1 are smooth if $F(\cdot \mid \theta_B)$ and $G(\cdot \mid \theta_B)$ are absolutely continuous given any θ_B and θ_S, and depend sufficiently regularly on the parameters. Similarly, it can be verified that if the parameters of the linear AR(1) processes of Example 3 are common knowledge, then the processes are smooth simply whenever $\gamma_i < \frac{1}{\delta}$ for $i \in \{B, S\}$. In contrast, the renewal model of Example 2 is clearly not smooth as the transitions have atoms (but it is regular in the sense of Definition A.2 by Lemma A.2).

5   Applications

We now turn to the main part of our analysis: putting Proposition 1 to work to explore, for a fixed discount rate, how the feasibility of efficient bargaining depends on persistence, private

information about the process parameters θ, and the frequency of interaction. To facilitate discussion, we say that *satisfactory trading mechanisms exist* if statements 2 through 4 in Proposition 1 are satisfied. By Proposition 1 this is the case in environments that have the payoff-equivalence property if and only if

$$\inf_{\theta_B, v_0} \mathbb{E}[Y \mid \theta_B, v_0] + \inf_{\theta_S, c_0} \mathbb{E}[Y \mid \theta_S, c_0] \geq \mathbb{E}[Y],$$

where $Y := \frac{1-\delta}{1-\delta^{T+1}} \sum_{t=0}^{T} \delta^t (V_t - C_t)^+$ denotes the first-best gains from trade.

5.1 Persistence

We say that the environment is *with known parameters* if the parameter space $\Omega = \Omega_B \times \Omega_S$ is a singleton, in which case we suppress all references to θ in the notation. We restrict attention to such environments in this subsection in order to focus on the persistence of values and costs. It is useful to start by reviewing the following well-known extreme cases:

Example 4 (IID with known parameters) *Consider the environment of Example 1 with known parameters. Then V and C are, respectively, iid draws from the known distributions F_0 and G_0, which we assume to have strictly positive continuous densities everywhere on their domains. This environment is smooth and hence has the payoff-equivalence property by Proposition A.1. For simplicity, take $T = \infty$ and suppose $\mathcal{V} = \mathcal{C}$. Then*

$$\inf_{v_0} \mathbb{E}[Y \mid v_0] = \inf_{c_0} \mathbb{E}[Y \mid c_0] = \delta \mathbb{E}[Y],$$

as the worst initial type of either agent does not trade in period 0 under the efficient allocation rule, but expects first-best trade from period 1 onwards. Hence, by Proposition 1, satisfactory trading mechanisms exist iff $2\delta \mathbb{E}[Y] \geq \mathbb{E}[Y]$, or equivalently, iff $\delta \geq \frac{1}{2}$. By statement 4 of Proposition 1 this replicates the finding by Athey and Miller (2007, Proposition 1), who derived the cutoff by considering recursive mechanisms.[10]

Example 5 (Full persistence with known parameters) *Consider the renewal model of Example 2 with known parameters and $\gamma_i = 1$ for $i \in \{B, S\}$. This environment is regular by Lemma A.2, and hence it has the payoff-equivalence property by Proposition A.1.*[11] *Thus*

---

[10]Their notion of ex ante budget balance requires budget to balance in expectation in every period, whereas our definition only considers the expected balance in period 0. However, with iid types the two are equivalent.

[11]Alternatively, because of the non-changing types, the standard static argument (e.g., Milgrom and Segal, 2002) can be used to establish the payoff-equivalence property as the environment can be viewed as a static problem with one-dimensional types and a $T + 1$-dimensional allocation.

Proposition 1 applies. Note that $Y = \frac{1-\delta}{1-\delta^{T+1}} \sum_{t=0}^{T} \delta^t (V_t - C_t)^+ = (V_0 - C_0)^+$ almost surely for all $T \in \mathbb{N} \cup \{\infty\}$. Hence, given any horizon T, inequality (1) takes the form

$$\inf_{v_0} \mathbb{E}[(v_0 - C_0)^+] + \inf_{c_0} \mathbb{E}[(V_0 - c_0)^+] \geq \mathbb{E}[(V_0 - C_0)^+],$$

which is the condition for the static model. By Remark 2 this implies that the subsidy (in discounted average terms) required for efficient contracting is independent of the horizon. For example, if V_0 and C_0 have the same, nondegenerate support, then the above inequality yields the contradiction $0 \geq \mathbb{E}[(V_0 - C_0)^+] > 0$, implying that inefficiency is inevitable, and that the minimum subsidy needed for efficiency equals the full expected gains from trade. More generally, by Myerson and Satterthwaite (1983), an inefficiency result obtains as long as the intersection of the supports of V_0 and C_0 (i.e., $\mathcal{V}_0 \cap \mathcal{C}_0$) has a nonempty interior. (This can of course also be derived directly by computing the terms in the above inequality.)

These two examples suggest that persistence of private information is detrimental to efficiency. In order to formalize this intuition, we need a notion of persistence that discriminates between Markov chains according to their short-run behavior as payoffs are discounted (or the horizon finite). The following definition provides one such notion in terms of a partial order on ergodic Markov chains.

Definition 7 *Let $Z = (Z_t)_{t=0}^{T}$ and $Z' = (Z'_t)_{t=0}^{T}$ be Markov chains on $\mathcal{Z} \subset \mathbb{R}$ with kernels $H(\cdot \mid \cdot)$ and $H'(\cdot \mid \cdot)$, respectively. We say that Z is* (weakly) more persistent *than Z' if there exists a distribution Φ satisfying the following conditions:*

1. *Φ is the unique invariant distribution admitted by the kernels $H(\cdot \mid \cdot)$ and $H'(\cdot \mid \cdot)$.*

2. *Φ is the initial distribution of the chains Z and Z'.*

3. *For all $(z_0, z) \in \mathcal{Z}^2$ and all $t \geq 1$, the t-step distributions satisfy*

$$\left| H^{(t)}(z \mid z_0) - \Phi(z) \right| \geq \left| H'^{(t)}(z \mid z_0) - \Phi(z) \right|.$$

The first two conditions ensure that an increase in persistence only affects the short-run properties of the Markov chain. The third condition captures the idea that a more persistent chain should take longer to converge to the invariant distribution. For example, if Z is any Markov chain started from its unique invariant distribution Φ, and Z' is a sequence of iid draws from Φ, then Z is more persistent than Z' according to the above definition. More generally, the family of chains $\{Z^\alpha\}_{\alpha \in [0,1]}$ with kernels $H^\alpha(\cdot \mid z) := \alpha H(\cdot \mid z) + (1 - \alpha)\Phi(\cdot)$, $z \in \mathcal{Z}$,

15

where $H(\cdot \mid \cdot)$ is the kernel of $Z = Z^1$, is ordered by persistence with Z^α more persistent than $Z^{\alpha'}$ iff $\alpha \geq \alpha'$. Note that taking Z to be a fully persistent chain yields the renewal model of Example 2, which thus provides one possible parameterization of persistence spanning from iid to permanent types.

We can now show that the kind of persistence captured by Definition 7 is harmful when the transitions of the type processes are order-preserving in the following sense:

Definition 8 Let $Z = (Z_t)_{t=0}^T$ be a Markov chain on $\mathcal{Z} \subset \mathbb{R}$ with kernel $H(\cdot \mid \cdot)$. We say that Z is stochastically monotone if for all $(z, z') \in \mathbb{R}^2$ with $z > z'$, $H(\cdot \mid z)$ first-order stochastically dominates $H(\cdot \mid z')$.

Proposition 2 Let (V, C) and (V', C') be two pairs of stochastically monotone Markov chains with compact sets of initial values (and known parameters). Suppose the following hold:

1. The first-best gains from trade under (V, C) satisfy inequality (1), and

2. V is more persistent than V', and C is more persistent than C'.

Then the first-best gains from trade under (V', C') satisfy inequality (1).[12]

Proof. Suppose (V, C) and (V', C') satisfy the assumptions, and let Y and Y' denote the first-best gains from trade under (V, C) and (V', C'), respectively. By conditions 1 and 2 of Definition 7, $\mathbb{E}[Y] = \mathbb{E}[Y']$. By inspection of (1) it thus suffices to show that

$$\inf_{v_0} \mathbb{E}[Y' \mid v_0] \geq \inf_{v_0} \mathbb{E}[Y \mid v_0], \quad \text{and} \quad \inf_{c_0} \mathbb{E}[Y' \mid c_0] \geq \inf_{c_0} \mathbb{E}[Y \mid c_0].$$

Consider the first inequality. The distributions of V_0 and V'_0 agree by Definition 7 so that by compactness there exists $w := \min \mathcal{V}_0 = \min \mathcal{V}'_0$. Note that the degenerate distribution at w, denoted μ_w, is (first-order stochastically) dominated by every other distribution on $\mathcal{V}_0$. As the t-step distributions of a stochastically monotone chain preserve dominance (see, e.g., Daley, 1968), $F^{(t)}(\cdot \mid w)$ and $F'^{(t)}(\cdot \mid w)$ are dominated by $F^{(t)}(\cdot \mid v_0)$ and $F'^{(t)}(\cdot \mid v_0)$, respectively, for all t and all v_0. Because the period-t gains from trade, $(v_t - c_t)^+$, increase in v_t, this

---

implies that the infima are achieved at $v_0 = w$. Furthermore, it suffices to show that for all t, $F'^{(t)}(\cdot \mid w)$ dominates $F^{(t)}(\cdot \mid w)$. To this end, note that by Condition 3 of Definition 7,

$$\left| F^{(t)}(v \mid w) - \Phi(v) \right| \geq \left| F'^{(t)}(v \mid w) - \Phi(v) \right| \quad \text{for all } v \in \mathcal{V},$$

where Φ is the common invariant distribution. Since Φ dominates μ_w, and the chains are stochastically monotone, Φ dominates $F^{(t)}(\cdot \mid w)$ and $F'^{(t)}(\cdot \mid w)$.[13] Thus we may dispense with the absolute value operator to get

$$F^{(t)}(v \mid w) \geq F'^{(t)}(v \mid w) \quad \text{for all } v \in \mathcal{V},$$

which is equivalent to saying that $F'^{(t)}(\cdot \mid w)$ dominates $F^{(t)}(\cdot \mid w)$. The second inequality involving the infima with respect to c_0 is established analogously. ∎

Taken together, Propositions 1 and 2 imply that satisfactory trading mechanisms are less likely to exist when the type processes are more persistent. The following example illustrates:

Example 6 (Renewals with known parameters) *Consider the renewal model of Example 2 with known parameters. By Lemma A.2 this environment is regular and hence it has the payoff-equivalence property by Proposition A.1. For any γ_B and γ_S the invariant distributions are simply F_0 and G_0, and for any $(v_0, c_0) \in \mathcal{V} \times \mathcal{C}$, the t-step distributions are*

$$F^{(t)}(\cdot \mid v_0) = \gamma_B^t \mathbf{1}_{[v_0, \infty)} + (1 - \gamma_B^t)F_0, \quad \text{and} \quad G^{(t)}(\cdot \mid c_0) = \gamma_S^t \mathbf{1}_{[c_0, \infty)} + (1 - \gamma_S^t)G_0.$$

Thus for all $(v, c) \in \mathcal{V} \times \mathcal{C}$, we have

$$\left| F^{(t)}(v \mid v_0) - F_0(v) \right| = \gamma_B^t \left| \mathbf{1}_{[v_0, \infty)}(v) - F_0(v) \right|,$$

and

$$\left| G^{(t)}(c \mid c_0) - G_0(c) \right| = \gamma_S^t \left| \mathbf{1}_{[c_0, \infty)}(c) - G_0(c) \right|.$$

Therefore, increasing γ_i for $i \in \{B, S\}$ results in a more persistent type process for agent i in the sense of Definition 7. Furthermore, the processes V and C are clearly stochastically monotone. Hence inequality (1) is harder to satisfy for higher values of γ_i by Proposition 2. For example, if $T = \infty$ and $\mathcal{V} = \mathcal{C}$, then straightforward calculations show that (1) is equivalent to

$$\sqrt{\rho_B \rho_S} \geq \frac{1 - \delta}{\delta}, \tag{2}$$

---

[13]To see this, note that the t-step distribution starting from the invariant distribution Φ is Φ itself.

where $\rho_i := 1 - \gamma_i$ is the probability of a renewal for agent i. That is, satisfactory trading mechanisms exist if and only if the geometric average of the agents' renewal probabilities $(\rho_B, \rho_S) \in [0,1]^2$ is high enough, with the threshold being decreasing in patience. (Note that $\delta \geq \frac{1}{2}$ is a necessary condition.) The fact that persistence is substitutable across agents in (2) is a manifestation of the joint restriction on the agents' processes embodied in (1).

