

Designing Efficient Mechanisms for Dynamic Bilateral Trading Games

Susan Athey and Ilya Segal*

This paper studies the problem of allocating a good among two players in each period of an infinite-horizon game. The players' valuations in each period are private information, and the valuations change over time. We analyze two special cases for the dynamics of valuations: “serially correlated valuations,” where players' valuations are exogenous but serially correlated, and “learning by doing,” where a player's past consumption improves his current distribution of valuations but his valuations are otherwise uncorrelated.

We analyze conditions under which there exists an efficient, Bayesian incentive-compatible (BIC), individually rational (IR), budget-balanced (BB) mechanism, when the mechanism designer has commitment power. We consider IR constraints where agents have the option to (permanently) exit the mechanism in each period. This captures situations where there are restrictions on the nature of long-term contracts agents can sign, such as at-will employment contracts, as well as restrictions on posting bonds. In addition, the tools of mechanism design are often used to model decentralized games, and period-by-period IR constraints may correspond to the need to deter certain types of deviations in a dynamic game. BB is a desirable property when there is no individual without private information who can serve as a source or sink of funds; we may also be interested in efficiency for a particular group of agents (e.g. colluding firms), so that “burning money” violates efficiency within the group.

In a static setting, Claude d'Aspremont and Louis-Andre Gerard-Varet (1979) (AGV) constructed an efficient, BIC, BB mechanism, in which given prior beliefs, each player receives the expected value (over opponent types) of opponent utilities. However, in a dynamic

model, today’s reports and allocations influence opponent prior beliefs in future periods, and so incentives for truthtelling are undermined.

In a general dynamic model, Susan Athey and Ilya Segal (2006) (AS) construct an efficient mechanism satisfying BIC and BB, and they provide sufficient conditions for IR to hold when players are very patient and players’ current reports have a vanishing impact on expected utility in the distant future. Here, we apply their construction to the bilateral trading game, interpreting the transfers and examining comparative statics. We also analyze IR constraints, finding sufficient conditions for IR constraints to hold for moderate patience or when there is a non-trivial long-run impact of player’s reports in a given period.

I. The Model

Time is indexed by $t = 1, \dots, \infty$. There are two players, the “buyer” (b) and the “seller” (s). Each player $i \in \{b, s\}$ has a privately observed valuation for consuming a single indivisible object in each period (his “type” in period t), denoted $\theta_{i,t} \in \Theta_{i,t}$, where $\Theta_{i,t} \subset \mathbb{R}_+$ is finite, and we let $\Theta_t = \Theta_{b,t} \times \Theta_{s,t}$. We consider several alternative models of the stochastic process over player types, denoted $(\tilde{\theta}_t)_{t=1}^\infty$, as detailed below. The allocation is denoted $(x_{b,t}, x_{s,t}) \in \{(0, 1), (1, 0)\}$. Payoffs are $x_{i,t}\theta_{i,t} + y_{i,t}$, where $y_{i,t} \in \mathbb{R}$ is the transfer to player i in period t .

For a sequence $(\theta_t)_{t=1}^\infty$, we use the notation $\theta^t = (\theta_1, \dots, \theta_t)$ and $\theta = (\theta_t)_{t=1}^\infty$.

We assume the existence of a mechanism designer who receives reports from the players in each period (which are publicly observed by both players), and can commit in advance to a history-contingent allocation and transfer plan. The allocation to player i in period t is denoted $(\chi_{b,t}, \chi_{s,t}) : \Theta^t \rightarrow \{(0, 1), (1, 0)\}$, while the transfers, restricted to satisfy budget balance, are $\psi_{b,t} : \Theta^t \rightarrow \mathbb{R}$ and $\psi_{s,t} = -\psi_{b,t}$. The strategy of each player i specifies player i ’s

reports in all periods t as functions of the history of his true types and all reports. We say that the strategy is truthtelling if the player reports his true type for all histories.

We consider “ex post” IR constraints within each period: a player can exit after reports are made but before allocations and transfers are made. For patient players, interim and ex post IR constraints are similar, so we focus on the more stringent constraint. Once either player exits, the seller always keeps the object. We require that ex post IR constraints hold for all possible histories of true types and reports, in order to deter deviations where a player engages in a series of misreports and plans to exit for some realizations of future types.

