

Mechanisms for Repeated Trade[†]

By ANDRZEJ SKRZYPACZ AND JUUSO TOIKKA*

How does feasibility of efficient repeated trade depend on the features of the environment such as persistence of values, private information about their evolution, or trading frequency? We derive a necessary and sufficient condition for efficient, unsubsidized, and voluntary trade, which implies that efficient contracting requires sufficient congruence of expectations. This translates to bounds on persistence of values and on private information about their evolution, and distinguishes increasing patience from more frequent interaction; the latter need not facilitate efficiency even when the former does. We also discuss second-best mechanisms and extend the characterization to general dynamic collective choice problems. (JEL C73, D82, D86)

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 interaction 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 can trade more than once.

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 multidimensional 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.

* Skrzypacz: Graduate School of Business, Stanford University, 655 Knight Way, Stanford, CA 94305 (e-mail: skrz@stanford.edu); Toikka: Department of Economics, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139 (e-mail: toikka@mit.edu). We thank Jeremy Bulow and Johannes Hörner for comments, and Yair Livne and Mikhail Panov for research assistance. This paper was formerly titled "Mechanisms for Repeated Bargaining."

[†] Go to <http://dx.doi.org/10.1257/mic.20140173> to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

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 extent 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.

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 of 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

¹In 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 a proof of the Myerson-Satterthwaite theorem due to Makowski and Mezzetti (1994), 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.

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 the characterization for static private value environments by Makowski and Mezzetti (1994), Williams (1999), and Krishna and Perry (2000).³

In terms of the model, this paper is most closely related to the work of Athey and Miller (2007) and Athey and Segal (2007, 2013). Athey and Miller study repeated bilateral trade when types are independent and identically distributed across periods, thus abstracting from the two features central to our analysis. Athey and Segal (2013) 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 multidimensional initial information by extending the necessary condition for incentive compatibility by Pavan, Segal, and Toikka (2014—henceforth PST) to our setting. Our discussion of second best also builds on their results. Athey and Segal (2013) 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

²This result is closely related to the contrasting findings by Battaglini (2005) and Boleslavsky and Said (2013), 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.

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

Athey and Segal's method of balancing transfers, which yields a dynamic version of the AGV mechanism of d'Aspremont and Gérard-Varet (1979).

I. The Environment

Consider the following dynamic bargaining environment with two-sided private information.⁴ There are two agents, a buyer (B) and a seller (S), who may trade a nonstorable good and a numeraire in each of countably many periods indexed by $t = 0, 1, \dots, T$, 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 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 $\{F(\cdot | \cdot; \theta_B)\}_{\theta_B \in \Omega_B}$ of Markov kernels $F(\cdot | \cdot; \theta_B) : \mathcal{V} \times \mathcal{V} \rightarrow [0, 1]$.⁵ Specifically, given any $(\theta_B, v_0) \in \Omega_B \times \mathcal{V}_0$, V follows the Markov chain $\langle v_0, F(\cdot | \cdot; \theta_B) \rangle$ with transitions $F(\cdot | \cdot; \theta_B)$ 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 $\langle v_0, F(\cdot | \cdot; \theta_B) \rangle$. The seller is only informed of the distribution F_0 and the family $\{F(\cdot | \cdot; \theta_B)\}_{\theta_B \in \Omega_B}$, 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 $\{G(\cdot | \cdot; \theta_S)\}_{\theta_S \in \Omega_S}$ 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 $\langle c_0, G(\cdot | \cdot; \theta_S) \rangle$ determined by his period-0 private information (θ_S, c_0) , whereas the buyer only knows the pair $(G_0, \{G(\cdot | \cdot; \theta_S)\}_{\theta_S \in \Omega_S})$. 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

⁴We discuss more general environments in Section VI.

⁵That is, $F(\cdot | v; \theta_B)$ is a cumulative distribution function for all $(v, \theta_B) \in \mathcal{V} \times \Omega_B$, and $F(v | \cdot; \cdot) : \mathcal{V} \times \Omega_B \rightarrow [0, 1]$ is (Borel) measurable for all $v \in \mathcal{V}$.

well-behaved, assume $E\left[\sum_{t=0}^T \delta^t |V_t| \mid \theta_B, v_0\right]$ and $E\left[\sum_{t=0}^T \delta^t |C_t| \mid \theta_S, c_0\right]$ are finite for all $(\theta, v_0, c_0) \in \Omega \times \mathcal{V}_0 \times \mathcal{C}_0$, where $\theta := (\theta_B, \theta_S) \in \Omega_B \times \Omega_S =: \Omega$.⁶

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

Example 1 (Conditionally Independent and Identically Distributed Types): Given any $\theta \in \Omega$, values and costs are independent and identically distributed 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 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 independent and identically distributed 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 independent and identically distributed 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 independent and identically distributed draws from a Normal distribution. The constants γ_i, m_i , and the parameters of the Normal distributions may be private information.

II. 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

⁶All 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.

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. While the latter is best viewed as a purely theoretical construct, it should be noted that real-world trading institutions are typically price-based indirect mechanisms, which entail an intermediate level of transparency. Thus, it is important to look beyond public mechanisms.

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

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

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 $0, \dots, t$.⁷ A decision rule μ induces a multi-stage game form where in every period $t = 0, 1, \dots, T$, given history of reports $(\hat{\theta}_B, \hat{\theta}_S, \hat{v}_0, \hat{c}_0, \dots, \hat{v}_{t-1}, \hat{c}_{t-1}) \in \Omega \times (\mathcal{V} \times \mathcal{C})^t$, 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 $\hat{v}_t \in \text{supp } F(\cdot | \hat{v}_{t-1}; \hat{\theta}_B)$ and the seller reports $\hat{c}_t \in \text{supp } G(\cdot | \hat{c}_{t-1}; \hat{\theta}_S)$ (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 $(\hat{\theta}_B, \hat{\theta}_S, \hat{v}_0, \hat{c}_0, \dots, \hat{v}_t, \hat{c}_t)$.

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 .

⁷Restricting attention to deterministic mechanisms is without loss for our results so we do so throughout to simplify notation.

The set of feasible period- t histories of agent i in a mechanism μ is denoted $\mathcal{H}_{i,t}^\mu$, or simply $\mathcal{H}_{i,t}$, if μ is clear from the context.⁸ A strategy for the buyer is then a sequence of (measurable) functions $\sigma_B = (\sigma_{B,t})_{t=0}^T$ where $\sigma_{B,0} : \mathcal{H}_{B,0} \rightarrow \Omega_B \times \mathcal{V}_0$ and $\sigma_{B,t} : \mathcal{H}_{B,t} \rightarrow \mathcal{V}$ for $t \geq 1$. The seller's strategy $\sigma_S = (\sigma_{S,t})_{t=0}^T$ is defined analogously. We say that agent i 's history $h_{i,t} \in \mathcal{H}_{i,t}$ is *truthful* if his own reports have been truthful in all periods $0, \dots, t - 1$. 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 $h_{B,t} \in \mathcal{H}_{B,t}$ and $h_{S,t} \in \mathcal{H}_{S,t}$, we denote the expected continuation utilities for the buyer and the seller, respectively, by

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

and

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

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 2013; Bergemann and Välimäki 2010; or PST):

DEFINITION 1: A mechanism $\mu = (x, p)$ is *efficient* (E) if for all t , $x_t = \mathbf{1}_{\{v_i \geq c_i\}}$.