When the agents have private information about the process parameters, the above forces are still at play for any given $\theta \in \Omega$. But since inequality (1) involves taking infima with respect to the parameters, what matters then is "worst-case persistence" rather than the persistence of the realized processes. For example, if the supports of the values and costs coincide, then an impossibility result obtains as soon as the least favorable type *may* be an absorbing state for one of the agents, thus generalizing the finding in Example 5.[14] However, as the kernels can in general depend on θ in complicated ways, obtaining clean predictions requires additional structure. A natural special case arises when θ simply parameterizes the persistence of the agents' processes. We consider this in the next subsection after some general remarks on the effects of private information about process parameters (see Example 8).

5.2 Private information about process parameters

We now turn to the possibility that the agents have at time zero private information about the distribution of their future types beyond the information contained in v_0 and c_0. This information is captured by the privately known parameters $\theta_i \in \Omega_i$, $i \in \{B, S\}$. By inspection of inequality (1), if we reduce the asymmetry of information by restricting the parameters to some subsets $\Omega'_B \subset \Omega_B$ and $\Omega'_S \subset \Omega_S$ without affecting the expected gains from trade, $\mathbb{E}[Y]$, then the inequality is easier to satisfy, and, consequently, satisfactory mechanisms are more likely to exist. In other words, efficient trade is harder to achieve when there is more asymmetric information about the type processes. For completeness, we record this observation in the form of a proposition:

Definition 9 *Fix a pair (V, C) of type process generated by $(F_0, \{F(\cdot \mid \cdot; \theta_B)\}_{\theta_B \in \Omega_B})$ and $(G_0, \{G(\cdot \mid \cdot; \theta_S)\}_{\theta_S \in \Omega_S})$. A surplus-neutral truncation of parameters is a pair (V', C') of type process generated by $(F'_0, \{F'(\cdot \mid \cdot; \theta_B)\}_{\theta_B \in \Omega'_B})$ and $(G'_0, \{G'(\cdot \mid \cdot; \theta_S)\}_{\theta_S \in \Omega'_S})$ such that*

1. $\Omega' \subset \Omega$,

---

[14]To see this, suppose the sets $\mathcal{V}_0 = \mathcal{V} = \mathcal{C}$ are closed and bounded from below, and there exists $\theta'_B \in \Omega_B$ such that $F(\min \mathcal{V} \mid \min \mathcal{V}; \theta'_B) = 1$. Then $\inf_{\theta_B, v_0} \mathbb{E}[Y \mid \theta_B, v_0] = \mathbb{E}[Y \mid \theta'_B, \min \mathcal{V}] = 0$, and hence inequality (1) is violated (unless the seller has no initial private information). An obvious sufficient condition for this is that the buyer's value is fully persistent given some $\theta'_B \in \Omega_B$.

2. $F_0' = F_0|_{\mathcal{V} \times \Omega_B'}$ and $G_0' = G_0|_{\mathcal{C} \times \Omega_S'}$,

3. $F'(\cdot \mid \cdot; \theta_B) = F(\cdot \mid \cdot; \theta_B)$ and $G'(\cdot \mid \cdot; \theta_S) = G(\cdot \mid \cdot; \theta_S)$ for all $\theta \in \Omega'$, and

4. $\mathbb{E}[Y] = \mathbb{E}[Y']$.

Proposition 3 *Fix an environment that has the payoff-equivalence property, and where inequality (1) is satisfied. Then any environment where the type processes are given by a surplus-neutral truncation of parameters has the payoff-equivalence property, and inequality (1) is satisfied under the truncated processes.*

The proof is immediate and hence omitted.

Taken together, Propositions 1 and 3 provide a sense in which private information about process parameters is detrimental to efficiency. Indeed, inequality (1) allows us to compute exactly when this is the case. In some specific examples this can even be done in closed form:

Example 7 (IID with private parameters) *Consider the environment of Example 1 with $T = \infty$. Specifically, assume that the distributions belong to the following single-parameter families of linear densities on the unit interval:*

$$
\begin{aligned}
f(v \mid \theta_B) &= \theta_B + 2(1 - \theta_B) v \text{ for } \theta_B \in [0, 2], \\
g(c \mid \theta_S) &= \theta_S + 2(1 - \theta_S) c \text{ for } \theta_S \in [0, 2].
\end{aligned}
$$

The case $\theta_i = 1$ corresponds to the uniform distribution. Note that $f(\cdot \mid \theta_B)$ and $g(\cdot \mid \theta_S)$ decrease in θ_i in the sense of first-order stochastic dominance (i.e., the distributions are the strongest when $\theta_i = 0$ and the weakest for $\theta_i = 2$), and all distributions have full support on $[0, 1]$. Suppose that each Θ_i is distributed over $[\underline{\theta}, 2 - \underline{\theta}]$ for some $\underline{\theta} \in [0, 1]$ according to a continuous distribution which is symmetric around 1 (the distributions can be different for the two agents but for simplicity we take $\underline{\theta}$ to be the same for both). By symmetry, the unconditional distributions of V_t and C_t are then uniform on $[0, 1]$ for all t. This implies that increasing $\underline{\theta}$ to some $\underline{\theta}' > \underline{\theta}$ induces a surplus-neutral truncation of parameters.

It is straightforward to verify that this environment is smooth, and hence Proposition 1 and Proposition 3 apply, implying that efficiency is easier to achieve for higher values of $\underline{\theta}$. In order to find the cutoff, note that the expected gains from trade always equal $\mathbb{E}[Y] = \frac{1}{6}$. The worst initial type of the buyer corresponds to having $v_0 = 0$ and $\theta_B = 2 - \underline{\theta}$, while for the seller it is $(1, \underline{\theta})$. Direct computation then yields

$$
\inf_{\theta_B, v_0} \mathbb{E}[Y \mid \theta_B, v_0] = \inf_{\theta_S, c_0} \mathbb{E}[Y \mid \theta_S, c_0] = \frac{\delta}{12}(1 + \underline{\theta}).
$$

19

This leads us to the following corollary: In the iid case with types drawn from the linear family, satisfactory mechanisms exist if and only if $\underline{\theta} \geq \frac{1-\delta}{\delta}$. In particular, if $\underline{\theta} = 0$, there is no $\delta < 1$ for which this condition is satisfied.

This example illustrates three general points:

1. Even though the parameters θ_i are fully persistent, if there is not "too much" uncertainty about the processes, satisfactory mechanisms exist. However, this requires the agents be more patient than in the case where the processes are known (which corresponds to $\underline{\theta} = 1$, and yields $\delta \geq \frac{1}{2}$ as in Example 4).

2. If there is enough uncertainty about the processes, satisfactory mechanisms may not exist even if the players are arbitrarily patient.

3. The distribution of the parameters θ_i affects feasibility of efficient trade only if it affects the unconditional distribution of (V, C), or the domain of possible parameters Ω (in the example, any distribution symmetric around 1 yields the same bound). To see this, note that if the unconditional distribution of (V, C) stays constant, so does $\mathbb{E}[Y]$. Furthermore, if we change the distribution of θ_B without changing its support, and keeping the unconditional distribution of V fixed, then $\inf_{\theta_S, c_0} \mathbb{E}[Y \mid \theta_S, c_0]$ is unaffected. Similarly, $\inf_{\theta_B, v_0} \mathbb{E}[Y \mid \theta_B, v_0]$ does not change either, as it depends only on Ω_B.

As a second example, we consider an environment where the agents have private information about persistence.

Example 8 (Privately known persistence) *Let $T < \infty$, and construct the type processes V and C as follows. Fix $\Delta > 0$. Define the "base kernels" $\bar{F}(\cdot \mid \cdot)$ and $\bar{G}(\cdot \mid \cdot)$ on $[0,1]^2$ by letting $\bar{F}(\cdot \mid z)$ and $\bar{G}(\cdot \mid z)$ be the time-Δ distribution of a Brownian motion starting from $z \in [0, 1]$ at time 0 and having reflecting boundaries at 0 and at 1. Let Φ denote the cdf of the uniform distribution on $[0, 1]$, which is the invariant distribution for this twice-reflected Brownian motion (see, e.g., Harrison, 1985). The families of kernels for the buyer and seller are then defined by setting for $v, c \in [0, 1] = \mathcal{V} = \mathcal{C}$ and $\theta_i \in \Omega_i \subset [0, 1]$, $i \in \{B, S\}$,*

$$F(\cdot \mid v; \theta_B) = \theta_B \bar{F}(\cdot \mid v) + (1 - \theta_B)\Phi(\cdot),$$

and

$$G(\cdot \mid c; \theta_S) = \theta_S \bar{G}(\cdot \mid c) + (1 - \theta_S)\Phi(\cdot).$$

Let V_0 and C_0 be distributed uniformly on $[0, 1]$ independently of Θ_B and Θ_S, so that conditional on any $\theta \in \Omega$, the distribution of V_t and C_t is simply Φ for all t. The parameters Θ_i are

distributed on Ω_i according to some continuous distribution with full support. Lemma A.3 in the Appendix shows that this environment is regular, and hence it has the payoff-equivalence property by Proposition A.1.

Conditional on any $\theta \in \Omega$, the processes V and C are stochastically monotone in the sense of Definition 8. Since increasing θ_i leads to a more persistent process in the sense of Definition 7,[15] an argument analogous to the proof of Proposition 2 shows that the worst case for each agent corresponds to having $\theta_i = \sup \Omega_i$, and the least favorable type in period 0 (i.e., $v_0 = 0$ or $c_0 = 1$). Moreover, we have $\mathbb{E}[Y] = \frac{1}{6}$ independently of the distribution or the support of Θ_i. Proposition 1 and Proposition 3 then imply the following:

- *Given any parameter spaces $\Omega_i \subset [0,1]$, $i \in \{B, S\}$, satisfactory mechanisms exist if and only if such mechanisms exist when it is common knowledge that each agent's type process is the most persistent one possible (i.e., that $\theta_i = \sup \Omega_i$ for $i \in \{B, S\}$).*

- *Any $\Omega' \subset \Omega$ induces a surplus-neutral truncation of parameters, which makes inequality (1) easier to satisfy (strictly so, if $\sup \Omega'_i < \sup \Omega_i$ for some $i \in \{B, S\}$).*

5.3 Trading Frequency

Let $T = \infty$ throughout this subsection, and denote by $\Delta > 0$ (real) time between periods. A natural modeling strategy that allows varying Δ is to fix an underlying pair of independent continuous-time processes, and think of the discrete-time processes V and C generated by sampling the continuous-time processes at Δ intervals. This implies that increasing frequency of interaction by reducing Δ has two realistic effects: It reduces discounting between interactions and increases correlation between an agent's types in adjacent periods. Based on our remarks on persistence, the former tends to be favorable for efficient bargaining whereas the latter tends to be detrimental. Depending on the parameters, either may dominate:

Example 9 *Consider the renewal model with known parameters considered in Example 6. Suppose that the type renewals of agent i are generated by Poisson arrivals at rate λ_i for $i \in \{B, S\}$. Then $\rho_i = 1 - e^{-\lambda_i \Delta}$, which implies that in the continuous-time limit (i.e., as $\Delta \to 0$), the necessary and sufficient condition (2) derived in Example 6 simplifies further to*

$$\sqrt{\lambda_B \lambda_S} \geq r,$$

---

[15]Note that the processes V and C are of the same form as the family $\{Z^\alpha\}$ discussed after Definition 7, only now the parameter α is private information.

where r is the continuous-time discount rate. That is, frequent interaction facilitates efficient, budget-balanced, and unsubsidized trade if and only if the geometric average of the renewal rates is higher than the discount rate.[16]

Note that, in contrast, taking $\delta = e^{-r\Delta} \to 1$ by sending $r \to 0$ always leads to inequality (2) being satisfied provided that neither agent has a fully persistent type (i.e., that $\rho_i > 0$, or equivalently, that $\lambda_i > 0$ for both i). Hence the two limits lead to qualitatively different results if we start from a situation with $0 < \sqrt{\lambda_B \lambda_S} < r$. This suggests that the efficiency results for high δ in repeated adverse selection models in the literature (e.g., in Athey and Miller, 2007, Athey and Segal, 2012, Escobar and Toikka, 2012, or Fudenberg, Levine, and Maskin, 1994) should be interpreted literally as low discounting results, and the findings will in general be different for the frequent-interaction case. Indeed, the contrast is particularly stark in case of stationary Gaussian types (see Proposition 6 below), but obtaining this result requires second-best analysis, which we turn to next.