II. Serially Correlated Valuations

Suppose that buyer and seller types are serially correlated, but that they are independent across players and the evolution of types is exogenous to the history of allocations. This captures the idea that a firm’s production technology and capabilities evolve slowly over time. The types follow a first-order Markov process (so that today’s type distribution depends only on yesterday’s type), and to simplify exposition we assume that the types are affiliated over time. The efficient policy allocates to the highest value player and so can be written $\chi_b^*(\theta_t) = 1\{\theta_{b,t} > \theta_{s,t}\}$. To simplify notation we let $\mathbb{E}_t [g(\tilde{\theta}_\tau)|\theta_t]$ denote the expectation of g over $\tilde{\theta}_\tau|\tilde{\theta}_t = \theta_t$, and $\mathbb{E}_{t,t-1} [g(\tilde{\theta}_\tau)|\theta_{i,t}^t, \theta_{-i}^{t-1}]$ denotes the expectation over $\tilde{\theta}_\tau|(\tilde{\theta}_{i,t}, \tilde{\theta}_{-i,t-1}) = (\theta_{i,t}, \theta_{-i,t-1})$. Following AS, the transfers are constructed as follows, given a constant $K \in \mathbb{R}$ that is used to transfer utility between the buyer and seller:

$$\begin{aligned} \psi_{b,t}(\theta_t, \theta_{t-1}) &= -\psi_{s,t}(\theta_t, \theta_{t-1}) = \gamma_{b,t}(\theta_{b,t}, \theta_{t-1}) - \gamma_{s,t}(\theta_{s,t}, \theta_{t-1}) + K, \text{ where} \\ \gamma_{i,t}(\theta_{i,t}, \theta_{t-1}) &= \sum_{\tau=t}^{\infty} \delta^{\tau-t} \left(\mathbb{E}_{t,t-1} \left[\tilde{\theta}_{-i,\tau} \chi_{-i}^*(\tilde{\theta}_\tau) | \theta_{i,t}, \theta_{-i,t-1} \right] \right) - \sum_{\tau=t}^{\infty} \delta^{\tau-t} \left(\mathbb{E}_{t-1} \left[\tilde{\theta}_{-i,\tau} \chi_{-i}^*(\tilde{\theta}_\tau) | \theta_{t-1} \right] \right). \end{aligned}$$

The buyer’s incentive payment $(\gamma_{b,t})$ is the difference between expected discounted seller

surplus given the buyer's report, and the expected seller surplus given reports in past periods.

The buyer incentive payment is large when serial correlation is high, and when yesterday's type was high and today's is low. Note that if (in an extension of this model) players observed only partially informative signals about their valuations prior to the allocation decision in each period, an increase in the accuracy of the signals would increase the overall variability of the signals as well as the magnitude of transfers, since each period's information would have a bigger effect on expectations about the future.

We show that there exists a truthtelling BNE. Consider BIC for the buyer. Because we impose BB, the buyer must not wish to misreport in order to manipulate the seller's future incentive payments; however, the seller's future incentive payments depend on the seller's beliefs about the buyer, which are in turn functions of today's reports.

First, we show that from the buyer's perspective, the incentive payments for the seller in each future period $\tau \geq t$, $\gamma_{s,\tau}$, have zero expectation no matter what reporting strategy the buyer uses, provided that the seller uses a truthful strategy. This follows because the seller's incentive payments give the seller the *change* in expected buyer utility due to the seller's current period report, but the buyer's expectation of this change is always zero. Formally, for every t and history of reports $\hat{\theta}_{t-1}$,

$$\mathbb{E}_{t-1} \left[\gamma_{s,t} \left(\tilde{\theta}_{s,t}, \hat{\theta}_{t-1} \right) \middle| \hat{\theta}_{s,t-1} \right] = 0. \quad (1)$$