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

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

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 $(h_{B,t}, h_{S,t}) \in \mathcal{H}_{B,t} \times \mathcal{H}_{S,t}$, 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 $j \neq i$.

⁸Formally, the sets of buyer's histories in a public mechanism are defined recursively by setting $\mathcal{H}_{B,0} := \Omega_B \times \mathcal{V}_0$, $\mathcal{H}_{B,1} := \mathcal{H}_{B,0} \times \Omega \times \mathcal{V}_0 \times \mathcal{C}_0 \times \{0, 1\} \times \mathbb{R}^2 \times \mathcal{V}$, and $\mathcal{H}_{B,t} := \mathcal{H}_{B,t-1} \times \mathcal{V} \times \mathcal{C} \times \{0, 1\} \times \mathbb{R}^2 \times \mathcal{V}$ 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 $\mathcal{H}_{B,0} := \Omega_B \times \mathcal{V}_0$, $\mathcal{H}_{B,1} := \mathcal{H}_{B,0} \times \Omega_B \times \mathcal{V}_0 \times \mathcal{V} \times \{0, 1\}$, and $\mathcal{H}_{B,t} := \mathcal{H}_{B,t-1} \times \mathcal{V} \times \mathcal{V} \times \{0, 1\}$ 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 $\mathcal{H}_{S,t}$ for the seller are defined analogously.

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 $(\theta, v_0, c_0) \in \Omega \times \mathcal{V}_0 \times \mathcal{C}_0$,

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

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

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

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

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

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 $\mathcal{H}_{i,t}^\eta$ in the blind mechanism η corresponds to a partition of the set of histories $\mathcal{H}_{i,t}^\mu$ 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.

III. 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 A1)—to the Appendix.

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 A2 (along with all other omitted proofs).

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

(i) *The first-best gains from trade, Y , satisfy*

$$(1) \quad \inf_{\theta_B, v_0} E[Y \mid \theta_B, v_0] + \inf_{\theta_S, c_0} E[Y \mid \theta_S, c_0] \geq E[Y].$$

(ii) *There exists a blind mechanism that is E, IC, IR₀, and BB₀.*

(iii) *There exists a public mechanism that is E, PIC, IR₀, and BB.*

(iv) *There exists a public mechanism that is E, EPIC, IR, and BB₀.*

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

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 statement (ii) 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, statements (iii) and (iv) 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 individually rational in period 0, or individually rational in every period and budget balanced ex ante. (In the latter case, incentive compatibility can be strengthened to hold within-period ex post.) However, (1) is not sufficient to simultaneously guarantee ex post budget balance and individual rationality in every period. Indeed, as we discuss below, when the horizon is finite, no such mechanism exists, and even with an infinite horizon, a stronger condition is needed.

For the special case of a static model, Proposition 1 follows from the characterizations by Makowski and Mezzetti (1994); Williams (1999); or Krishna and Perry (2000).⁹ However, in a static model, inequality (1) is satisfied only in the trivial case where there is common knowledge of positive gains from trade, which yields an alternative proof of Myerson and Satterthwaite's (1983) theorem. In contrast, in dynamic settings the existence of future surplus allows (1) to be satisfied in many cases where 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, $E[Y]$, on each agent's initial private information.¹⁰ 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 distribution 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 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.

⁹See also Athey and Miller (2007) for an exploration of different combinations of properties in the case of independent and identically distributed types drawn from known distributions.

¹⁰Trivially, if at most one agent has private information in period 0 (i.e., if $\Omega_B \times \mathcal{Y}_0$ or $\Omega_S \times \mathcal{C}_0$ is a singleton), then at least one of the infimum terms equals $E[Y]$, 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 $\inf E[Y | \theta_B, v_0] \wedge \inf E[Y | \theta_S, c_0] = 0$), then the inequality is never satisfied.

Since the payoff-equivalence property is taken as given, the proof of Proposition 1 parallels that of the static case, with a repeated Groves' scheme and the balancing trick of Athey and Segal (2013) 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 Groves' scheme, 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.

As suggested by the above proof sketch, Proposition 1 immediately generalizes to other dynamic bargaining problems. We state one such extension to n -person problems in the concluding remarks (Section VI).

One may wonder about the possibility of satisfying both individual rationality and budget balance period by period. With a finite horizon, no mechanism satisfies E, IC, IR, and BB. This follows by noting that under these conditions, the last period T essentially reduces to a static problem because the budget has to be balanced at the end of period $T - 1$.¹¹ The ability of finite long-term contracts to support efficient trade thus rests on third party financing (BB_0) or contract enforcement (IR_0). With an infinite horizon, a general characterization appears to be out of reach as it would require tracking the agents' beliefs. However, for *public* mechanisms, it can be shown by an argument similar to the proof of Proposition 1 that if the kernels have full support, then E, PIC, IR, and BB can be satisfied if and only if inequality (1) holds and $\inf_{v_0} E[Y | \theta, v_0, c_{-1}] + \inf_{c_0} E[Y | \theta, v_{-1}, c_0] \geq E[Y | \theta, v_{-1}, c_{-1}]$ for all (θ, v_{-1}, c_{-1}) , where the expectation $E[Y | \theta, v_0, c_{-1}]$ is taken with respect to the type processes with parameters θ , with c_0 drawn from $G(\cdot | c_{-1}; \theta_S)$ (and similarly for $E[Y | \theta, v_{-1}, c_0]$ and $E[Y | \theta, v_{-1}, c_{-1}]$).¹²

A. Conditions for Payoff Equivalence

We provide here an informal discussion of sufficient conditions for the payoff-equivalence property, relegating a formal treatment to the Appendix. 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 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

¹¹The formal result requires that with positive probability, the period- T type distribution satisfies the assumptions of Myerson-Satterthwaite theorem. For example, this holds if for all $(v, c, \theta_B, \theta_S)$, the supports of $F(\cdot | v; \theta_B)$ and $G(\cdot | c; \theta_S)$ are intervals whose intersection has a nonempty interior.

¹²Athey and Segal (2007, 2013) show that a public mechanism that is E, PIC, IR, and BB exists in a class of ergodic Markov environments with known parameters and finitely many types, if agents are sufficiently patient. We conjecture that the same is true in our model whenever (1) is satisfied for some $\delta < 1$.