6   On Second Best

So far we have restricted attention to studying whether there exists an incentive compatible mechanism with the efficient, or *first-best*, allocation rule, and which satisfies some form of individual rationality and budget balance. When the necessary and sufficient condition (1) for this fails, it is natural to look for a *second-best* mechanism, which we take to mean a mechanism that maximizes the expected gains from trade subject to incentive compatibility, individual rationality, and budget balance.

Unfortunately, the second-best problem appears highly intractable for the general model. First of all, optimal mechanism design is notoriously difficult with multi-dimensional types even in a static setting. Hence we are lead to consider environments where new private information is one-dimensional in each period. However, even with this restriction, the existing methods for characterizing optimal dynamic mechanisms require introducing additional structure (see PST, as well as Battaglini and Lamba, 2012). Given that a second-best result as general as Proposition 1 is thus out of reach, we focus here on two special cases, which allow us to illustrate the arguments employed in the second-best analysis, and deliver clean results.

---

[16]More precisely, it is straightforward to verify that reducing Δ always helps in the sense of making inequality (2) easier to satisfy, but the inequality is satisfied in the limit iff $\sqrt{\lambda_B \lambda_S} \geq r$. We conjecture that this "comparative static" with respect to Δ extends to the generalized renewal processes considered in Example 8.

6.1 Limits of Second Best under Stationary Gaussian Types

Consider the linear AR(1) processes of Example 3 with known parameters. Assume right away that the horizon is infinite, and that the processes are generated by sampling independent Ornstein-Uhlenbeck processes $\tilde{V}$ and $\tilde{C}$ defined by the stochastic differential equations

$$
\begin{aligned}
d\tilde{V}_\tau &= -\alpha_B(\tilde{V}_\tau - m_B)d\tau + \sigma_B dW_\tau^B, \\
d\tilde{C}_\tau &= -\alpha_S(\tilde{C}_\tau - m_S)d\tau + \sigma_S dW_\tau^S,
\end{aligned}
$$

where $\tau \geq 0$ denotes real time, W^i are independent copies of standard one-dimensional Brownian motion, and $\alpha_i > 0$, m_i, and $\sigma_i > 0$ are parameters. This induces discrete-time processes

$$
\begin{aligned}
v_t &= \gamma_B v_{t-1} + (1 - \gamma_B)m_B + \varepsilon_{B,t}, \\
c_t &= \gamma_S c_{t-1} + (1 - \gamma_S)m_S + \varepsilon_{S,t},
\end{aligned}
$$

where $\gamma_i = e^{-\alpha_i \Delta}$, and the distribution of the independent shocks $\varepsilon_{i,t}$ is $N\left(0, (1 - \gamma_i^2)\frac{\sigma_i^2}{2\alpha_i}\right)$. We assume that $\tilde{V}_0$ and $\tilde{C}_0$ (and hence V_0 and C_0) are distributed according to the invariant distributions $N(m_i, \frac{\sigma_i^2}{2\alpha_i})$. *Stationary Gaussian types* are thus defined by a collection $\{r, \Delta, (\alpha_i, \sigma_i, m_i)_{i \in \{B,S\}}\}$, where $r > 0$, $\Delta > 0$, $\alpha_i > 0$, $\sigma_i > 0$, and $m_i \in \mathbb{R}$ (for $i \in \{B, S\}$). Such processes are smooth, and hence Proposition 1 can be applied to obtain the following striking result, which is the starting point for our second-best analysis:[17]

Proposition 4 *If μ is an E, IC, IR_0 mechanism in an environment with stationary Gaussian types, then μ does not satisfy BB_0. In particular,*

$$
(1 - \delta)\mathbb{E}^\mu \left[\sum_{t=0}^\infty \delta^t \left(P_{B,t} + P_{S,t} \right) \right] = \mathbb{E}[Y].
$$

That is, efficient, individually rational trade requires a subsidy equal to the expected first-best gains from trade for any choice of the mean reversion and discount rates, the long-run means, or the length of the period! We relegate the proof into the Appendix along all other proofs omitted from this section, but the argument is simple: By Proposition 1 and Remark 2, it suffices to establish that $\inf_{v_0} \mathbb{E}[Y \mid v_0] = \inf_{c_0} \mathbb{E}[Y \mid c_0] = 0$. This in turn follows from the unbounded supports. Namely, given any $\{r, \Delta, (\alpha_i, \sigma_i, m_i)_{i \in \{B,S\}}\}$ and $\varepsilon > 0$, we may choose v_0 small enough and c_0 large enough so that conditional on v_0 or c_0, convergence to the

---

[17]The result extends a priori to the case where some or all of the parameters $(\alpha_i, \sigma_i, m_i)$ are private information.

invariant distribution takes arbitrarily long, which results in the expected gains from trade being less than ε despite the fact that α_i and $m_B - m_S$ may be large and r may be small.[18]

As the first step towards second-best mechanisms, we extend a part of the characterization of static IC, IR, and BB trading mechanisms by Myerson and Satterthwaite (1983).

Lemma 1 *Consider an environment with stationary Gaussian types. If μ is an IC, IR_0, and BB_0 mechanism, then*

$$\mathbb{E}^\mu \left[\sum_{t=0}^\infty \delta^t X_t \left(V_t - \frac{1 - F_0(V_0)}{f_0(V_0)} \gamma_B^t - C_t - \frac{G_0(C_0)}{g_0(C_0)} \gamma_S^t \right) \right] \geq 0. \tag{3}$$

The left-hand side of (3) is the expected dynamic virtual surplus, which is shown to be non-negative by an argument exactly analogous to the static case. Note that the result only invokes the weakest versions of the three conditions.

By Lemma 1, the expected gains from trade in any IC, IR_0, and BB_0 mechanism are bounded from above by

$$y^{**} := \sup \left\{ (1 - \delta)\mathbb{E}^\mu \left[\sum_{t=0}^\infty \delta^t X_t \left(V_t - C_t \right) \right] \middle| \mu \text{ satisfies (3)} \right\}, \tag{4}$$

where both the objective function and the constraint depend on μ only through the allocation rule. We say that μ is a *second-best mechanism* if μ is IC, IR_0, and BB_0, and the gains from trade under μ achieve y^{**}, which we refer to correspondingly as the *expected second-best gains from trade*. We show next that such mechanisms exist, and that we can ask them to satisfy stronger conditions at no cost:

Proposition 5 *The following hold in every environment with stationary Gaussian types:*

1. *There exist a public second-best mechanism that is PIC, IR_0, and BB.*

2. *There exist a public second-best mechanism that is EPIC, IR, and BB_0.*

Furthermore, the allocation rule in any second-best mechanism is almost surely given by the allocation rule x^{**} defined by setting (for all t) $x_t^{**} = 1$ if and only if

$$v_t - c_t \geq \frac{\lambda}{1 + \lambda} \left(\frac{1 - F_0(v_0)}{f_0(v_0)} \gamma_B^t + \frac{G_0(c_0)}{g_0(c_0)} \gamma_S^t \right), \tag{5}$$

---

[18]It can be shown that if the distributions of V_0 and C_0 are truncated, respectively, from below and above, then inequality (1) is satisfied for any r small enough. Thus Proposition 4 relies on the noncompact supports, which are a consequence of our assumption that the processes be started from the stationary distribution.

where $\lambda > 0$ is the Lagrange multiplier on the constraint (3) in the optimization problem (4).

To sketch the proof, we observe first that the supremum in (4) is achieved by some mechanism μ since both the expected gains from trade and the expected dynamic virtual surplus are continuous in the allocation rule, and the set of allocation rules is compact. As both the objective function and the constraint in (4) are linear, the allocation rule in the mechanism μ must almost surely be equal to the allocation rule x^{**} defined in the proposition for some Lagrange multiplier $\lambda > 0$, which is strictly positive by Proposition 4. Note that x^{**} is "strongly monotone" as for all $s \leq t$, increasing v_s or decreasing c_s weakly increases x_t^{**}. Since the processes are stochastically monotone, Corollary 2 of PST implies that there exists a transfer rule p^{**} such that the public mechanism $\mu^{**} := (x^{**}, p^{**})$ is PIC. The other properties are established using arguments resembling the proofs of Proposition 1 and Lemma 1.

Analogously to the static case, trade occurs in a second-best mechanism only if the buyer's value exceeds the seller's cost by a sufficient margin, which in the current setting depends on the agents' (reported) first-period types. Recalling that $\gamma_i = e^{-\alpha_i \Delta} < 1$, we see by inspection of (5) that this margin converges to zero as $t \to 0$, and hence distortions vanish over time. As we discuss further in the next subsection, this result is best viewed as a consequence of the fact that the impulse response of each agents' AR(1) process, which is given by γ_i^t, decays over time.

Inequality (5) features the Lagrange multiplier λ, and hence it is not immediately obvious how y^{**} varies in relation to $\mathbb{E}[Y]$ as we vary the parameters. However, it is possible to use approximation arguments to show that the findings for the limits $r \to 0$ and $\Delta \to 0$ are qualitatively different. In order to state the result, let $y^* := \mathbb{E}[(V_0 - C_0)^+] = \mathbb{E}[Y]$, where the second equality follows because initial types are drawn from the stationary distribution.

Proposition 6 *Let $y^{**}(r, \Delta)$ denote the expected second-best gains from trade in an environment with stationary Gaussian types given discount rate $r > 0$ and period length $\Delta > 0$. Then*

1. *for all $\Delta > 0$, $\lim_{r \to 0} y^{**}(r, \Delta) = y^*$,*

2. *for all $r > 0$, $\lim_{\Delta \to 0} y^{**}(r, \Delta) < y^*$.*

The first part gives a limit efficiency result for patient agents, which is somewhat surprising given the negative result of Proposition 4. The reason for the seeming discrepancy is that the large information rents under the first-best rule are in part due to the types far in the tails of the distributions, which contribute little to the expected gains from trade. In particular, a mechanism where trade breaks down permanently given a very low value of v_0 or a very high

value of c_0, but where trading is efficient otherwise, results in a small loss in surplus but yields a large reduction in information rents. As $r \to 0$, we may move the truncations arbitrarily far out in the tails to obtain an approximately efficient mechanism.[19]

In contrast, for the frequent-interaction limit $\Delta \to 0$, the discount rate r is held constant, and hence the lack of uniform bounds on the convergence of the t-step distributions $F^{(t)}(\cdot \mid v_0)$ across v_0 yields the inefficiency result in the second part of Proposition 6. While the formal arguments differ, the intuition for the finding is similar to that for Proposition 4.