For each t and $\theta_{s,t-1}$, let $(\tilde{\theta}_{s,t}, \tilde{\theta}_{s,t-1})$ be equal in distribution to $(\tilde{\theta}_{s,t}, \tilde{\theta}_{s,t-1})$. Equation (1) holds because for each $\tau \geq t$, the Law of Iterated Expectations (LIE) implies

$$\mathbb{E}_{t-1} \left[\mathbb{E}_{t,t-1} \left[\tilde{\theta}_{b,\tau} \chi_b^*(\tilde{\theta}_\tau) \middle| \tilde{\theta}_{s,t}, \hat{\theta}_{b,t-1} \right] \middle| \hat{\theta}_{s,t-1} \right] = \mathbb{E}_{t-1} \left[\tilde{\theta}_{b,\tau} \chi_b^*(\tilde{\theta}_\tau) \middle| \hat{\theta}_{t-1} \right].$$

Since this holds for all possible histories, it must hold when we take expectations given the buyer's period t beliefs, even if the buyer does not follow a truthful reporting strategy.

Second, we show that in any period t' , the expected value of $\sum_{\tau=t'}^{\infty} \delta^{\tau-t'} \gamma_{b,\tau}$ provides the correct incentives for the buyer. By the one-stage deviation principle, it suffices to verify that for any history of reports $\hat{\theta}^{t'-1}$ and true types for the buyer $\theta_b^{t'}$, the buyer has the incentive to report truthfully in period t' when he anticipates truthful reporting in future periods. Note that $\hat{\theta}_{b,t'}$ only enters the first term of $\gamma_{b,t'}$ and the second term of $\delta\gamma_{b,t'+1}$. We proceed by establishing that the first term of $\gamma_{b,t'}$ plus the buyer's expectation of the second term of $\delta\gamma_{b,t'+1}$ equals

$$\mathbb{E}_{t'-1} \left[\tilde{\theta}_{s,t'} \chi_s^*(\hat{\theta}_{b,t'}, \tilde{\theta}_{s,t'}) \mid \hat{\theta}_{s,t'-1} \right]. \quad (2)$$

Critically, the second term of $\delta\gamma_{b,t'+1} \left(\tilde{\theta}_{b,t'+1}, \left(\hat{\theta}_{b,t'}, \tilde{\theta}_{s,t'} \right) \right)$ does not depend on $\tilde{\theta}_{b,t'+1}$, so that the buyer's expectation of it conditional on $(\theta_{b,t'}, \hat{\theta}_{s,t'-1})$ does not depend on $\theta_{b,t'}$ and can be written as

$$\begin{aligned} & - \sum_{\tau=t'+1}^{\infty} \delta^{\tau-t'} \left(\mathbb{E}_{t'-1} \left[\mathbb{E}_{t'} \left[\tilde{\theta}_{s,\tau} \chi_s^*(\tilde{\theta}_{b,\tau}, \tilde{\theta}_{s,\tau}) \mid \hat{\theta}_{b,t'}, \tilde{\theta}_{s,t'} \right] \mid \hat{\theta}_{s,t'-1} \right] \right) \\ & = - \sum_{\tau=t'+1}^{\infty} \delta^{\tau-t'} \left(\mathbb{E}_{t',t'-1} \left[\tilde{\theta}_{s,\tau} \chi_s^*(\tilde{\theta}_{b,\tau}, \tilde{\theta}_{s,\tau}) \mid \hat{\theta}_{b,t'}, \hat{\theta}_{s,t'-1} \right] \right), \end{aligned} \quad (3)$$

using LIE. Adding (3) to the first term of $\gamma_{b,t'} \left(\hat{\theta}_{b,t'}, \hat{\theta}_{s,t'-1} \right)$, yields (2). Since (2) represents the expected externality of the buyer's report on the seller, we conclude that the anticipation of future transfers induces the buyer to report truthfully. Notice that this argument exploits the fact that the first term of $\gamma_{b,t'}$ and the second term of $\delta\gamma_{b,t'+1}$ differ in that the latter incorporates the arrival of new information about seller types, and the buyer's private information about past misreports does not affect his beliefs about the seller's information.