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 immediately covered by the result of PST as here the agents' initial information is multidimensional. However, by combining the standard payoff-equivalence argument for static multidimensional 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 multidimensional initial information.¹³ We provide general sufficient conditions for this in the Appendix by introducing the notion of a *regular* environment (Definition A2) and proving that such environments have the payoff-equivalence property (Proposition A1). 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 A1), 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 $\{H(\cdot | \cdot; \theta_i)\}_{\theta_i \in \Omega_i}$ 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:

- (i) Every $H(\cdot | z; \theta_i)$, $(z, \theta_i) \in \mathcal{Z} \times \Omega_i$, is absolutely continuous with density $h(\cdot | z; \theta_i)$ strictly positive on \mathcal{Z} .
- (ii) The kernel $H(\cdot | \cdot; \cdot) : \mathbb{R}^{k+2} \rightarrow \mathbb{R}$ 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 $\theta_i \in \text{int } \Omega_i$,

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

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

¹³ Alternatively, one can consider a fictitious environment where the initial period is divided into subperiods and the agents observe and report one dimension of their multidimensional initial type in each subperiod. Any mechanism that is incentive compatible in the original environment remains incentive compatible in the fictitious environment. As private information is one dimensional in the latter, under appropriate regularity conditions one can apply the ex post payoff equivalence result of PST (Theorem 2) to establish a multidimensional payoff equivalence result analogous to our Proposition A1.

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 independent and identically distributed types in Example 1 are smooth if $F(\cdot | \theta_B)$ and $G(\cdot | \theta_S)$ 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 A2 by Lemma A2).

IV. Applications

We now use Proposition 1 to explore, for a fixed discount rate, how the feasibility of efficient contracting 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 (ii)–(iv) in Proposition 1 are satisfied. By Proposition 1, this is the case in environments that have the payoff-equivalence property if and only if

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

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.

A. 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 independent and identically distributed case.

Example 4 (Independent and Identically Distributed with Known Parameters): Consider the environment of Example 1 with known parameters. Then V and C are, respectively, independent and identically distributed 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 A1. For simplicity, take $T = \infty$ and suppose $\mathcal{V} = \mathcal{C}$. Then,

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

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 onward. Hence, by Proposition 1, satisfactory trading mechanisms exist if and only if $2\delta E[Y] \geq E[Y]$, or equivalently, if and only if $\delta \geq \frac{1}{2}$. 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.¹⁴

At the other extreme, when types are perfectly persistent and do not change over time, the analysis reduces to the static case. The Myerson-Satterthwaite (1983) theorem then implies that gains from trade are bounded away from first-best for any level of the discount factor.¹⁵

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 (and the horizon may be 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 $H(\cdot | \cdot)$ and $H'(\cdot | \cdot)$, respectively. We say that Z is (weakly) more persistent than Z' if there exists a distribution Φ satisfying the following conditions:

- (i) Φ is the unique invariant distribution admitted by the kernels $H(\cdot | \cdot)$ and $H'(\cdot | \cdot)$.
- (ii) Φ is the initial distribution of the chains Z and Z' .
- (iii) For all $(z_0, z) \in \mathcal{Z}^2$ and all $t \geq 1$, the t -step distributions satisfy

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

Conditions (i) and (ii) ensure that increasing persistence only affects the short-run properties of the Markov chain. Condition (iii) captures the idea that a more persistent chain converges slower to the invariant distribution. For example, if Z is a Markov chain started from its unique invariant distribution Φ , and Z' is a sequence of independent and identically distributed 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 $H^\alpha(\cdot | z) := \alpha H(\cdot | z) + (1 - \alpha)\Phi(\cdot)$, $z \in \mathcal{Z}$, where $H(\cdot | \cdot)$ is the kernel of $Z = Z^1$, is ordered by persistence with Z^α more persistent than $Z^{\alpha'}$ if and only if $\alpha \geq \alpha'$. Note that taking Z to be a fully persistent chain yields the renewal model

¹⁴Their 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 independent and identically distributed types the two are equivalent.

¹⁵Formally, this can be shown by applying our Proposition 1 to the renewal model of Example 2 with known parameters and $\gamma_i = 1$ for $i \in \{B, S\}$. This environment is regular by Lemma A2, and hence it has the payoff-equivalence property by Proposition A1. (Because of the non-changing types, payoff equivalence can also be established using the standard static argument, see, e.g., Milgrom and Segal 2002.)

of Example 2, which thus provides one possible parameterization of persistence spanning from independent and identically distributed 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 $H(\cdot | \cdot)$. We say that Z is stochastically monotone if for all $(z, z') \in \mathbb{R}^2$ with $z > z'$, $H(\cdot | z)$ first-order stochastically dominates $H(\cdot | z')$.

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:

- (i) The first-best gains from trade under (V, C) satisfy inequality (1), and
- (ii) 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).¹⁶

The key to the proof is the observation that for monotone chains with known parameters, inequality (1) reduces to $E[Y | \min \mathcal{V}_0] + E[Y | \max \mathcal{C}_0] \geq E[Y]$. To see this, note that the lowest possible value for the buyer is not only the worst current type, but also results in the lowest possible distribution of values in every future period in the sense of first-order stochastic dominance because the chain is assumed to be monotone. (The seller is treated analogously.) Furthermore, since any two comparable pairs of chains (V, C) and (V', C') necessarily have the same invariant distribution, $E[Y]$ is the same under both pairs of processes. Hence, it remains to show that increasing persistence in the sense of Definition 7 decreases $E[Y | \min \mathcal{V}_0]$ and $E[Y | \max \mathcal{C}_0]$. Heuristically, this follows because a monotone chain started from an extreme value (e.g., at $\min \mathcal{V}_0$) converges monotonically to the invariant distribution so that increasing persistence means that the distribution remains concentrated around the unfavorable extreme value for longer.

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 5 (Renewals with Known Parameters): Consider the renewal model of Example 2 with known parameters. By Lemma A2 this environment is regular and

¹⁶To see that stochastic monotonicity is needed, consider a two-period problem with $\mathcal{V} = \mathcal{C} = [0, 1]$. If types are independent and identically distributed from the uniform distribution, then (1) holds if and only if $\delta = 1$. However, if types are perfectly negatively correlated (i.e., $v_1 = 1 - v_0$ and $c_1 = 1 - c_0$) with period-0 types distributed uniformly, then for $\delta = 1$ the worst initial types are $v_0 = c_0 = \frac{1}{2}$ with average expected gains from trade one-eighth each. As $E[Y] = \frac{1}{6}$, this implies that inequality (1) is slack, and, hence, the critical discount factor is lower than in the independent and identically distributed case by continuity. By slightly perturbing the latter process we obtain an ergodic process which is more persistent than the independent and identically distributed process in the sense of Definition 7, yet for which the critical discount factor is strictly less than 1.

hence it has the payoff-equivalence property by Proposition A1. 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

$$F^{(t)}(\cdot | v_0) = \gamma_B^t \mathbf{1}_{[v_0, \infty)} + (1 - \gamma_B^t) F_0, \quad \text{and}$$

$$G^{(t)}(\cdot | c_0) = \gamma_S^t \mathbf{1}_{[c_0, \infty)} + (1 - \gamma_S^t) G_0.$$

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

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

and

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

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

$$(2) \quad \sqrt{\rho_B \rho_S} \geq \frac{1 - \delta}{\delta},$$

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 negative finding from the case of perfectly persistent types.¹⁷ 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 7).