6.2   Second Best with Private Parameters

As a second example, we consider a setting where the initial private information is about a parameter of the type process. Specifically, suppose that V and C are the linear AR(1) processes of Example 3 as above, but now the long-run means m_B and m_S are private information, i.e., $\theta_B = m_B$ and $\theta_S = m_S$. To keep initial private information one-dimensional, we assume that v_0 and c_0 are known (i.e., $\mathcal{V}_0$ and $\mathcal{C}_0$ are singletons). It is straightforward to verify that this environment is smooth.[20]

Lemma 1 immediately extends to the current setting with inequality (3) replaced by

$$\mathbb{E}^\mu \left[\sum_{t=0}^\infty \delta^t X_t \left(V_t - \frac{1 - F_0(\Theta_B)}{f_0(\Theta_B)}(1 - \gamma_B^t) - C_t - \frac{G_0(\Theta_S)}{g_0(\Theta_S)}(1 - \gamma_S^t) \right) \right] \geq 0. \qquad (6)$$

The proof is the same and hence omitted.[21] Second-best mechanisms are then defined analogously to the previous subsection by substituting inequality (6) for constraint (3) in the second-best problem (4).

There are two differences in the expected dynamic virtual surpluses (3) and (6). The first is due to the hazard rates, which simply reflect what is assumed to be private information in period 0. The second difference is due to the impulse responses of the type processes to changes in the agents' initial information. For the linear AR(1) processes they can be derived simply by writing out the moving-average representation of the process. For example, for the

---

[19]This is closely related to the observation in footnote 18.

[20]It is even easier to verify regularity in the sense of Definition A.2: Put $\psi(\theta_B, v, \varepsilon_B) = \gamma_B v + (1 - \gamma_B)\theta_B + \varepsilon_B$. Then $\partial_v \psi = \gamma_B$ and $\nabla_{\theta_B} \psi = 1 - \gamma_B$, which are bounded in the desired sense.

[21]The only difference is in the expressions for the derivatives. For example, formula (11) now becomes

$$U_0'(\theta_B) = (1 - \delta)\mathbb{E}^\mu \left[\sum_{t=0}^\infty \delta^t (1 - \gamma_B^t) X_t \mid \theta_B \right] \quad \text{a.e. } \theta_B,$$

where we have omitted conditioning on the known constant v_0.

buyer we have

$$v_t = \gamma_B^t v_0 + (1 - \gamma_B^t) m_B + \sum_{s=1}^{t} \gamma_B^{t-s} \varepsilon_{B,s}.$$

Thus, when m_B is common knowledge as in the previous subsection, we have the impulse response $\partial_{v_0} v_t = \gamma_B^t$, which decays over time. In contrast, when m_B is private information (and v_0 is known), the impulse response becomes $\partial_{m_B} v_t = 1 - \gamma_B^t$, which is increasing over time whenever $\gamma_B > 0$. Note that $\gamma_B = 0$ corresponds to values being drawn iid from a distribution with a privately known mean—a special case of Example 1—in which case the impulse response is constant over time.

   We refer the reader to PST for the general definition and discussion of impulse responses, and their role in optimal mechanism design. For our purposes, the relevant observation is that the dynamics of the impulse responses translate to dynamics of distortions in the second-best mechanisms:

Proposition 7 *Consider the above environment with Gaussian linear AR(1) processes where the long-run means are private information, and the sets $\mathcal{V}_0$ and $\mathcal{C}_0$ are singletons. Assume that the maps $\theta_B \mapsto \frac{1 - F_0(\theta_B)}{f_0(\theta_B)}$ and $\theta_S \mapsto -\frac{G_0(\theta_S)}{g_0(\theta_S)}$ are non-increasing. Then*

1. *There exist a public second-best mechanism that is PIC, IR_0, and BB.*

2. *There exist a public second-best mechanism that is EPIC, IR, and BB_0.*

Furthermore, the allocation rule in any second-best mechanism is almost surely given by the allocation rule x^{***} defined by setting (for all t) $x_t^{***} = 1$ if and only if

$$v_t - c_t \geq \frac{\lambda}{1 + \lambda} \left[\frac{1 - F_0(\theta_B)}{f_0(\theta_B)} (1 - \gamma_B^t) + \frac{G_0(\theta_S)}{g_0(\theta_S)} (1 - \gamma_S^t) \right],$$

where $\lambda \geq 0$ is the Lagrange multiplier on inequality (6) in the second-best problem.

   The proof is essentially the same as for Proposition 5, the only subtlety rising from having to establish implementability of an inefficient allocation rule in a non-Markov environment.

   By inspection, whenever inequality (1) is not satisfied so that the second-best mechanism differs from the first best, the buyer's value has to exceed the seller's cost by some margin for trade to take place. In period 0 trade is actually efficient given the commonly known values v_0 and c_0. Thereafter, the margin stays constant in the case of conditionally iid types, but increases over time whenever types are autocorrelated. In contrast, by Proposition 4, distortions decrease over time when the private information is about v_0 and c_0. Heuristically,

27

the difference is due to the fact that distortions are introduced to screen the agents based on their initial information. Hence it is efficient to distort more in periods where types are more sensitive to changes in the agents' initial information, the relevant sense of stochastic dependence being captured by the impulse responses—see PST for further discussion.

The above findings about distortions with privately known means mirror the results by Boleslavsky and Said (2012) who study monopolistic screening of an agent who is privately informed about a parameter of his value process. Exploring the properties of optimal dynamic mechanisms when agents have private information about the parameters of their type processes appears to be an interesting direction for future research.

7   Concluding Remarks

We explore the feasibility of efficient dynamic contracting by deriving, interpreting, and applying a necessary and sufficient condition for the existence of efficient, unsubsidized, and individually rational contracts in a setting where agents may have private information about the evolution of their personal uncertainty. The condition is given by inequality (1), which corresponds to an upper bound on the sensitivity of the expected gains from trade to the agents' initial private information. As illustrated in Section 5, the effects of considerations such as persistence, patience, number and frequency of interactions, or asymmetric information about process parameters on the prospects of efficient contracting can be either simply read off of, or computed from, inequality (1). The result relies on a dynamic payoff-equivalence theorem for settings with multi-dimensional initial information, which may turn out to be useful elsewhere as well.

Our methods apply as such to general dynamic Bayesian collective choice problems in quasi-linear environments. To illustrate this, consider the following class of n-agent problems: In each period $t = 0, 1, \ldots, T$, with $T \in \mathbb{N} \cup \{\infty\}$, a decision is chosen from a measurable space $\mathcal{X}$. If the decision in period t is $x_t \in \mathcal{X}$ and agent $i \in I := \{1, \ldots, n\}$ receives $p_{i,t}$ units of the numeraire, then the resulting flow payoff to agent i is $u_i(x_t, z_{i,t}) + p_{i,t}$ for some (measurable) $u_i : \mathcal{X} \times \mathcal{Z}_i \to \mathbb{R}$. The *type* $z_{i,t}$ is private information of agent i and evolves on the interval $\mathcal{Z}_i \subset \mathbb{R}$ according to a parameterized Markov process Z_i generated by $(H_{i,0}, \{H(\cdot \mid \cdot; \theta_i)\}_{\theta_i \in \Omega_i})$, which is thus of the same form as the value and cost processes V and C in the bilateral trade problem. The agents evaluate streams of flow payoffs according to their discounted average using a common discount factor $\delta \in [0, 1]$ with $\delta < 1$ if $T = \infty$.

We normalize the outside option of each agent to zero, and assume that there exists a decision $x' \in \mathcal{X}$ such that $u_i(x', z_i) = 0$ for all $i \in I$, all $z_i \in \mathcal{Z}_i$. Applications fitting this

framework include repeated versions of allocation problems such as double auctions, sharing a common resource within a team, and the provision of excludable public goods.

Direct mechanisms and their properties are defined for the above dynamic collective choice problems as in Section 3, and the payoff-equivalence property can be defined analogously to Definition 5. Our proof of payoff equivalence (Proposition A.1) extends verbatim to show that a sufficient condition for the latter is that (1) the type processes are smooth in the sense of Definition 6 (or, more generally, regular as in Definition A.2), and (2) for each agent $i \in I$ and every decision $x \in \mathcal{X}$, $u_i(x, \cdot)$ is differentiable and the family $\{u_i(x, \cdot)\}_{x \in \mathcal{X}}$ is equi-Lipschitz.

Denote the *first-best social surplus* by

$$S := \frac{1-\delta}{1-\delta^{T+1}} \sum_{t=0}^{T} \delta^t \sum_{i=1}^{n} u_i(\chi^*(Z_{1,t}, \ldots Z_{n,t}), Z_{i,t}),$$

where χ^* is a static first-best allocation rule.[22] We then have the following generalization of Proposition 1, which provides a dynamic version of Williams' (1999) characterization.[23]

Proposition 1′ *Suppose that the dynamic Bayesian collective choice problem defined above has the payoff-equivalence property. Then the following are equivalent:*

1. *The first-best social surplus, S, satisfies*

$$\sum_{i=1}^{n} \inf_{\theta_i, z_{i,0}} \mathbb{E}[S \mid \theta_i, z_{i,0}] \geq (n-1)\mathbb{E}[S]. \tag{7}$$

2. *There exists a blind mechanism that is E, IC, IR_0, and BB_0.*

3. *There exists a public mechanism that is E, PIC, IR_0, and BB.*

4. *There exists a public mechanism that is E, EPIC, IR, and BB_0.*

Proposition 1' permits an analysis analogous to Section 5 for any problem in the above class. For example, it allows exploring how the performance of markets organized as double auctions is affected by trading frequency, persistence of valuations, or asymmetric information about the processes generating the valuations.

---

[22]I.e., $\chi^*(z_1, \ldots, z_n) \in \arg\max_{x \in \mathcal{X}} \sum_{i=1}^{n} u_i(x, z_i)$ for all type profiles $(z_1, \ldots, z_n) \in \mathcal{Z}_1 \times \cdots \times \mathcal{Z}_n$.

[23]Define the mechanism $\mu = (x^*, p)$ in the beginning of the proof of Proposition 1 to consist of the repetition of the static Groves' scheme (instead of the static n-agent Pivot mechanism) so that $p_{i,t} = \sum_{j \neq i} u_j(x_t^*, z_{i,t})$. Then μ is E, EPIC, and IR with budget-deficit equal to $(n-1)\mathbb{E}[S]$. The rest of the proof now goes through with the obvious adjustments.

Another application of Proposition 1' comes from repeated Bayesian games. Namely, while the above collective choice problems assume transferable utility, inequality (7) obviously remains a necessary condition for the existence of an equilibrium that maximizes the sum of the players' payoffs even if utility is non-transferable. This observation can be used, for example, to put bounds on firms' ability to collude when their cost structures are private information, thus providing a way of extending the results of Miller (2012), who shows for the case of iid costs that first-best collusion (or E) is unattainable under ex post incentives and ex post budget balance (or EPIC and BB).

Remark 3 *Proposition 1' is by no means the most general possible statement. Indeed, we assume payoffs to be additively separable across time and evolution of types to be independent of decisions for the ease of exposition, and because for such environments our proof of payoff equivalence goes through verbatim. But since payoff equivalence is simply an assumption for the result, Proposition 1' immediately extends—with first-best social surplus S appropriately defined—to the general environments studied by Athey and Segal (2012) and Bergemann and Välimäki (2010) as we may take their efficient dynamic mechanisms as the starting point in the proof. Sufficient conditions for the payoff-equivalence property to hold can then be obtained as in this paper by applying or extending the results of PST.*

Finally, we have abstracted from institutional detail throughout the paper in order to focus on the informational problems inherent in dynamic contracting. A natural question for future research is to investigate to what extent the results are affected by the introduction of additional concerns. For example, as noted above, inequality (1) is in general not enough to simultaneously guarantee ex post budget balance and individual rationality in every period. Hence, while (1) remains a necessary condition for any environment, it need not be sufficient in some institutional settings. An example is provided by relational contracting with third-party financing if only bounded credit lines are available. (See Athey and Miller, 2007, for an exploration of these issues in the iid case, or Athey and Segal, 2007, 2012, for positive limit results in settings with serial dependence.) The design of second-best mechanisms when initial information is multi-dimensional is another natural, but challenging, next step.