Now consider IR. Transfers potentially grow without bound as δ approaches 1, so that we cannot necessarily appeal to arbitrary patience to satisfy IR constraints. Which player's IR constraint is most stringent in period t depends on history. However, the only degree of freedom the transfers offer in transferring utility across players is a fixed constant K (if it varied with history, it would affect incentives). Thus, we look for a K that allows IR to be satisfied in the “worst-case” scenario for the buyer, and separately in the worst-case scenario for the seller. To simplify, suppose seller valuations are serially independent. Let

$$\begin{aligned}\varphi(\theta_t, \theta_{b,t-1}) &= \mathbb{E} \left[\tilde{\theta}_{s,t} \chi_s^*(\theta_{b,t}, \tilde{\theta}_{s,t}) \right] - \mathbb{E}_{t-1} \left[\tilde{\theta}_{s,t} \chi_s^*(\tilde{\theta}_{b,t}, \tilde{\theta}_{s,t}) \middle| \theta_{b,t-1} \right] \\ &\quad - \left(\mathbb{E}_{t-1} \left[\tilde{\theta}_{b,t} \chi_b^*(\tilde{\theta}_{b,t}, \theta_{s,t}) \middle| \theta_{b,t-1} \right] - \mathbb{E}_{t-1} \left[\tilde{\theta}_{b,t} \chi_b^*(\tilde{\theta}_{b,t}, \tilde{\theta}_{s,t}) \middle| \theta_{b,t-1} \right] \right), \\ \bar{\varphi} &= \max_{\theta_t, \theta_{b,t-1}} \varphi(\theta_t, \theta_{b,t-1}), \text{ and } \underline{\varphi} = \min_{\theta_t, \theta_{b,t-1}} \varphi(\theta_t, \theta_{b,t-1}).\end{aligned}$$

The function φ is the component of period t 's transfer that corresponds to period t payoffs, and $\underline{\varphi}$ and $\bar{\varphi}$ are the smallest and largest possible values of these. Let $\bar{\theta}_i = \max_t \max \Theta_{i,t}$, and $\underline{\theta}_i = \min_t \min \Theta_{i,t}$.

Proposition 1 *Suppose seller valuations are serially independent. A sufficient condition for IR to hold in each period is:*

$$\begin{aligned}& \sum_{\tau=t+1}^{\infty} \delta^{\tau-t} \mathbb{E}_t \left[\tilde{\theta}_{b,\tau} \chi_b^*(\tilde{\theta}_\tau) \middle| \underline{\theta}_b \right] - \sum_{\tau=t+1}^{\infty} \delta^{\tau-t} \mathbb{E}_t \left[\tilde{\theta}_{s,\tau} \chi_s^*(\tilde{\theta}_\tau) \middle| \bar{\theta}_b \right] - \bar{\theta}_s \quad (4) \\ & \geq \bar{\varphi} - \underline{\varphi} + \sum_{\tau=t+1}^{\infty} \delta^{\tau-t} \left(\begin{aligned} & \mathbb{E}_{t-1} \left[\tilde{\theta}_{s,\tau} \chi_s^*(\tilde{\theta}_\tau) \middle| \underline{\theta}_b \right] - \mathbb{E}_{t-1} \left[\tilde{\theta}_{s,\tau} \chi_s^*(\tilde{\theta}_\tau) \middle| \bar{\theta}_b \right] \\ & + \mathbb{E}_t \left[\tilde{\theta}_{s,\tau} \chi_s^*(\tilde{\theta}_\tau) \middle| \underline{\theta}_b \right] - \mathbb{E}_t \left[\tilde{\theta}_{s,\tau} \chi_s^*(\tilde{\theta}_\tau) \middle| \bar{\theta}_b \right] \end{aligned} \right).\end{aligned}$$

The left-hand side of (4) is the sum of the worst-case buyer utility from consumption and the worst case of the difference between the seller utility from consumption and his outside option. The right-hand side is the sum of worst-case transfers.

When both players' types are fully persistent, we essentially have a static model, in which, by the logic of the theorem of Roger Myerson and Mark Satterthwaite (1983), we know that IRs in general cannot be satisfied (they require continuous types, but the result holds for many cases of discrete types as well).