¹⁷To 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 $F(\min \mathcal{V} | \min \mathcal{V}; \theta'_B) = 1$. Then $\inf_{\theta_B, v_0} E[Y | \theta_B, v_0] = E[Y | \theta'_B, \min] = 0$, 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 $\theta'_B \in \Omega_B$.

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

- (i) $\Omega' \subset \Omega$,
- (ii) $F'_0 = F_0|_{V \times \Omega'_B}$ and $G'_0 = G_0|_{C \times \Omega'_S}$,
- (iii) $F'(\cdot | \cdot; \theta_B) = F(\cdot | \cdot; \theta_B)$ and $G'(\cdot | \cdot; \theta_S) = G(\cdot | \cdot; \theta_S)$ for all $\theta \in \Omega'$, and
- (iv) $E[Y] = E[Y']$.

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 6 (Independent and Identically Distributed 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:

$$f(v | \theta_B) = \theta_B + 2(1 - \theta_B)v \text{ for } \theta_B \in [0, 2],$$

$$g(c | \theta_S) = \theta_S + 2(1 - \theta_S)c \text{ for } \theta_S \in [0, 2].$$

The case $\theta_i = 1$ corresponds to the uniform distribution. Note that $f(\cdot | \theta_B)$ and $g(\cdot | \theta_S)$ 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 $E[Y] = \frac{1}{6}$. 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

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

This leads us to the following corollary: In the independent and identically distributed 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:

- 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).
- If there is enough uncertainty about the processes, satisfactory mechanisms may not exist even if the players are arbitrarily patient.
- 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 $E[Y]$. Furthermore, if we change the distribution of θ_B without changing its support, and keeping the unconditional distribution of V fixed, then $\inf_{\theta_S, c_0} E[Y | \theta_S, c_0]$ is unaffected. Similarly, $\inf_{\theta_B, v_0} E[Y | \theta_B, v_0]$ 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 7 (Privately Known Persistence): Let $T < \infty$, and construct the type processes V and C as follows. Fix $\Delta > 0$. Define the “base kernels” $\bar{F}(\cdot | \cdot)$ and $\bar{G}(\cdot | \cdot)$ on $[0, 1]^2$ by letting $\bar{F}(\cdot | z)$ and $\bar{G}(\cdot | z)$ 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 cumulative distribution function 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 $\theta_i \in \Omega_i \subset [0, 1]$, $i \in \{B, S\}$,

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

and

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

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 A3 in the Appendix shows that this environment is regular, and hence it has the payoff-equivalence property by Proposition A1.

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 7,¹⁸ 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 $E[Y] = \frac{1}{6}$ 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 $i \in \{B, S\}$).

C. 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 agent's types in adjacent periods. Based on our remarks on persistence, the

¹⁸Note 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.

former tends to be favorable for efficient bargaining whereas the latter tends to be detrimental. Depending on the parameters, either may dominate.

Example 8: Consider the renewal model with known parameters considered in Example 5. 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 $\Delta \rightarrow 0$), the necessary and sufficient condition (2) derived in Example 5 simplifies further to

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

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.¹⁹

Note that, in contrast, taking $\delta = e^{-r\Delta} \rightarrow 1$ by sending $r \rightarrow 0$ 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 2013; Escobar and Toikka 2013; 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.

V. 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 multidimensional types even in a static setting. Hence, we are led 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

¹⁹More 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 if and only if $\sqrt{\lambda_B \lambda_S} \geq r$. We conjecture that this “comparative static” with respect to Δ extends to the generalized renewal processes considered in Example 7.

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.

A. 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

$$d\tilde{V}_\tau = -\alpha_B(\tilde{V}_\tau - m_B) d\tau + \sigma_B dW_\tau^B,$$

$$d\tilde{C}_\tau = -\alpha_S(\tilde{C}_\tau - m_S) d\tau + \sigma_S dW_\tau^S,$$

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

$$v_t = \gamma_B v_{t-1} + (1 - \gamma_B)m_B + \varepsilon_{B,t}$$

$$c_t = \gamma_S c_{t-1} + (1 - \gamma_S)m_S + \varepsilon_{S,t}$$

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 $N\left(m_i, \frac{\sigma_i^2}{2\alpha_i}\right)$. *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.²⁰

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,*

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

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 $\inf_{v_0} E[Y | v_0] = \inf_{c_0} E[Y | c_0] = 0$. 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

²⁰The result extends a priori to the case where some or all of the parameters $(\alpha_i, \sigma_i, m_i)$ are private information.

enough so that conditional on v_0 or c_0 , convergence to the 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.²¹

As the first step toward 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, IR₀, and BB₀ mechanism, then*

$$(3) \quad E^\mu \left[\sum_{t=0}^{\infty} \delta^t X_t \left(V_t - \frac{1 - F_0(V_0)}{f_0(V_0)} \gamma'_B - C_t - \frac{G_0(C_0)}{g_0(C_0)} \gamma'_S \right) \right] \geq 0.$$

The left-hand side of (3) is the expected dynamic virtual surplus, which is shown to be nonnegative 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, IR₀, and BB₀ mechanism are bounded from above by

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

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, IR₀, and BB₀, and the gains from trade under μ achieve y^{**} , which we refer to correspondingly as the *expected second-best gains from trade*. Next, we show 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:*

- (i) *There exist a public second-best mechanism that is PIC, IR₀, and BB.*
- (ii) *There exist a public second-best mechanism that is EPIC, IR, and BB₀.*

*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*

$$(5) \quad v_t - c_t \geq \frac{\lambda}{1 + \lambda} \left(\frac{1 - F_0(v_0)}{f_0(v_0)} \gamma'_B + \frac{G_0(c_0)}{g_0(c_0)} \gamma'_S \right),$$

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

²¹It 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 non-compact supports, which are a consequence of our assumption that the processes be started from the stationary distribution.