A Appendix

A.1 Sufficient Conditions for Payoff Equivalence

We introduce regular environments and show that they have the payoff-equivalence property by extending the "first-order approach" of PST to our setting, where the agents' initial

information is multi-dimensional.

The following definition adapts the concept of a state representation from PST to our environment (see also Eső and Szentes, 2007). For definiteness we use the notation for the buyer's type process throughout this section; the seller's process is treated analogously.

Definition A.1 *A* state representation *of the process V is a triple $(\mathcal{E}, Q, \psi)$, where $\mathcal{E}$ is a measurable space, Q is a probability distribution on $\mathcal{E}$, and $\psi : \Omega_B \times \mathcal{V} \times \mathcal{E} \to \mathcal{V}$ is a (measurable) function such that, for all $(\theta_B, v) \in \Omega_B \times \mathcal{V}$, $\psi(\theta_B, v, \cdot) : \mathcal{E} \to \mathcal{V}$ is a random variable with distribution $F(\cdot \mid v; \theta_B)$.*

Given a state representation $(\mathcal{E}, Q, \psi)$, we can think of the buyer's values as being generated as follows: Draw the initial information (θ_0, v_0) according to F_0, and draw a sequence $(\varepsilon_t)_{t=1}^T \in \mathcal{E}^T$ of "independent shocks" according to the product measure $\Pi_{t=1}^T Q$. Values for periods $t > 0$ are then obtained by iterating $v_t = \psi(\theta_B, v_{t-1}, \varepsilon_t)$. Note that this amounts to simply extending the standard construction of a Markov chain in terms of iid random variables (see, e.g., Williams, 1991, p. 209) to a mixture over the parameterized collection of Markov chains $\{\langle v_0, F(\cdot \mid v; \theta_B)\rangle\}_{(\theta_B, v_0) \in \Omega_B \times \mathcal{V}_0}$, and hence a state representation exists. For example, the *canonical representation* obtains by taking $\mathcal{E} = [0, 1]$ and $\psi(\theta_B, v, \cdot) = F^{-1}(\cdot \mid v; \theta_B)$ for all (θ_B, v), and letting Q to be the uniform distribution.

Definition A.2 *The process V is* regular *if there exists a state representation $(\mathcal{E}, Q, \psi)$ of V and constants $b < \frac{1}{\delta}$, $d < \infty$ such that, for all $(\theta_B, v, \varepsilon) \in \Omega_B \times \mathcal{V} \times \mathcal{E}$, $\psi(\theta_B, v, \varepsilon)$ is a differentiable function of (θ_B, v) satisfying*

   1. $\|\nabla_{\theta_B}\psi\| \leq d$, and

   2. $|\partial_v \psi| \leq b$.

If T is finite, it suffices that the constant b be finite.

   The environment is regular *if the processes V and C are regular.*

An easy-to-check sufficient condition for regularity is provided by the notion of smoothness introduced in Definition 6 in the main text:

Lemma A.1 *Every smooth environment is regular.*

Proof. It suffices to show that a smooth process is a regular process. We do this by showing that the canonical representation of a smooth process satisfies the conditions of Definition A.2. For definiteness, suppose the buyer's process V is smooth, and consider the canonical

representation $([0,1], Q, F^{-1})$, where Q is the uniform distribution on $[0,1]$. Smoothness implies that the kernel $F(\cdot \mid \cdot; \cdot)$ is continuously differentiable, and for all (v, θ_B), there is a density $f(\cdot \mid v; \theta_B)$ strictly positive on $\mathcal{V}$. Therefore, for all $(\varepsilon, v, \theta_B)$ and $v' := F^{-1}(\varepsilon \mid v; \theta_B)$, the Implicit Function Theorem implies

$$\partial_v F^{-1}(\varepsilon \mid v; \theta_B) = -\frac{\partial_v F(v' \mid v; \theta_B)}{f(v' \mid v; \theta_B)},$$

where the right-hand side is bounded by some $b < \frac{1}{\delta}$ in absolute value by smoothness. Thus $([0,1], Q, F^{-1})$ satisfies the second condition in Definition A.2. Similarly, we have

$$\nabla_{\theta_B} F^{-1}(\varepsilon \mid v; \theta_B) = -\frac{\nabla_{\theta_B} F(v' \mid v; \theta_B)}{f(v' \mid v; \theta_B)},$$

where the right-hand side is bounded in the norm by some constant d by smoothness. Thus $([0,1], Q, F^{-1})$ satisfies also the first condition of Definition A.2. ∎

An example of a non-smooth regular environment is provided by the renewal model:

Lemma A.2 *Consider the renewal model of Example 2. Suppose that the inverses of the initial conditional distributions, $F_0^{-1}(\cdot \mid \theta_B)$ and $G_0^{-1}(\cdot \mid \theta_S)$, are differentiable functions of θ_B and θ_S, respectively, with uniformly bounded gradients (i.e., $\exists d < \infty : \left\| \nabla_{\theta_B} F_0^{-1} \right\| \vee \left\| \nabla_{\theta_S} G_0^{-1} \right\| \leq d$). Then the environment is regular.*

Note that the assumed differentiability of the inverses is satisfied, e.g., if the environment is with known parameters, or if $F_0(\cdot \mid \theta_B)$ and $G_0(\cdot \mid \theta_S)$ have strictly positive densities given any $\theta \in \Omega$, and the ratios $\frac{\left\| \nabla_{\theta_B} F_0(v|\theta_B) \right\|}{f_0(v|\theta_B)}$ and $\frac{\left\| \nabla_{\theta_S} G_0(c|\theta_S) \right\|}{g_0(c|\theta_S)}$ are bounded uniformly in $(\theta, v, c) \in \Omega \times \mathcal{V} \times \mathcal{C}$.

Proof. Consider the buyer's process V. Define the state representation $(\mathcal{E}, Q, \psi)$ as follows: Put $\mathcal{E} = [0,1] \times \{0,1\}$, and let $Q = Q_1 \times Q_2$, where Q_1 is the uniform distribution on $[0,1]$, and the distribution Q_2 on $\{0,1\}$ is defined by $\Pr\{\varepsilon_2 = 1\} = \gamma_B$. Define ψ by setting

$$\psi(\theta_B, v, \varepsilon) = \varepsilon_2 v + (1 - \varepsilon_2) F_0^{-1}(\varepsilon_1 \mid \theta_B).$$

Verifying that this indeed defines a state representation is straightforward. Moreover, $\psi(\theta_B, v, \varepsilon)$ is clearly a differentiable function of (θ_B, v), and we have $\left\| \nabla_{\theta_B} \psi \right\| \leq \left\| \nabla_{\theta_B} F_0^{-1} \right\| \leq d$ for some $d < \infty$, and $|\partial_v \psi| = \varepsilon_2 \leq 1 < \frac{1}{\delta}$. Therefore, the buyer's process V is regular. The seller's process C is treated similarly. ∎

We may now establish our payoff-equivalence result:

Proposition A.1 *Every regular environment has the payoff-equivalence property.*

In the special case where each agent's initial private information is one-dimensional, Proposition A.1 follows from the results of PST. The extension to the multi-dimensional case presented here is novel. The proof, which combines the standard payoff-equivalence argument from static multi-dimensional models can be sketched as follows: Fix an IC mechanism, and let α be a smooth path (e.g., a line segment) between two initial types of the buyer, say, (v_0^0, θ_B^0) and (v_0^1, θ_B^1). Consider an auxiliary problem where a buyer, whose true initial type is in α, is restricted to report a type in α in period 0 (but may report any v_t in periods $t > 0$). In this problem, the agent's initial type is one-dimensional. Since the mechanism is IC, a truthful strategy is still optimal for the buyer and results in the same payoff as in the original model. Furthermore, if the environment is regular, then this auxiliary problem satisfies the assumptions of Theorem 1 of PST. This implies that the payoff difference between any two types in α is pinned down by the allocation rule alone. As (v_0^0, θ_B^0) and (v_0^1, θ_B^1) were arbitrary, the result follows.

Proof. Fix some IC mechanism $\mu = (x, p)$, and two initial buyer types (θ_B^0, v_0^0), $(\theta_B^1, v_0^1) \in \Omega_B \times \mathcal{V}_0$ with $(\theta_B^0, v_0^0) \neq (\theta_B^1, v_0^1)$. (The seller is treated analogously.) The proposition is proven by establishing that the equilibrium-payoff difference $U^\mu(\theta_B^1, v_0^1) - U^\mu(\theta_B^0, v_0^0)$ depends only on the allocation rule x if V is regular.[24]

Fix a smooth path $\alpha : [0, 1] \to \Omega_B \times \mathcal{V}_0$ from (θ_B^0, v_0^0) to (θ_B^1, v_0^1), i.e., a continuously differentiable map $\lambda \mapsto \alpha(\lambda) = (\alpha_\theta(\lambda), \alpha_v(\lambda)) \in \Omega_B \times \mathcal{V}_0$, where $\lambda \in [0, 1]$, such that $\alpha(0) = (\theta_B^0, v_0^0)$ and $\alpha(1) = (\theta_B^1, v_0^1)$. (Such paths exist by the convexity of $\Omega_B \times \mathcal{V}_0$.) In the remainder of the proof we restrict attention to the path α and hence (abusing terminology) refer to λ as the buyer's initial type.

Given a strategy profile σ and an initial type $\lambda \in [0, 1]$, define $W^\sigma(\lambda) := U^{\mu, \sigma}(\alpha(\lambda))$ with $W(\lambda) := W^{\sigma^*}(\lambda)$. By IC, we have for all $\lambda \in [0, 1]$,

$$
\begin{aligned}
W(\lambda) &= \max_{\sigma_B} W^{(\sigma_B, \sigma_S^*)}(\lambda) \\
&= \max_{\sigma_B} \{ W^{(\sigma_B, \sigma_S^*)}(\lambda) : \sigma_{B,0}(\alpha(\lambda)) = \alpha(\lambda') \text{ for some } \lambda' \in [0, 1] \},
\end{aligned}
$$

where the second equality follows since the optimal truthful report remains feasible. By inspection of the second line, we can view W as the value function to a family of dynamic optimization problems parameterized by the initial type $\lambda \in [0, 1]$, where in period 0 the buyer with true initial type λ is restricted to report some initial type $\lambda' \in [0, 1]$ (and makes a report $\hat{v}_t \in \mathcal{V}$ in periods $t > 0$, as usual). As the initial type λ is one-dimensional, this auxiliary problem is amenable to the first-order approach of PST. In particular, their Theo-

---

[24]To see this, fix an arbitrary $a \in \Omega_B \times \mathcal{V}_0$. Given any IC mechanisms μ, η with the same allocation rule x, let $k := U^\mu(a) - U^\eta(a)$. Then for all $b \in \Omega_B \times \mathcal{V}_0$, $U^\mu(b) - U^\mu(a) = U^\eta(b) - U^\eta(a)$, or $U^\mu(b) - U^\eta(b) = k$.

rem 1 implies that under certain conditions, W is Lipschitz-continuous with a derivative W' independent of p, so the independence of $U^\mu(\theta_B^1, v_0^1) - U^\mu(\theta_B^0, v_0^0)$ of p then follows from the Fundamental Theorem of Calculus by observing that

$$U^\mu(\theta_B^1, v_0^1) - U^\mu(\theta_B^0, v_0^0) = W(1) - W(0) = \int_0^1 W'(\lambda)d\lambda.$$

Thus to complete the proof, it suffices to verify that if the process V is regular, then the auxiliary problem satisfies the assumptions of Theorem 1 of PST.

The buyer's type in the auxiliary problem is given by a sequence $(\lambda, v_1, v_2, \ldots, v_T)$, and his payoff takes the time-separable form

$$\frac{1-\delta}{1-\delta^{T+1}}[x_0\alpha_v(\lambda) + p_{B,0} + \sum_{t=1}^{T} \delta^t(x_t v_t + p_{B,t})].$$

Thus the two conditions on the utility function (U-D and U-ELC), which require differentiability and equi-Lipschitz continuity in types in the appropriate sense, are clearly satisfied (see the example after Condition 2 in PST). Similarly, the condition requiring the expected discounted type to be finite conditional on the initial type (F-BE$_0$) follows immediately from our assumption that $\mathbb{E}[\sum_{t=0}^{T} \delta^t|V_t| \mid \theta_B, v_0]$ be finite for all $(\theta_B, v_0) \in \Omega_B \times \mathcal{V}_0$.