With limited persistence, IRs are easier to satisfy. AS show that so long as the type process has a unique ergodic set (and thus a unique invariant distribution), IR constraints are satisfied in this model for δ close enough to 1. Consider two examples with moderate patience. First, suppose that there are two possible "states" at the start of period t , $z_t = 0$ and $z_t = 1$, and a constant c , such that $\tilde{z}_t = 1\{\tilde{\theta}_{t-1} > c\}$. Let $p_z = \Pr_{t-1}(\tilde{z}_t = 0|z)$ and $\pi_{i,j,z} = \mathbb{E}_t \left[\tilde{\theta}_{i,t} \chi_j^*(\tilde{\theta}_t) | z \right]$. Then, (4) becomes

$$\begin{aligned} & \frac{\delta}{1-\delta} (p_0 \pi_{b,b,0} + (1-p_0) \pi_{b,b,1} - (p_1 \pi_{s,b,0} + (1-p_1) \pi_{s,b,1}) - \delta(\pi_{b,b,0} - \pi_{s,b,1})(p_0 - p_1)) \quad (5) \\ & \geq \delta (1 + p_0 - p_1) (p_0 - p_1) (\pi_{s,s,0} - \pi_{s,s,1}) + (\bar{\theta}_s + \bar{\varphi} - \underline{\varphi}) (1 - \delta(p_0 - p_1)). \end{aligned}$$

Notice that so long as the impact of states on distributions is small enough and $p_0 - p_1$ is small enough, the term in parentheses on the left-hand side is positive, and thus the left-hand side grows without bound as δ approaches 1, while the right-hand side remains bounded. For particular parameter values, we can find a critical discount factor such that (5) holds.

In a second example, suppose $\Theta_t \subset [0, 1]^2$ for all t . Suppose there exist $\theta'_b, \theta''_b \in (0, 1)$ such that for all $\theta_{b,t} \in [\theta''_b, 1]$, $\Pr_t(\tilde{\theta}_{b,t+1} = 1 | \theta_{b,t}) = 1$, and for all $\theta_{b,t} \in [0, \theta'_b]$, $\Pr_t(\tilde{\theta}_{b,t+1} = 0 | \theta_{b,t}) = 1$. The left-hand side of (4) is negative, while the right-hand side is positive. The inequality will be impossible to satisfy.

We emphasize that (4) is sufficient but not necessary. In a model of repeated trade with serially independent types, Susan Athey and David Miller (2006) show how reallocating

transfers between states where IR constraints are binding and those where IR constraints do not bind, while preserving the expected transfers for each player, can reduce the critical discount factor required for sustaining efficient trade. We conjecture that the qualitative conclusions about when (4) can be satisfied will also apply to the question of whether IR constraints can be satisfied with any transfers.

III. Learning by Doing

Consider the following model of learning by doing. There is a finite set of possible states, $Z \subset \mathbb{R}$, with highest and lowest elements \bar{Z} and \underline{Z} . These states represent the stock of learning: in each period t , “nature” selects period t types using a probability distribution that depends on the state but not directly on time, and player types are independent conditional on the state z_t . The states evolve according to $z_t = \sigma(z_{t-1}, x_{b,t})$. (To simplify exposition we restrict attention to deterministic transitions). We assume that higher states imply weakly higher (by First Order Stochastic Dominance) valuations for the buyer and weakly lower valuations for the seller, and that σ is weakly monotone in z_{t-1} . This is a Markov decision problem, and there exists a efficient policy of the Markov form $(\chi_b^*, \chi_s^*) : \Theta_t \times Z \rightarrow \{(0, 1), (1, 0)\}$, which in turn induces a Markov process over states. It is well known that there exists a “Blackwell” policy, χ^B , that is optimal for all δ sufficiently high.

In general, the efficient policy may be biased away from the static efficient allocation depending on the current state. In the special case where $z_t = x_{t-1}$, the efficient policy is stationary, since history is irrelevant after types are realized in period t . If, in addition, the two players are ex ante symmetric and learning by doing is also symmetric, the efficient policy is just the static efficient policy.