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 Proposition 5 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 1 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 $t \rightarrow 0$, 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 $E[Y]$ as we vary the parameters. However, it is possible to use approximation arguments to show that the findings for the limits $r \rightarrow 0$ and $\Delta \rightarrow 0$ are qualitatively different. In order to state the result, let $y^* := E[(V_0 - C_0)^+] = E[Y]$, 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,*

$$(i) \text{ for all } \Delta > 0, \lim_{r \rightarrow 0} y^{**}(r, \Delta) = y^*,$$

$$(ii) \text{ for all } r > 0, \lim_{\Delta \rightarrow 0} y^{**}(r, \Delta) < y^*.$$

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 $r \rightarrow 0$, we may move the truncations arbitrarily far out in the tails to obtain an approximately efficient mechanism.²²

²²This is closely related to the observation in footnote 21.

In contrast, for the frequent-interaction limit $\Delta \rightarrow 0$, the discount rate r is held constant, and, hence, the lack of uniform bounds on the convergence of the t -step distributions $F^{(t)}(\cdot | v_0)$ 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.

B. *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.²³

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

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

The proof is the same and, hence, omitted.²⁴ 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 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 $\gamma_B = 0$

²³ 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.

²⁴ The only difference is in the expressions for the derivatives. For example, formula (A4) now becomes

$$U'_0(\theta_B) = (1 - \delta) E^\mu \left[\sum_{t=0}^{\infty} \delta^t (1 - \gamma_B^t) X_t \mid \theta_B \right] \quad \text{a.e. } \theta_B,$$

where we have omitted conditioning on the known constant v_0 .

corresponds to values being drawn independent and identically distributed 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 nonincreasing. Then,*

(i) *There exist a public second-best mechanism that is PIC, IR₀, and BB.*

(ii) *There exist a public second-best mechanism that is EPIC, IR, and BB₀.*

*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 independent and identically distributed 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, 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 (2013) 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.

VI. 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 IV, 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 multidimensional initial information, which may be useful elsewhere.

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 $t = 0, 1, \dots, T$, 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 $i \in I := \{1, \dots, n\}$ 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) $u_i : \mathcal{X} \times \mathcal{Z}_i \rightarrow \mathbb{R}$. 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 $(H_{i,0}, \{H(\cdot | \cdot; \theta_i)\}_{\theta_i \in \Omega_i})$, 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 II, and the payoff-equivalence property can be defined analogously to Definition 5. Our proof of payoff equivalence (Proposition A1) 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 A2), 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

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

where χ^* is a static first-best allocation rule.²⁵ We then have the following generalization of Proposition 1, which provides a dynamic version of the characterization by Makowski and Mezzetti (1994).²⁶

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

(i) *The first-best social surplus, S , satisfies*

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

(ii) *There exists a blind mechanism that is E, IC, IR₀, and BB₀.*

(iii) *There exists a public mechanism that is E, PIC, IR₀, and BB.*

(iv) *There exists a public mechanism that is E, EPIC, IR, and BB₀.*

Proposition 1' permits an analysis analogous to Section IV 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.

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 nontransferable. 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 independent and identically distributed 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 (2013) and Bergemann and Välimäki (2010) as we may take their efficient dynamic mechanisms as the starting point in

²⁵I.e., $\chi^*(z_1, \dots, z_n) \in \operatorname{argmax}_{x \in X} \sum_{i=1}^n u_i(x, z_i)$ for all type profiles $(z_1, \dots, z_n) \in \mathcal{Z}_1 \times \dots \times \mathcal{Z}_n$.

²⁶Define the mechanism $\mu = (x^*, p)$ in the beginning of the proof of Proposition 1 to consist of the repetition of the static Groves' scheme $p_{i,t} = \sum_{j \neq i} u_j(x_t^*, z_{j,t})$. Then μ is E, EPIC, and IR with budget deficit equal to $(n-1)E[S]$. The rest of the proof now goes through with the obvious adjustments.

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 independent and identically distributed case, or Athey and Segal 2007, 2013, for positive limit results in settings with serial dependence.) The design of second-best mechanisms when initial information is multidimensional is another natural but challenging next step.

APPENDIX

A. 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 multidimensional.

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 A1: A state representation of the process V is a triple (\mathcal{E}, Q, ψ) , where \mathcal{E} is a measurable space, Q is a probability distribution on \mathcal{E} , and $\psi : \Omega_B \times \mathcal{V} \times \mathcal{E} \rightarrow \mathcal{V}$ is a (measurable) function such that, for all $(\theta_B, v) \in \Omega_B \times \mathcal{V}$, $\psi(\theta_B, v, \cdot) : \mathcal{E} \rightarrow \mathcal{V}$ is a random variable with distribution $F(\cdot | v; \theta_B)$.

Given a state representation (\mathcal{E}, Q, ψ) , 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 independent and identically distributed random variables (see, e.g., Williams (1991, 209)) to a mixture over the parameterized collection of Markov chains $\left\{ \left\{ v_0, F(\cdot | v; \theta_B) \right\} \right\}_{(\theta_B, v_0) \in \Omega_B \times \mathcal{V}}$, and hence a state representation exists. For example, the *canonical representation* obtains by taking $\mathcal{E} = [0, 1]$ and $\psi(\theta_B, v, \cdot) = F^{-1}(\cdot | v; \theta_B)$ for all (θ_B, v) , and letting Q to be the uniform distribution.

DEFINITION A2: *The process V is regular if there exists a state representation (\mathcal{E}, Q, ψ) 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*

$$(i) \quad \|\nabla_{\theta_B} \psi\| \leq d, \text{ and}$$

$$(ii) \quad |\partial_v \psi| \leq b.$$

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 A1: *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 A2. 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 $F(\cdot | \cdot; \cdot)$ is continuously differentiable, and for all (v, θ_B) , there is a density $f(\cdot | v; \theta_B)$ strictly positive on \mathcal{V} . Therefore, for all $(\varepsilon, v, \theta_B)$ and $v' := F^{-1}(\varepsilon | v; \theta_B)$, the Implicit Function Theorem implies

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

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 A2. Similarly, we have

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

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 A2.

An example of a nonsmooth regular environment is provided by the renewal model:

LEMMA A2: *Consider the renewal model of Example 2. Suppose that the inverses of the initial conditional distributions, $F_0^{-1}(\cdot | \theta_B)$ and $G_0^{-1}(\cdot | \theta_S)$, are differentiable functions of θ_B and θ_S , respectively, with uniformly bounded gradients i.e., $\exists d < \infty : \|\nabla_{\theta_B} F_0^{-1}\| \vee \|\nabla_{\theta_S} G_0^{-1}\| \leq d$. 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 $F_0(\cdot | \theta_B)$ and $G_0(\cdot | \theta_S)$ have strictly positive densities given any $\theta \in \Omega$, and the ratios $\frac{\|\nabla_{\theta_B} F_0(v | \theta_B)\|}{f_0(v | \theta_B)}$ and $\frac{\|\nabla_{\theta_S} G_0(c | \theta_S)\|}{g_0(c | \theta_S)}$ 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, ψ) 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

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

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 $\|\nabla_{\theta_B} \psi\| \leq \|\nabla_{\theta_B} F_0^{-1}\| \leq d$ 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 A1: *Every regular environment has the payoff-equivalence property.*

In the special case where each agent's initial private information is one dimensional, Proposition A1 follows from the results of PST. The extension to the multidimensional case presented here is novel. The proof, which combines the standard payoff-equivalence argument from static multidimensional 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 OF PROPOSITION A1:

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.²⁷

²⁷To 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 $U^\mu(b) - U^\eta(b) = k$.