It remains to show that the type process in the auxiliary problem has "bounded impulse responses" (F-BIR$_0$). Given initial type λ, future types are distributed as follows:

$$\begin{aligned} v_1 &\sim F(\cdot \mid \alpha_v(\lambda); \alpha_\theta(\lambda)), \\ v_t &\sim F(\cdot \mid v_{t-1}; \alpha_\theta(\lambda)) \quad \text{for } t = 2, \ldots, T. \end{aligned} \tag{8}$$

As V is regular, we can take a state representation $(\mathcal{E}, Q, \psi)$ satisfying our Definition A.2, and define the functions $z_1 : [0,1] \times \mathcal{E} \to \mathbb{R}$, and $z_t : \mathcal{V} \times [0,1] \times \mathcal{E} \to \mathbb{R}$, $t = 2, \ldots, T$, by

$$\begin{aligned} z_1(\lambda, \varepsilon) &= \psi(\alpha_\theta(\lambda), \alpha_v(\lambda), \varepsilon), \\ z_t(v, \lambda, \varepsilon) &= \psi(\alpha_\theta(\lambda), v, \varepsilon) \quad \text{for } t = 2, \ldots, T. \end{aligned}$$

We note for future reference that, given the properties of ψ, we have the global existence of the partial derivatives

$$\begin{aligned} \partial_\lambda z_1 &= \nabla_{(\theta,v)}\psi \cdot \alpha', \\ \partial_\lambda z_t &= \nabla_\theta \psi \cdot \alpha_\theta' \quad \text{for } t = 2, \ldots, T, \\ \partial_v z_t &= \partial_v \psi \quad \text{for } t = 2, \ldots, T. \end{aligned}$$

Furthermore, given any constants b and d that satisfy Definition A.2, we may put $K := (d+b)\max_\lambda \|\alpha'(\lambda)\|$ to obtain the bounds

$$
\begin{aligned}
|\partial_\lambda z_t| &\leq K &&\text{for } t = 1,\dots,T, \\
|\partial_v z_t| &\leq b &&\text{for } t = 2,\dots,T.
\end{aligned}
$$

By construction, the collection $(\mathcal{E}, Q, z)$, where $z := (z_t)_{t=1}^T$, defines an S-representation of the kernels (8) in the sense of Definition 2 of PST.[25] By inspection of their Condition 4 and their equation (9), it suffices to show the finiteness of the sum $\sum_{t=0}^T \delta^t |I_t|$, where $I_0 \equiv 1$ and

$$
I_t := \sum_{\tau=1}^t \partial_\lambda z_\tau \prod_{s=\tau+1}^t \partial_v z_s \quad \text{for } t = 1,\dots,T.
$$

By the above bounds on the partials of the functions $(z_t)_{t=1}^T$, we have for all t,

$$
|I_t| \leq \sum_{\tau=1}^t K \prod_{s=\tau+1}^t b = K \sum_{\tau=1}^t b^{t-\tau} \leq K t b^t,
$$

where the last inequality follows, since we may assume without loss that $b \geq 1$. Therefore,

$$
\sum_{t=0}^T \delta^t |I_t| \leq K \sum_{t=0}^T t(\delta b)^t.
$$

Because either $b < \frac{1}{\delta}$ or $T < \infty$, the last sum is finite, which implies the result. ∎

Finally, we show that the environment in Example 8 is regular and hence has the payoff-equivalence property.

Lemma A.3 *The environment in Example 8 is regular.*

Proof. By symmetry, it suffices to consider the buyer's process V. By the properties of the twice-reflected Brownian motion (see, e.g., Harrison, 1985), the kernel $\bar{F}(v' \mid v)$ is a differentiable function of (v', v) with a density $\bar{f}(v' \mid v)$ bounded away from zero uniformly in (v', v) and with a uniformly bounded partial derivative $\partial_v \bar{F}$. Therefore, there exists $\varepsilon > 0$ and $K < \infty$ such that for all $(v', v, \theta_B) \in [0, 1]^3$,

$$
\frac{|\partial_v F(v' \mid v; \theta_B)|}{f(v' \mid v; \theta_B)} = \frac{|\partial_v \bar{F}(v' \mid v)|}{\theta_B \bar{f}(v' \mid v) + 1 - \theta_B} \leq \frac{K}{\theta_B \varepsilon + 1 - \theta_B} \leq \frac{K}{\varepsilon}.
$$

---

[25]Note that we can put $\mathcal{E}_t = \mathcal{E}$ and $G_t = Q$ for all t.

Similarly, we have

$$\frac{|\nabla_{\theta_B} F(v' \mid v; \theta_B)|}{f(v' \mid v; \theta_B)} = \frac{|\bar{F}(v' \mid v) - v'|}{\theta_B \bar{f}(v' \mid v) + 1 - \theta_B} \leq \frac{1}{\varepsilon}.$$

Thus the process V is smooth, save for the constant $\frac{K}{\varepsilon}$ being possibly greater than $\frac{1}{\delta}$. However, as $T < \infty$, V is regular by inspection of Definition A.2 and the proof of Lemma A.1. $\blacksquare$

A.2 Omitted Proofs

Proof of Proposition 1. The implication $3 \vee 4 \Rightarrow 2$ follows immediately from the Definitions 1–4 (see Remark 1). Hence it suffices to show $1 \Rightarrow 3 \wedge 4$, and $2 \Rightarrow 1$.

We establish $1 \Rightarrow 4$ by showing that if (1) holds, there is a simple mechanism that has the desired properties. Consider first the public mechanism $\mu = (x^*, p)$, which consists of running the static Pivot mechanism in every period. I.e., the allocation rule x^* and the payment rule p are defined by setting, for all t, $x_t^* = \mathbf{1}_{\{v_t \geq c_t\}}$, $p_{B,t} = -\mathbf{1}_{\{v_t \geq c_t\}} c_t$, and $p_{S,t} = \mathbf{1}_{\{v_t \geq c_t\}} v_t$.[26] By construction, μ is E, and each player's payoff equals the first-best gains from trade $(V_t - C_t)^+$ in each period t. Thus μ is IR, and period-0 payoffs are given by $U_0^\mu(\theta_B, v_0) = \mathbb{E}[Y \mid \theta_B, v_0]$ and $\Pi_0^\mu(\theta_S, c_0) = \mathbb{E}[Y \mid \theta_S, c_0]$ for all $(\theta, v_0, c_0) \in \Omega \times \mathcal{V}_0 \times \mathcal{C}_0$. Furthermore, in each period t, the agents' reports only affect the current allocation and transfers, and thus μ is EPIC by the usual static argument. Finally, note that μ runs an expected budget deficit equal to the expected gains from trade, or $\frac{1-\delta}{1-\delta^{T+1}} \mathbb{E}^\mu \left[\sum_{t=0}^{T} \delta^t (P_{B,t} + P_{S,t}) \right] = \mathbb{E}[Y]$.

In order to recover the budget deficit, we add "participation fees" to the mechanism μ by constructing a new transfer rule p^* from p by setting $p_t^* = p_t$ for all $t > 0$, and defining the new period-0 transfers by

$$p_{B,0}^* := p_{B,0} - \frac{1 - \delta^{T+1}}{1 - \delta} \inf_{\theta_B', v_0'} \mathbb{E}[Y \mid \theta_B', v_0'],$$

and

$$p_{S,0}^* := p_{S,0} - \frac{1 - \delta^{T+1}}{1 - \delta} \inf_{\theta_S', c_0'} \mathbb{E}[Y \mid \theta_S', c_0'].$$

Denote the public mechanism so obtained by $\mu^* = (x^*, p^*)$. As we just subtracted constants,

---

[26] In our environment, this mechanism coincides with the Team mechanism of Athey and Segal (2012) and the Dynamic Pivot mechanism of Bergemann and Välimäki (2010).

μ^* is E and EPIC. It is IR, since periods $t > 0$ are unaffected, and

$$\inf_{\theta_B, v_0} U_0^{\mu^*}(\theta_B, v_0) = \inf_{\theta_B, v_0} \left(U_0^\mu(\theta_B, v_0) - \inf_{\theta'_B, v'_0} \mathbb{E}[Y \mid \theta'_B, v'_0] \right)$$

$$= \inf_{\theta_B, v_0} \mathbb{E}[Y \mid \theta_B, v_0] - \inf_{\theta'_B, v'_0} \mathbb{E}[Y \mid \theta'_B, v'_0] = 0, \tag{9}$$

and similarly for the seller. Note that

$$\frac{1-\delta}{1-\delta^{T+1}} \mathbb{E}^{\mu^*} \left[\sum_{t=0}^T \delta^t (P_{B,t} + P_{S,t}) \right] = \frac{1-\delta}{1-\delta^{T+1}} \mathbb{E}^\mu \left[\sum_{t=0}^T \delta^t (P_{B,t} + P_{S,t}) \right]$$

$$- \left(\inf_{v_0, \theta_B} \mathbb{E}[Y \mid v_0, \theta_B] + \inf_{c_0, \theta_S} \mathbb{E}[Y \mid c_0, \theta_S] \right) \tag{10}$$

$$= \mathbb{E}[Y] - \left(\inf_{v_0, \theta_B} \mathbb{E}[Y \mid v_0, \theta_B] + \inf_{c_0, \theta_S} \mathbb{E}[Y \mid c_0, \theta_S] \right).$$

Thus, if (1) is satisfied, then μ^* is also BB_0, and hence $1 \Rightarrow 4$.

We show then that $2 \Rightarrow 1$. Let $\eta = (x', q)$ be an E, IC, IR_0, and BB_0 blind mechanism, and let $\mu^* = (x^*, p^*)$ be the public "Pivot mechanism with participation fees" constructed above. Since η and μ^* are both E, we have $x' = x^*$, and hence $U_0^\eta = U_0^{\mu^*} + a$ and $\Pi_0^\eta = \Pi_0^{\mu^*} + b$ for some constants $a, b \in \mathbb{R}$ by IC and the payoff-equivalence property. Since η is IR_0, we have

$$0 \le \inf_{v_0, \theta_B} U_0^\eta(v_0, \theta_B) + \inf_{c_0, \theta_S} \Pi_0^\eta(c_0, \theta_S) = \inf_{v_0, \theta_B} U_0^{\mu^*}(v_0, \theta_B) + a + \inf_{c_0, \theta_S} \Pi_0^{\mu^*}(c_0, \theta_S) + b = a + b,$$

where the last equality follows by (9). Thus, $U_0^\eta + \Pi_0^\eta \ge U_0^{\mu^*} + \Pi_0^{\mu^*}$, which in turn implies $\mathbb{E}^\eta \left[\sum_{t=0}^T \delta^t (P_{B,t} + P_{S,t}) \right] \ge \mathbb{E}^{\mu^*} \left[\sum_{t=0}^T \delta^t (P_{B,t} + P_{S,t}) \right]$ as $x' = x^*$. But η is BB_0, and hence

$$0 \ge \frac{1-\delta}{1-\delta^{T+1}} \mathbb{E}^\eta \left[\sum_{t=0}^T \delta^t (P_{B,t} + P_{S,t}) \right]$$

$$\ge \frac{1-\delta}{1-\delta^{T+1}} \mathbb{E}^{\mu^*} \left[\sum_{t=0}^T \delta^t (P_{B,t} + P_{S,t}) \right]$$

$$= \mathbb{E}[Y] - \left(\inf_{v_0, \theta_B} \mathbb{E}[Y \mid v_0, \theta_B] + \inf_{c_0, \theta_S} \mathbb{E}[Y \mid c_0, \theta_S] \right),$$

where the last equality follows by (10). Thus (1) is satisfied.