Fix δ , with efficient policy (χ_b^*, χ_s^*) . Let $(\tilde{z}_\tau^*, \tilde{\theta}_\tau^*) | \tilde{z}_t = z_t$ be the random vector equal to the

state and the type in period τ when the state in t is z_t and when the efficient policy is used from t to $\tau - 1$. To conserve notation, for a function g we write $\mathbb{E}_t \left[g(\tilde{\theta}_\tau^*, \tilde{z}_\tau) | z_t \right]$ to indicate the expected value of g with respect to $(\tilde{z}_\tau^*, \tilde{\theta}_\tau^*) | z_t$. Transfers in period t are:

$$\begin{aligned} \psi_{b,t}(\theta_t, z_t) &= -\psi_{s,t}(\theta_t, z_t) = \gamma_{b,t}(\theta_{b,t}, z_t) - \gamma_{s,t}(\theta_{s,t}, z_t) + K, \text{ where } K \in \mathbb{R}, \text{ and} \\ \gamma_{i,t}(\theta_{i,t}, z_t) &= \mathbb{E}_t \left[\tilde{\theta}_{-i,t} \chi_{-i}^*(\tilde{\theta}_{-i,t}, \theta_{i,t}; z_t) | z_t \right] - \mathbb{E}_t \left[\tilde{\theta}_{-i,t} \chi_{-i}^*(\tilde{\theta}_t; z_t) | z_t \right] \\ &+ \begin{pmatrix} \Pr_t \left(\chi_{-i}^*(\tilde{\theta}_{-i,t}, \theta_{i,t}; z_t) = 1 | z_t \right) \\ - \Pr_t \left(\chi_{-i}^*(\tilde{\theta}_t; z_t) = 1 | z_t \right) \end{pmatrix} \cdot \sum_{\tau=t+1}^{\infty} \delta^{\tau-t} \begin{pmatrix} \mathbb{E}_{t+1} \left[\tilde{\theta}_{-i,\tau} \chi_{-i}^*(\tilde{\theta}_\tau^*; \tilde{z}_\tau^*) | \sigma(z_t, 1) \right] \\ - \mathbb{E}_{t+1} \left[\tilde{\theta}_{-i,\tau} \chi_{-i}^*(\tilde{\theta}_\tau^*; \tilde{z}_\tau^*) | \sigma(z_t, 0) \right] \end{pmatrix}. \end{aligned}$$

The incentive payment $\gamma_{i,t}$ incorporate the change, starting from state z_t , to the expectation of player $-i$'s discounted payoffs that results from player i 's period- t report. The results of AS imply that with these transfers, truthtelling is a BNE. The arguments are similar to those in the last section, appropriately modified to account for the fact that different reporting strategies lead to different distributions over future states.

Now consider IR. Each player's incentive payment is large and positive (negative) when the player's type is lower (higher) than expected, and when period t 's allocation can have (through its effect on z_{t+1}) a large impact on expected future utility for the opponent. AS show that the following is sufficient for IR constraints to hold for δ sufficiently close to 1:

Condition M For the Blackwell policy χ^B , the induced Markov process over states has a unique ergodic set (with a possibly empty set of transient states).

There will be a unique ergodic distribution if, for example, $\chi^*(\theta)$ can be either zero or one with positive probability for all histories, and if σ is strictly increasing in $x_{b,t}$. We now

consider sufficient conditions for IR to be satisfied for a fixed δ or if Condition M fails. Define (where $\bar{\theta}_s$ is the highest seller type)

$$\begin{aligned}\varphi(\theta_t, z_t) &= \mathbb{E}_t \left[\tilde{\theta}_{s,t} \chi_s^*(\theta_{b,t}, \tilde{\theta}_{s,t}; z_t) \middle| z_t \right] - \mathbb{E}_t \left[\tilde{\theta}_{s,t} \chi_s^*(\tilde{\theta}_t; s_t) \middle| z_t \right] \\ &\quad - \left(\mathbb{E}_t \left[\tilde{\theta}_{b,t} \chi^*(\tilde{\theta}_{b,t}, \theta_{s,t}; z_t) \middle| z_t \right] - \mathbb{E}_t \left[\tilde{\theta}_{b,t} \chi^*(\tilde{\theta}_{b,t}, \tilde{\theta}_{s,t}; z_t) \middle| z_t \right] \right), \\ \underline{\varphi} &= \min_{\theta_t, z_t} \varphi(\theta_t, z_t), \quad \bar{\varphi} = \max_{\theta_t, z_t} \varphi(\theta_t, z_t),\end{aligned}$$