Fix a smooth path $\alpha : [0, 1] \rightarrow \Omega_B \times \mathcal{V}_0$ 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 $W^\sigma(\lambda) := U^{\mu, \sigma}(\alpha(\lambda))$ and $W(\lambda) := W^{\sigma^*}(\lambda)$. 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 Theorem 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 $(\lambda, v_1, v_2, \dots, v_T)$, 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 utility function is clearly differentiable and equi-Lipschitz continuous in types in the sense of PST’s conditions U-D and U-ELC (see PST, section 2.1). Similarly, condition F-BE, which requires the expected discounted type to be finite conditional on any history, follows immediately from our assumption that $E[\sum_{t=0}^T \delta^t | V_t | | \theta_B, v_0]$ 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” and thus satisfies condition F-BIR. Since we will apply the envelope formula only in the initial period, it suffices to verify this for period 0 (i.e., that PST’s condition F-BIR holds for $s = 0$). Given initial type λ , future types are distributed as follows:

$$(A1) \quad \begin{aligned} v_1 &\sim F(\cdot | \alpha_v(\lambda); \alpha_\theta(\lambda)), \\ v_t &\sim F(\cdot | v_{t-1}; \alpha_\theta(\lambda)) \quad \text{for } t = 2, \dots, T. \end{aligned}$$

As V is regular, we can take a state representation (\mathcal{E}, Q, ψ) satisfying our Definition A2, and define the functions $z_1 : [0, 1] \times \mathcal{E} \rightarrow \mathbb{R}$, and $z_t : \mathcal{V} \times [0, 1] \times \mathcal{E} \rightarrow \mathbb{R}, t = 2, \dots, T$, by

$$z_1(\lambda, \varepsilon) = \psi(\alpha_\theta(\lambda), \alpha_v(\lambda), \varepsilon),$$

$$z_t(v, \lambda, \varepsilon) = \psi(\alpha_\theta(\lambda), v, \varepsilon) \quad \text{for } t = 2, \dots, T.$$

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

$$\partial_\lambda z_1 = \nabla_{(\theta, v)} \psi \cdot \alpha',$$

$$\partial_\lambda z_t = \nabla_\theta \psi \cdot \alpha'_\theta \quad \text{for } t = 2, \dots, T,$$

$$\partial_v z_t = \partial_v \psi \quad \text{for } t = 2, \dots, T.$$

Furthermore, given any constants b and d that satisfy Definition A2, we may put $K := (d + b) \max_\lambda \|\alpha'(\lambda)\|$ to obtain the bounds

$$|\partial_\lambda z_t| \leq K \quad \text{for } t = 1, \dots, T,$$

$$|\partial_v z_t| \leq b \quad \text{for } t = 2, \dots, T.$$

By construction, the collection (\mathcal{E}, Q, z) , where $z := (z_t)_{t=1}^T$, defines a state representation of the kernels (A1) in the sense of PST, section 2.²⁸ By their equation (4), 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.

²⁸Note that we can put $\mathcal{E}_t = \mathcal{E}$ and $G_t = Q$ for all t .

Finally, we show that the environment in Example 7 is regular and hence has the payoff-equivalence property.

LEMMA A3: *The environment in Example 7 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 $\bar{F}(v' | v)$ is a differentiable function of (v', v) with a density $\bar{f}(v' | v)$ 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 $(v', v, \theta_B) \in [0, 1]^3$,

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

Similarly, we have

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

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 A2 and the proof of Lemma A1.

B. Omitted Proofs

PROOF OF PROPOSITION 1:

The implication $(iii) \vee (iv) \Rightarrow (ii)$ follows immediately from Definitions 1–4 (see Remark 1). Hence, it suffices to show $(i) \Rightarrow (iii) \wedge (iv)$, and $(ii) \Rightarrow (i)$.

We establish $(i) \Rightarrow (iv)$ 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 a static Groves’ scheme 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 $p_{S,t} = \mathbf{1}_{\{v_t \geq c_t\}} v_t$.²⁹ 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 $U_0^\mu(\theta_B, v_0) = E[Y | \theta_B, v_0]$ and $\Pi_0^\mu(\theta_S, c_0) = E[Y | \theta_S, c_0]$ 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 $\frac{1 - \delta}{1 - \delta^{T+1}} E^\mu \left[\sum_{t=0}^T \delta^t (P_{B,t} + P_{S,t}) \right] = E[Y]$.

²⁹In our environment, this mechanism coincides with the team mechanism of Athey and Segal (2013) and the Dynamic Pivot mechanism of Bergemann and Välimäki (2010).

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

$$p_{B,0}^* := p_{B,0} - \frac{1 - \delta^{T+1}}{1 - \delta} \inf_{\theta'_B, v'_0} E[Y | \theta'_B, v'_0], \quad \text{and}$$

$$p_{S,0}^* := p_{S,0} - \frac{1 - \delta^{T+1}}{1 - \delta} \inf_{\theta'_S, c'_0} E[Y | \theta'_S, c'_0].$$

Denote the public mechanism so obtained by $\mu^* = (x^*, p^*)$. As we just subtracted constants, μ^* is E and EPIC. It is IR, since periods $t > 0$ are unaffected, and

$$(A2) \quad \inf_{\theta_B, v_0} U_0^{\mu^*}(\theta_B, v_0) = \inf_{\theta_B, v_0} \left(U_0^\mu(\theta_B, v_0) - \inf_{\theta'_B, v'_0} E[Y | \theta'_B, v'_0] \right)$$

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

and similarly for the seller. Note that

$$(A3) \quad \frac{1 - \delta}{1 - \delta^{T+1}} E^{\mu^*} \left[\sum_{t=0}^T \delta^t (P_{B,t} + P_{S,t}) \right] = \frac{1 - \delta}{1 - \delta^{T+1}} E^\mu \left[\sum_{t=0}^T \delta^t (P_{B,t} + P_{S,t}) \right]$$

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

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

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

We show then that (ii) \Rightarrow (i). Let $\eta = (x', q)$ be an E, IC, IR_0 , and BB_0 blind mechanism, and let $\mu^* = (x^*, p^*)$ be the public mechanism constructed above by adding participation fees to the repetition of the static Groves’ scheme. 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

$$0 \leq \inf_{v_0, \theta_B} U_0^\eta(v_0, \theta_B) + \inf_{c_0, \theta_S} \Pi_0^\eta(c_0, \theta_S) =$$

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

where the last equality follows by (A2). Thus, $U_0^\eta + \Pi_0^\eta \geq U_0^{\mu^*} + \Pi_0^{\mu^*}$, which in turn implies $E^\eta \left[\sum_{t=0}^T \delta^t (P_{B,t} + P_{S,t}) \right] \geq E^{\mu^*} \left[\sum_{t=0}^T \delta^t (P_{B,t} + P_{S,t}) \right]$ as $x' = x^*$.