It remains to establish $1 \Rightarrow 3$. Note that the public mechanism $\mu^* = (x^*, p^*)$ constructed above is PIC. Hence by the "balancing trick" of Athey and Segal (2012, Proposition 2), there exists a public mechanism $\bar{\mu}^* = (x^*, \bar{p}^*)$ that is PIC and BB. Since the allocation rule is

37

unchanged, $\bar{\mu}^*$ is also E, and the payoff-equivalence property implies that $U_0^{\bar{\mu}^*} = U_0^{\mu^*} + a$ and $\Pi_0^{\bar{\mu}^*} = \Pi_0^{\mu^*} + b$ for some constants $a, b \in \mathbb{R}$. By (1) and (10) we then have

$$
\begin{aligned}
a + b &= \mathbb{E}\left[U_0^{\bar{\mu}^*}(V_0, \Theta_B) + \Pi_0^{\bar{\mu}^*}(C_0, \Theta_S)\right] - \mathbb{E}\left[U_0^{\mu^*}(V_0, \Theta_B) + \Pi_0^{\mu^*}(C_0, \Theta_S)\right] \\
&= \mathbb{E}[Y] - \left[2\mathbb{E}[Y] - \left(\inf_{v_0, \theta_B} \mathbb{E}[Y \mid v_0, \theta_B] + \inf_{c_0, \theta_S} \mathbb{E}[Y \mid c_0, \theta_S]\right)\right] \\
&= -\mathbb{E}[Y] + \left(\inf_{v_0, \theta_B} \mathbb{E}[Y \mid v_0, \theta_B] + \inf_{c_0, \theta_S} \mathbb{E}[Y \mid c_0, \theta_S]\right) \geq 0.
\end{aligned}
$$

In particular, this implies that

$$
\inf_{v_0, \theta_B} U_0^{\bar{\mu}^*}(v_0, \theta_B) + \inf_{c_0, \theta_S} \Pi_0^{\bar{\mu}^*}(c_0, \theta_S) = \inf_{v_0, \theta_B} U_0^{\mu^*}(v_0, \theta_B) + a + \inf_{c_0, \theta_S} \Pi_0^{\mu^*}(c_0, \theta_S) + b = a + b \geq 0.
$$

Thus $\bar{\mu}^*$ can be made to satisfy IR_0 by adding a type-independent transfer between the agents in period 0. We conclude that $1 \Rightarrow 3$. ∎

Proof of Proposition 4. Fix $\{r, \Delta, (\alpha_i, \sigma_i, m_i)_{i \in \{B, S\}}\}$. By Proposition 1 and Remark 2, it suffices to show that $\inf_{v_0} \mathbb{E}[Y \mid v_0] = \inf_{c_0} \mathbb{E}[Y \mid c_0] = 0$. By symmetry, it is enough to consider $\inf_{v_0} \mathbb{E}[Y \mid v_0]$. Let $v_0 \leq m_B$, and note that for all t, conditional on $V_0 = v_0$, $V_t - C_t$ is Normally distributed with mean $\gamma_B^t(v_0 - m_B) + m_B - m_S$ and variance $(1 - \gamma_B^{2t})\frac{\sigma_B^2}{2\alpha_B} + \frac{\sigma_S^2}{2\alpha_S}$. Thus,

$$
\mathbb{E}[(V_t - C_t)^+ \mid v_0] \leq \mathbb{E}[Z_t^+],
$$

where Z_t is distributed $N(\gamma_B^t(v_0 - m_B) + m, \sigma^2)$, where $m := m_B - m_S$ and $\sigma^2 := \frac{\sigma_B^2}{2\alpha_B} + \frac{\sigma_S^2}{2\alpha_S}$. Letting ϕ and Φ denote the pdf and the cdf for $N(0, 1)$, we have by standard formulas

$$
\frac{\mathbb{E}[Z_t^+]}{1 - \Phi(\frac{-\gamma_B^t(v_0 - m_B) - m}{\sigma})} = \mathbb{E}[Z_t \mid Z_t \geq 0] = \gamma_B^t(v_0 - m_B) + m + \frac{\phi(\frac{-\gamma_B^t(v_0 - m_B) - m}{\sigma})}{1 - \Phi(\frac{-\gamma_B^t(v_0 - m_B) - m}{\sigma})}\sigma,
$$

and hence,

$$
\begin{aligned}
\mathbb{E}[Y \mid v_0] \leq (1 - \delta) \sum_{t=0}^{\infty} \delta^t \Bigg[&\left(1 - \Phi(\frac{-\gamma_B^t(v_0 - m_B) - m}{\sigma})\right)(\gamma_B^t(v_0 - m_B) + m) \\
&+ \phi(\frac{-\gamma_B^t(v_0 - m_B) - m}{\sigma})\sigma \Bigg].
\end{aligned}
$$

As $v_0 \leq m_B$, the summand on the right is dominated by $\delta^t(m_B + m + \sigma)$. Thus, taking the limit $v_0 \to -\infty$, we may pass the limit through the sum by the dominated convergence theorem to find that the sum converges to zero. Therefore, $0 \leq \inf_{v_0} \mathbb{E}[Y \mid v_0] \leq 0$. ∎

Proof of Lemma 1. Suppose that μ is an IC, BB_0, IR_0 mechanism. Since the environment is smooth and μ is IC, the dynamic envelope theorem of PST (Theorem 1) implies that $U_0(v_0) := U_0^\mu(v_0)$ is Lipschitz continuous in v_0 with derivative

$$U_0'(v_0) = (1-\delta)\mathbb{E}^\mu\left[\sum_{t=0}^\infty (\delta\gamma_B)^t X_t \mid v_0\right] \quad \text{a.e. } v_0, \tag{11}$$

where we have omitted θ_B as the environment is with known parameters. By inspection, U_0 is nondecreasing. Thus $\lim_{v\to-\infty} U_0(v)$ is well-defined, and for all v_0,

$$\lim_{v\to-\infty} U_0(v) = U_0(v_0) - \lim_{v\to-\infty}\int_v^{v_0} U_0'(r)dr = U_0(v_0) - \int_{-\infty}^{v_0} U_0'(r)dr,$$

where the last equality follows by the Monotone Convergence Theorem since $U_0' \geq 0$. An analogous result obtains for the seller's equilibrium payoff Π_0, which is seen to be nonincreasing. Therefore, we have

$$\begin{aligned}
0 &\leq \lim_{v\to-\infty} U_0(v) + \lim_{c\to\infty} \Pi_0(c) \\
&= \mathbb{E}^\mu\left[U_0(V_0) - \int_{-\infty}^{V_0} U_0'(r)dr + \Pi_0(C_0) - \int_{C_0}^\infty \Pi_0'(y)dy\right] \\
&= \mathbb{E}^\mu\left[U_0(V_0) - \frac{1-F_0(V_0)}{f_0(V_0)}U_0'(V_0) + \Pi_0(C_0) - \frac{G_0(C_0)}{g_0(C_0)}\Pi_0'(C_0)\right] \\
&\leq (1-\delta)\mathbb{E}^\mu\left[\sum_{t=0}^\infty \delta^t X_t\left(V_t - \frac{1-F_0(V_0)}{f_0(V_0)}\gamma_B^t - C_t - \frac{G_0(C_0)}{g_0(C_0)}\gamma_S^t\right)\right],
\end{aligned} \tag{12}$$

where the first line follows by IR_0, the second by the fundamental theorem of calculus, the third by Fubini's theorem, and the last by BB_0, the law of iterated expectations, and the envelope formula (11) (and its analog for the seller). $\blacksquare$

Proof of Proposition 5. The argument in the text gives the form of the allocation rule x^{**} and the existence of a PIC mechanism $\mu^{**} = (x^{**}, p^{**})$ that solves the second-best problem (4). By balancing the transfers as in Athey and Segal (2012, Proposition 2), we may take μ^{**} to be BB. Since μ^{**} satisfies (3), IR_0 follows by reversing the steps in (12) once we note that the inequality on the last line holds as equality due to BB. This establishes the first claim.

In the interest of space, we only sketch the proof of the second claim. Note that since x^{**} is ex post monotone, Corollary 3 of PST shows that there exist transfers p such that $\mu = (x^{**}, p)$ is EPIC (see the opus cited for the definition of ex post monotonicity). In periods $t > 0$, given (not necessarily truthful) first-period reports v_0 and c_0, the transfers p_t can be simply taken to be the static Pivot transfers from the proof of Proposition 1 adjusted to account for

the wedge in (5). Hence we have IR for periods $t > 0$. Suppose then than we add constant participation fees π_B and π_S in period 0 such that $\pi_B + \pi_S = \frac{1-\delta}{1-\delta^T}\mathbb{E}^\mu\left[\sum_{t=0}^T \delta^t \left(P_{B,t} + P_{S,t}\right)\right]$. We then have a mechanism $\mu' = (x^{**}, p')$ that is EPIC, BB$_0$, and IR for $t > 0$. So it remains to check IR in period 0. But since BB$_0$ holds as equality, this follows again by reversing the steps in (12) because x^{**} satisfies (3). $\blacksquare$

Proof of Proposition 6. We start by establishing the first claim. To this end, fix $\Delta > 0$, and note that then $r \to 0$ iff $\delta \to 1$, so we may work with the latter.

For $b \in \mathbb{R}_+$, define the allocation rule x^b by

$$x_t^b = 1 \text{ iff } v_t \geq c_t, \ v_0 \geq -b, \text{ and } c_0 \leq b.$$

Let $y(\delta, b)$ denote the expected gains from trade under x^b. Then

$$y(\delta, b) = (1 - F_0(-b))G_0(b)\mathbb{E}[Y \mid -v_0, c_0 \leq b].$$

Note that $\lim_{\delta \to 1} \mathbb{E}[Y \mid -v_0, c_0 \leq b] = y^*$.

Now fix $\varepsilon \in (0, y^*)$. For any b, let $\delta_b < 1$ be such that $\mathbb{E}[Y \mid -v_0, c_0 \leq b] > y^* - \frac{\varepsilon}{2}$ for all $\delta > \delta_b$. Pick $\bar{b}$ large enough such that for all $b > \bar{b}$, we have

$$(1 - F_0(-b))G_0(b) > \frac{y^* - \varepsilon}{y^* - \frac{\varepsilon}{2}}.$$

Then for all $b > \bar{b}$ and $\delta > \delta_b$,

$$y(\delta, b) > (1 - F_0(-b))G_0(b)(y^* - \frac{\varepsilon}{2}) > y^* - \varepsilon.$$

To finish the proof, we show that for b and δ large enough, x^b satisfies (3). Define the "expected information rents" under x^b as

$$r(\delta, b) := (1 - \delta)\mathbb{E}^{x^b}\left[\sum_{t=0}^\infty \delta^t X_t \left(\frac{1 - F_0(V_0)}{f_0(V_0)}\gamma_B^t + \frac{G_0(C_0)}{g_0(C_0)}\gamma_S^t\right)\right].$$

We have

$$
\begin{aligned}
r(\delta, b) \ &\leq \ (1 - \delta)\mathbb{E}^{x^b}\left[\sum_{t=0}^\infty \delta^t X_t \left(\frac{1 - F_0(V_0)}{f_0(V_0)}\gamma_B^t + \frac{G_0(C_0)}{g_0(C_0)}\gamma_S^t\right)\bigg| V_0 \geq -b, C_0 \leq b\right] \\
&\leq \ \frac{1 - \delta}{1 - \delta\gamma_S}\mathbb{E}\left[\frac{1 - F_0(V_0)}{f_0(V_0)}\bigg| V_0 \geq -b\right] + \frac{1 - \delta}{1 - \delta\gamma_S}\mathbb{E}\left[\frac{G_0(C_0)}{g_0(C_0)}\bigg| C_0 \leq b\right],
\end{aligned}
$$

where for all b, the conditional expectations on the second line are finite. Thus for any b, there exists $\delta'_b < 1$ such that $r(\delta, b) < y^* - \varepsilon$ for all $\delta > \delta'_b$. Therefore, if we let $\bar{\delta}_b = \max\{\delta_b, \delta'_b\}$, then for all $b > \bar{b}$ and $\delta > \bar{\delta}_b$, we have the desired gains from trade as

$$y(\delta, b) > y^* - \varepsilon,$$

and, condition (3) is satisfied as

$$y(\delta, b) - r(\delta, b) > y^* - \varepsilon - (y^* - \varepsilon) = 0.$$

Hence, $y^{**}(\frac{-\log \delta}{\Delta}, \Delta) \geq y(\delta, b) > y^* - \varepsilon$ for δ large enough, or equivalently, for r small enough. Since $\varepsilon > 0$ was arbitrary, this establishes the first part of Proposition 6.