$$\begin{aligned}\kappa_i(z_t, x_{b,t}) &= \sum_{\tau=t+1}^{\infty} \delta^{\tau-t} \mathbb{E}_{t+1} \left[\tilde{\theta}_{i,\tau} \chi_{i,\tau}^*(\tilde{\theta}_\tau^*; \tilde{z}_\tau) \middle| \sigma(z_t, x_{b,t}) \right] - 1\{i = s\} \left(\sum_{\tau=t+1}^{\infty} \delta^{\tau-t} \mathbb{E}_{t+1} \left[\tilde{\theta}_{s,\tau} \middle| \sigma(z_t, x_{b,t}) \right] \right) \\ &\quad - \Pr \left(\chi_{-i}^*(\tilde{\theta}_t; z_t) = 1 \middle| z_t \right) \cdot \sum_{\tau=t+1}^{\infty} \delta^{\tau-t} \left(\begin{array}{c} \mathbb{E}_{t+1} \left[\tilde{\theta}_{-i,\tau} \chi_{-i}^*(\tilde{\theta}_\tau^*; \tilde{z}_\tau) \middle| \sigma(z_t, 1) \right] \\ - \mathbb{E}_{t+1} \left[\tilde{\theta}_{-i,\tau} \chi_{-i}^*(\tilde{\theta}_\tau^*; \tilde{z}_\tau) \middle| \sigma(z_t, 0) \right] \end{array} \right) \\ &\quad - \Pr \left(\chi_{-i}^*(\tilde{\theta}_t; z_t) = 0 \middle| z_t \right) \cdot \sum_{\tau=t+1}^{\infty} \delta^{\tau-t} \left(\begin{array}{c} \mathbb{E}_{t+1} \left[\tilde{\theta}_{i,\tau}^* \chi_{i,t}^*(\tilde{\theta}_\tau^*; \tilde{z}_\tau) \middle| \sigma(z_t, 1) \right] \\ - \mathbb{E}_{t+1} \left[\tilde{\theta}_{i,\tau}^* \chi_{i,t}^*(\tilde{\theta}_\tau^*; \tilde{z}_\tau) \middle| \sigma(z_t, 0) \right] \end{array} \right),\end{aligned}$$

The expressions $\underline{\varphi}$ and $\bar{\varphi}$ are analogous to those in the serially correlated valuations model, while $\kappa_i(z_t, x_{b,t})$ contains all terms in a player's expected utility that depend on patience, representing the difference between expected discounted future payoffs and the portion of transfers corresponding to future periods in state z_t . The following result provides sufficient conditions for there to exist a K such that the worst-case buyer expected utility is greater than K and the worst-case seller utility is greater than $-K$.

Proposition 2 *A sufficient condition for IR to hold in each period is*

$$\min_{z_t, x_{b,t}} \kappa_s(z_t, x_{b,t}) + \min_{z_t, x_{b,t}} \kappa_b(z_t, x_{b,t}) \geq \bar{\theta}_s + \bar{\varphi} - \underline{\varphi}. \quad (6)$$

When will (6) be satisfied? Condition (M) implies that from every initial state, convergence to the ergodic distribution occurs at a geometric rate. For any two states z_t, z'_t and

any x_{bt}, x'_{bt} combining the first two terms of $\kappa_b(z_t, x_{b,t})$ and $\kappa_s(z'_t, x'_{b,t})$ yields

$$\sum_{\tau=t+1}^{\infty} \delta^{\tau-t} \left(\mathbb{E}_{t+1} \left[\tilde{\theta}_{b,\tau} \chi_{b,\tau}^*(\tilde{\theta}_\tau^*; \tilde{z}_\tau) \middle| \sigma(z_t, x_{b,t}) \right] - \mathbb{E}_{t+1} \left[\tilde{\theta}_{s,\tau} \chi_{b,\tau}^*(\tilde{\theta}_\tau^*; \tilde{z}_\tau) \middle| \sigma(z'_t, x'_{b,t}) \right] \right), \quad (7)$$

which would be the gains from efficient trade if there was no effect of the state. Since Condition (M) guarantees that the long-run effect of z_{t+1} is negligible in expectation, for high enough patience the sum will be large. On the other hand, under condition (M), the last two terms of κ_s and κ_b converge to a finite bound as δ increases (the long-run impact of today's decision is negligible; see AS for details).