But η is BB_0 by assumption, and hence

$$\begin{aligned} 0 &\geq \frac{1 - \delta}{1 - \delta^{T+1}} E^\eta \left[\sum_{t=0}^T \delta^t (P_{B,t} + P_{S,t}) \right] \\ &\geq \frac{1 - \delta}{1 - \delta^{T+1}} E^{\mu^*} \left[\sum_{t=0}^T \delta^t (P_{B,t} + P_{S,t}) \right] \\ &= E[Y] - \left(\inf_{v_0, \theta_B} E[Y | v_0, \theta_B] + \inf_{c_0, \theta_S} E[Y | c_0, \theta_S] \right), \end{aligned}$$

where the last equality follows by (A3). Thus (1) is satisfied.

It remains to establish (i) \Rightarrow (iii). Note that the public mechanism $\mu^* = (x^*, p^*)$ constructed above is PIC. Hence, by the “balancing trick” of Athey and Segal (2013, proposition 2), there exists a public mechanism $\bar{\mu}^* = (x^*, \bar{p}^*)$ that is PIC and BB. Since the allocation rule is unchanged, $\bar{\mu}^*$ is also E, and the pay-off-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 (A3) we then have

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

In particular, this implies that

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

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 (i) \Rightarrow (iii).

PROOF OF PROPOSITION 2:

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 (i) and (ii) of Definition 7, $E[Y] = E[Y']$. By inspection of (1) it thus suffices to show that

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

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), $F^{(t)}(\cdot | w)$ and $F'^{(t)}(\cdot | w)$ are dominated by $F^{(t)}(\cdot | v_0)$ and $F'^{(t)}(\cdot | v_0)$, 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 , $F^{(t)}(\cdot | w)$ dominates $F'^{(t)}(\cdot | w)$. To this end, note that by condition (iii) of Definition 7,

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

where Φ is the common invariant distribution. Since Φ dominates μ_w , and the chains are stochastically monotone, Φ dominates $F^{(t)}(\cdot | w)$ and $F'^{(t)}(\cdot | w)$.³⁰ Thus, we may dispense with the absolute value operator to get

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

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

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 $\inf_{v_0} E[Y | v_0] = \inf_{c_0} E[Y | c_0] = 0$. By symmetry, it is enough to consider $\inf_{v_0} E[Y | v_0]$. 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,

$$E[(V_t - C_t)^+ | v_0] \leq E[Z_t^+],$$

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 cumulative distribution function for $N(0, 1)$, we have by standard formulae,

$$\begin{aligned} \frac{E[Z_t^+]}{1 - \Phi\left(\frac{-\gamma_B^t(v_0 - m_B) - m}{\sigma}\right)} &= E[Z_t | Z_t \geq 0] \\ &= \gamma_B^t(v_0 - m_B) + m + \frac{\phi\left(\frac{-\gamma_B^t(v_0 - m_B) - m}{\sigma}\right)}{1 - \Phi\left(\frac{-\gamma_B^t(v_0 - m_B) - m}{\sigma}\right)}\sigma, \end{aligned}$$

³⁰To see this, note that the t -step distribution starting from the invariant distribution Φ is Φ itself.

and hence,

$$E[Y | v_0] \leq (1 - \delta) \sum_{t=0}^{\infty} \delta^t \left[\left(1 - \Phi \left(\frac{-\gamma_B^t(v_0 - m_B) - m}{\sigma} \right) \right) (\gamma_B^t(v_0 - m_B) + m) + \phi \left(\frac{-\gamma_B^t(v_0 - m_B) - m}{\sigma} \right) \sigma \right]$$

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

PROOF OF LEMMA 1:

Suppose that μ is an IC, BB_0 , IR_0 mechanism. Since the environment is smooth and μ is IC, PSTs Theorem 1 implies that $U_0(v_0) := U_0^\mu(v_0)$ is Lipschitz continuous in v_0 with derivative given by (omitting the known parameter θ in the notation)

$$(A4) \quad U_0'(v_0) = (1 - \delta) E^\mu \left[\sum_{t=0}^{\infty} (\delta \gamma_B)^t X_t | v_0 \right] \quad \text{a.e. } v_0.$$

By inspection, U_0 is nondecreasing. Thus, $\lim_{v \rightarrow -\infty} U_0(v)$ is well-defined, and for all v_0 ,

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

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

$$(A5) \quad 0 \leq \lim_{v \rightarrow -\infty} U_0(v) + \lim_{c \rightarrow \infty} \Pi_0(c) \\ = E^\mu \left[U_0(V_0) - \int_{-\infty}^{V_0} U_0'(r) dr + \Pi_0(C_0) - \int_{C_0}^{\infty} \Pi_0'(y) dy \right] \\ = E^\mu \left[U_0(V_0) - \frac{1 - F_0(V_0)}{f_0(V_0)} U_0'(V_0) + \Pi_0(C_0) - \frac{G_0(C_0)}{g_0(C_0)} \Pi_0'(C_0) \right] \\ \leq (1 - \delta) E^\mu \left[\sum_{t=0}^{\infty} \delta^t X_t \left(V_t - \frac{1 - F_0(V_0)}{f_0(V_0)} \gamma_B^t - C_t - \frac{G_0(C_0)}{g_0(C_0)} \gamma_S^t \right) \right],$$

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 (A4) (and its analog for the seller).

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 (2013, proposition 2), we may take μ^{**} to be BB. Since μ^{**} satisfies (3), IR_0 follows by reversing the steps in (A5) 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 in the sense of Corollary 1 of PST, there exist transfers p such that $\mu = (x^{**}, p)$ is EPIC (see PST, Remark 5). 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 Groves' scheme from the proof of Proposition 1, adjusted to account for the wedge in (5). Hence, we have IR for periods $t > 0$. Suppose then that we add constant participation fees π_B and π_S in period 0 such that $\pi_B + \pi_S = \frac{1 - \delta}{1 - \delta^T} E^\mu \left[\sum_{t=0}^T \delta^t (P_{B,t} + P_{S,t}) \right]$. 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 (A5) because x^{**} satisfies (3).

PROOF OF PROPOSITION 6:

We start by establishing the first claim. To this end, fix $\Delta > 0$, and note that then $r \rightarrow 0$ if and only if $\delta \rightarrow 1$, so we may work with the latter.