For the second part, fix $r > 0$, and let $y^{**}(\Delta) := y^{**}(r, \Delta)$. To simplify notation, we present the proof for the symmetric case where $m_B = m_S = 0$, $\alpha_B = \alpha_S =: \alpha$, and $\sigma_B = \sigma_S =: \sigma$. The general case follows by an analogous argument.

Given allocation rule x, denote the expected gains from trade by

$$y(x, \Delta) := (1 - e^{-r\Delta})\mathbb{E}^x \left[\sum_{t=0}^{\infty} e^{-r\Delta t} X_t (V_t - C_t) \right],$$

and denote the expected information rents by

$$r(x, \Delta) := (1 - e^{-r\Delta})\mathbb{E}^x \left[\sum_{t=0}^{\infty} e^{-(r+\alpha)\Delta t} X_t \left(\frac{1 - F_0(V_0)}{f_0(V_0)} + \frac{G_0(C_0)}{g_0(C_0)} \right) \right].$$

Then for any $\Delta > 0$, the second-best problem (4) becomes

$$y^{**}(\Delta) = \max \{ y(x, \Delta) \mid y(x, \Delta) - r(x, \Delta) \geq 0 \}.$$

Observe that $y(\cdot, \cdot)$ and $r(\cdot, \cdot)$ are continuous on $\{x : x_t \leq x_t^* \; \forall t\} \times (0, \infty)$. (Note that the restriction $x_t \leq x_t^*$ all t is not binding as by inspection of (5), $x_t^{**} \leq x_t^*$ all t, but it guarantees that $y(x, \Delta) - r(x, \Delta)$ is finite and allows for a straightforward continuity proof by the dominated convergence theorem.) Thus $y^{**}(\Delta)$ is a continuous function of Δ by the Theorem of the Maximum, and $\lim_{\Delta \to 0} y^{**}(\Delta) \leq y^*$ exists. Furthermore, by Proposition 5, we have $y^{**}(\Delta) = y(x^{**}(\Delta), \Delta) = r(x^{**}(\Delta), \Delta)$ for all $\Delta > 0$, and hence we have $\lim_{\Delta \to 0} r(x^{**}(\Delta), \Delta) = \lim_{\Delta \to 0} y^{**}(\Delta)$.

Suppose towards contradiction that $\lim_{\Delta \to 0} y^{**}(\Delta) = y^*$. By inspection of (5) this requires that $\liminf_{\Delta \to 0} \lambda(\Delta) = 0$, where $\lambda(\Delta)$ is the Lagrange multiplier on the constraint (3).

For $q \geq 0$, define the allocation rule x^q by setting $x^q_t = 1$ iff

$$v_t - c_t \geq \frac{q}{1+q} \left(\frac{1 - F_0(v_0)}{f_0(v_0)} + \frac{G_0(c_0)}{g_0(c_0)} \right).$$

Observe that $x^0 = x^*$.

Claim A.1 *For all $\Delta > 0$, all $q' > q$, $r(x^{q'}, \Delta) \leq r(x^q, \Delta)$ and $r(x^{\lambda(\Delta)}, \Delta) \leq r(x^{**}(\Delta), \Delta)$.*

Let $\Delta > 0$. The first inequality follows by noting that $q' > q$ implies $x^q_t \leq x^{q'}_t$ for all t, and the second by noting that $x^{\lambda(\Delta)}_t \leq x^{**}_t$ for all t. This establishes the claim.

Recalling that $\lim_{\Delta \to 0} r(x^{**}(\Delta), \Delta) = \lim_{\Delta \to 0} y^{**}(\Delta)$, we now obtain the desired contradiction from the following claim:

Claim A.2 *If $\liminf_{\Delta \to 0} \lambda(\Delta) = 0$, then $\lim_{\Delta \to 0} r(x^{**}(\Delta), \Delta) \geq 2y^*$.*

To establish Claim A.2, note that if $\liminf_{\Delta \to 0} \lambda(\Delta) = 0$, then there exists a monotone sequence (Δ_n) with $\Delta_n \to 0$ such that $\lambda_n := \lambda(\Delta_n)$ defines a monotone sequence (λ_n) with $\lambda_n \to 0$. By Claim A.1, for all n and k, with $n > k$, we have

$$r(x^{**}(\Delta_n), \Delta_n) \geq r(x^{\lambda_n}, \Delta_n) \geq r(x^{\lambda_k}, \Delta_n).$$

Thus for all k,

$$\lim_{n \to \infty} r(x^{**}(\Delta_n), \Delta_n) \geq \lim_{n \to \infty} r(x^{\lambda_k}, \Delta_n).$$

Note that

$$\lim_{n \to \infty} r(x^{\lambda_k}, \Delta_n)$$
$$= \lim_{\Delta \to 0} (1 - e^{-r\Delta}) \sum_{t=0}^{\infty} e^{-(r+\alpha)\Delta t} \mathbb{E} \left[\Phi \left(\frac{e^{-\alpha \Delta t}(V_0 - C_0) - \frac{\lambda_k}{1+\lambda_k} H(V_0, C_0)}{\sqrt{(1 - e^{-2\alpha \Delta t}) \frac{\sigma^2}{\alpha}}} \right) H(V_0, C_0) \right]$$
$$= r \int_0^{\infty} e^{-(r+\alpha)t} \mathbb{E} \left[\Phi \left(\frac{e^{-\alpha t}(V_0 - C_0) - \frac{\lambda_k}{1+\lambda_k} H(V_0, C_0)}{\sqrt{(1 - e^{-2\alpha t}) \frac{\sigma^2}{\alpha}}} \right) H(V_0, C_0) \right] dt =: r(x^{\lambda_k}, 0),$$

where Φ denotes the cdf of $N(0,1)$ and $H(V_0, C_0) := \frac{1 - F_0(V_0)}{f_0(V_0)} + \frac{G_0(C_0)}{g_0(C_0)}$, and where the third line follows by Karamata's (1930) generalization of a theorem of Hardy and Littlewood. By

the Monotone Convergence Theorem, we have,

$$\lim_{k \to \infty} r(x^{\lambda_k}, 0) = r \int_0^\infty e^{-(r+\alpha)t} \mathbb{E}\left[\Phi\left(\frac{e^{-\alpha t}(V_0 - C_0)}{\sqrt{(1 - e^{-2\alpha t})\frac{\sigma^2}{\alpha}}} \right) H(V_0, C_0) \right] dt$$
$$= \lim_{\Delta \to 0} r(x^*, \Delta) = 2y^*,$$

where the second equality obtains by another application of Karamata's theorem, and the last equality follows since by Proposition 4, $r(x^*, \Delta) = 2y^*$ for all $\Delta > 0$. Collecting from above, we have $\lim_{n \to \infty} r(x^{**}(\Delta_n), \Delta_n) \geq \lim_{k \to \infty} \lim_{n \to \infty} r(x^{\lambda_k}, \Delta_n) = 2y^*$. ∎

Proof of Proposition 7. The form of the allocation rule x^{***} is established by an argument analogous to that given in the text after Proposition 5. The existence of a public EPIC mechanism $\mu^{***} = (x^{***}, p)$ then follows by Proposition 8 of PST. For statement 1 we may use the balancing argument of Athey and Segal (2012, Proposition 2) to get a public PIC and BB mechanism, which is shown to be IR$_0$ verbatim as in the proof of Proposition 5. Statement 2 can then be proven starting from the EPIC mechanism $\mu^{***} = (x^{***}, p)$ exactly the same way as the corresponding claim in Proposition 5. ∎

References

[1] Abreu, Dilip, Paul Milgrom, and David Pearce (1991). "Information and Timing in Repeated Partnerships," *Econometrica*, 59(6): 1713–33.

[2] d'Aspremont, C., and L. Gerard-Varet (1979). "Incentives and Incomplete Information," *Journal of Public Economics*, 11:25–45.

[3] Athey, Susan, and David A. Miller (2007). "Efficiency in Repeated Trade with Hidden Valuations," *Theoretical Economics*, 2(3): 299–354.

[4] Athey, Susan, and Ilya Segal (2007). "Designing Dynamic Mechanisms," *American Economic Review Papers and Proceedings*, 97(2): 131–136.

[5] Athey, Susan, and Ilya Segal (2012). "An Efficient Dynamic Mechanism," working paper.

[6] Battaglini, Marco, (2005). "Long-Term Contracting with Markovian Consumers," *American Economic Review*, 95(3): 637–658.

[7] Battaglini, Marco, and Rohini Lamba (2012). "Optimal Dynamic Contracting," working paper.

[8] Baron, David, and D. Besanko (1984). "Regulation and Information in a Continuing Relationship," *Information Economics and Policy*, 1: 447–470.

[9] Bergemann, Dirk, and Juuso Välimäki (2010). "The Dynamic Pivot Mechanism," *Econometrica*, 78(2): 771–789.

[10] Boleslavsky, Raphael, and Maher Said (2012). "Progressive Screening: Long-Term Contracting with a Privately Known Stochastic Process," working paper.

[11] Courty, Pascal, and Hao Li (2000). "Sequential Screening," *Review of Economic Studies*, 67: 697–717.

[12] Cramton, Peter, Robert Gibbons, and Paul Klemperer (1987). "Dissolving a Partnership Efficiently," *Econometrica*, 55(3): 615–32.

[13] Daley, D. J. (1968). "Stochastically monotone Markov chains," *Probability Theory and Related Fields*, 10(4): 305–317.

[14] Escobar, Juan, and Juuso Toikka (2012). "Efficiency in Games with Markovian Private Information," working paper.

[15] Eső, Péter, and Balász Szentes (2007). "Optimal Information Disclosure in Auctions and the Handicap Auction," *Review of Economic Studies*, 74(3): 705–731.

[16] Fudenberg, D., D. Levine, and E. Maskin (1994). "The Folk Theorem with Imperfect Public Information," *Econometrica*, 62(5): 997–1039.

[17] Harrison, Michael (1985). *Brownian Motion and Stochastic Flow Systems*. John Wiley & Sons, New York.

[18] Holmström, Bengt (1979). "Groves' Scheme on Restricted Domains," *Econometrica*, 47(5): 1137–1144.

[19] Karamata, J. (1930). "Über die Hardy-Littlewoodschen Umkehrungen des Abelschen Stetigkeitssatzes," *Mathematische Zeitschrift*, 32(1): 319–320.

[20] Krishna, Vijay, and Motty Perry (2000). "Efficient Mechanism Design," working paper.

[21] Miller, David (2012). "Robust Collusion with Private Information," *Review of Economic Studies*, 79(2): 778–811.

[22] Myerson, Roger (1986). "Multistage Games with Communication," *Econometrica*, 54(2): 323–58.

[23] Myerson, Roger B., and Mark A. Satterthwaite (1983). "Efficient Mechanisms for Bilateral Trading," *Journal of Economic Theory*, 29(2): 265–281.

[24] Pavan, Alessandro, Ilya Segal, and Juuso Toikka (2012). "Dynamic Mechanism Design," working paper.

[25] Segal, Ilya, and Michael Whinston (2011). "A Simple Status Quo that Ensures Participation (with Application to Efficient Bargaining)," *Theoretical Economics*, 6(1):109–125.

[26] Williams, D. (1991). *Probability with Martingales.* Cambridge University Press, Cambridge, UK.

[27] Williams, S. (1999). "A Characterization of Efficient, Bayesian Incentive Compatible Mechanisms," *Economic Theory*, 14, 155–180.