For $b \in R_+$, define the allocation rule x^b by

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

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

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

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

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

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

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

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

We have

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

$$\leq \frac{1 - \delta}{1 - \delta\gamma_S} E \left[\frac{1 - F_0(V_0)}{f_0(V_0)} \middle| V_0 \geq -b \right] + \frac{1 - \delta}{1 - \delta\gamma_S} E \left[\frac{G_0(C_0)}{g_0(C_0)} \middle| C_0 \leq b \right],$$

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, $y^{**}\left(\frac{-\log \delta}{\Delta}, \Delta\right) \geq y(\delta, b) > y^* - \varepsilon$ 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

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

and denote the expected information rents by

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

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 $\lim_{\Delta \rightarrow 0} y^{**}(\Delta) \leq y^*$ exists. Furthermore, by Proposition 5, we have $y^{**}(\Delta) = y(x^{**}(\Delta), \Delta) = r(x^{**}(\Delta), \Delta)$ for all $\Delta > 0$, and, hence, $\lim_{\Delta \rightarrow 0} r(x^{**}(\Delta), \Delta) = \lim_{\Delta \rightarrow 0} y^{**}(\Delta)$.

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

For $q \geq 0$, define the allocation rule x^q by setting $x_t^q = 1$ if and only if

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

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

CLAIM A1: 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 $x_t^{q'} \leq x_t^q$ for all t , and the second by noting that $x_t^{\lambda(\Delta)} \leq x_t^{**}$ for all t . This establishes the claim.

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

CLAIM A2: If $\liminf_{\Delta \rightarrow 0} \lambda(\Delta) = 0$, then $\lim_{\Delta \rightarrow 0} r(x^{**}(\Delta), \Delta) \geq 2y^*$.

To establish Claim A2, note that if $\liminf_{\Delta \rightarrow 0} \lambda(\Delta) = 0$, then there exists a monotone sequence (Δ_n) with $\Delta_n \rightarrow 0$ such that $\lambda_n := \lambda(\Delta_n)$ defines a monotone sequence (λ_n) with $\lambda_n \rightarrow 0$. By Claim A1, 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 ,

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

Note that

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

where Φ denotes the cumulative distribution function 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,

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

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 $\lim_{n \rightarrow \infty} r(x^{**}(\Delta_n), \Delta_n) \geq \lim_{k \rightarrow \infty} \lim_{n \rightarrow \infty} r(x^{\lambda_k}, \Delta_n) = 2y^*$.

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 S2 of PST. For statement 1 we may use the balancing argument of Athey and Segal (2013, 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

- Abreu, Dilip, Paul Milgrom, and David Pearce.** 1991. "Information and Timing in Repeated Partnerships." *Econometrica* 59 (6): 1713–33.
- Athey, Susan, and David A. Miller.** 2007. "Efficiency in repeated trade with hidden valuations." *Theoretical Economics* 2 (3): 299–354.
- Athey, Susan, and Ilya Segal.** 2007. "Designing Dynamic Mechanisms for Dynamic Bilateral Trading Games?" *American Economic Review* 97 (2): 131–36.
- Athey, Susan, and Ilya Segal.** 2013. "An Efficient Dynamic Mechanism." *Econometrica* 81 (6): 2463–85.
- Baron, David P., and David Besanko.** 1984. "Regulation and Information in a Continuing Relationship." *Information Economics and Policy* 1 (3): 267–302.
- Battaglini, Marco.** 2005. "Long-Term Contracting with Markovian Consumers." *American Economic Review* 95 (3): 637–58.
- Battaglini, Marco, and Rohini Lamba.** 2012. "Optimal Dynamic Contracting." http://econ.as.nyu.edu/docs/IO/31409/optm_dyn_contract.pdf.
- Bergemann, Dirk, and Juuso Välimäki.** 2010. "The Dynamic Pivot Mechanism." *Econometrica* 78 (2): 771–89.
- Boleslavsky, Raphael, and Maher Said.** 2013. "Progressive Screening: Long-Term Contracting with a Privately Known Stochastic Process." *Review of Economic Studies* 80 (1): 1–34.
- Courty, Pascal, and Hao Li.** 2000. "Sequential Screening." *Review of Economic Studies* 67 (4): 697–717.
- Daley, D. J.** 1968. "Stochastically Monotone Markov Chains." *Probability Theory and Related Fields* 10 (4): 305–17.
- d'Aspremont, Claude, and Louis-André Gérard-Varet.** 1979. "Incentives and Incomplete Information." *Journal of Public Economics* 11 (1): 25–45.

- Escobar, Juan F., and Juuso Toikka.** 2013. "Efficiency in Games with Markovian Private Information." *Econometrica* 81 (5): 1887–1934.
- Eső, Péter, and Balász Szentes.** 2007. "Optimal Information Disclosure in Auctions and the Handicap Auction." *Review of Economic Studies* 74 (3): 705–31.
- Fudenberg, Drew, David Levine, and Eric Maskin.** 1994. "The Folk Theorem with Imperfect Public Information." *Econometrica* 62 (5): 997–1039.
- Harrison, J. Michael.** 1985. *Brownian Motion And Stochastic Flow Systems*. New York: John Wiley and Sons.
- Holmström, Bengt.** 1979. "Groves' Scheme on Restricted Domains." *Econometrica* 47 (5): 1137–44.
- Karamata, J.** 1930. "Über die Hardy-Littlewoodschen Umkehrungen des Abelschen Stetigkeitssatzes." *Mathematische Zeitschrift* 32 (1): 319–20.
- Krishna, Vijay, and Motty Perry.** 2000. "Efficient Mechanism Design." <http://grizzly.la.psu.edu/~vkrishna/papers/vcg20.pdf>.
- Makowski, Louis, and Claudio Mezzetti.** 1994. "Bayesian and Weakly Robust First Best Mechanisms: Characterizations." *Journal of Economic Theory* 64 (2): 500–519.
- Miller, David A.** 2012. "Robust Collusion with Private Information." *Review of Economic Studies* 79 (2): 778–811.
- Myerson, Roger B.** 1986. "Multistage Games with Communication." *Econometrica* 54 (2): 323–58.
- Myerson, Roger, and Mark Satterthwaite.** 1983. "Efficient Mechanisms for Bilateral Trading." *Journal of Economic Theory* 29 (2): 265–81.
- Pavan, Alessandro, Ilya Segal, and Juuso Toikka.** 2014. "Dynamic Mechanism Design: A Myersonian Approach." *Econometrica* 82 (2): 601–53.
- Segal, Ilya, and Michael D. Whinston.** 2011. "A Simple Status Quo that Ensures Participation (with Application to Efficient Bargaining)." *Theoretical Economics* 6 (1): 109–25.
- Williams, David.** 1991. *Probability with Martingales*. Cambridge, U.K.: Cambridge University Press.
- Williams, Steven R.** 1999. "A characterization of efficient, Bayesian incentive compatible mechanisms." *Economic Theory* 14 (1): 155–80.