However, when Condition (M) fails, (6) may not hold. Suppose that there are two ergodic sets of states. For example, suppose that under the Blackwell policy there is “increasing dominance”: once one player gets sufficiently far ahead, it is efficient to continue to allocate to that player often enough to keep him far ahead. Let z_t^* be a critical state where $\sigma(z_t^*, 0)$ is in the low ergodic set, while there is positive probability of entering the high ergodic set from $\sigma(z_t^*, 1)$. Then the last two terms in $\kappa_b(z_t^*, x_{b,t})$ grow without bound in δ . In an extreme case where the seller's (respectively buyer's) valuations are always above the buyer's (respectively seller's) in the low (resp. high) ergodic set, the first term in $\kappa_b(z_t^*, x_{b,t})$ is zero, while the negative terms increase with δ . (Recall, however, that (6) is sufficient but not necessary.)

On the other hand, if the states have only a small impact on the distribution over valuations, then even if we have absorbing states, the sufficient conditions can still be satisfied. For example, suppose that $\sigma(\bar{Z}, 0) = \bar{Z}$ and $\sigma(\underline{Z}, 1) = \underline{Z}$, and that players are ex ante symmetric, so that starting from a “middle” state z_0 , with the Blackwell policy, $\chi_{b,\tau}^B(\tilde{\theta}_\tau^*; z_0)$ has support $\{0, 1\}$, so that both absorbing states are possible. Even so, $\kappa_i(z_t, x_{b,t})$ grows without bound in δ so long as for each i , the effect of the states is sufficiently small.

IV. Conclusions

Dynamic mechanisms contend with more stringent incentive compatibility constraints, since players can react to information that they learn in the course of the game. Efficient, BB, BIC mechanisms can be constructed by compensating players for the change in expected opponent utilities due to the report of today's information. We show that IR can be satisfied when the long-run impact of today's information is always small relative to the worst-case player utilities, since the presence of future surplus that is unrelated to today's reports induces players to make large transfers and continue with the mechanism.

References

- d'Aspremont, Claude and Louis-Andre Gerard-Varet. 1979. "Incentives and incomplete information." *Journal of Public Economics*, 11: 25–45.
- Athey, Susan and Ilya Segal. 2006. "An Efficient Dynamic Mechanism." Mimeo, Stanford University.
- Athey, Susan, and David A. Miller. 2006. "Efficiency in repeated games with private information and moderate patience." Mimeo, University of California at San Diego.
- Myerson, Roger and Mark Satterthwaite. 1983. "Efficient Mechanisms for Bilateral Trade," *Journal of Economic Theory*, 29 (2): 265-281.

Footnotes

*Harvard University, Department of Economics, Littauer Center, 1875 Cambridge St., Cambridge, MA, 02138. email: athery@fas.harvard.edu; Stanford University, Department of Economics, Stanford, CA 94305-6072. email: ilya.segal@stanford.edu.

Supplementary Material for “Designing Efficient Mechanisms for Dynamic

Bilateral Trading Games,” by Susan Athey and Ilya Segal

American Economic Review Papers and Proceedings, May 2007

Proof of Proposition 1: Recall that we consider IR constraints after any history of reports, including misrepresentations, since a player might choose to misrepresent for several periods and then (depending on the realizations of types) exit. However, for the simple model where types follow a first-order Markov process, after a buyer sees today’s type, the fact that past reports were different than true types is irrelevant, since today’s type provides all relevant information for predicting the future. Expected discounted payoffs in period t , given reports θ_t are

$$\begin{aligned} & \theta_{i,t} \chi_i^*(\theta_t) + \sum_{\tau=t+1}^{\infty} \delta^{\tau-t} \mathbb{E}_t \left[\tilde{\theta}_{i,\tau} \chi_i^*(\tilde{\theta}_\tau) \middle| \theta_{b,t} \right] + (1\{i = b\} - 1\{i = s\}) \cdot \\ & \left(\varphi(\theta_t, \theta_{b,t-1}) - K + \sum_{\tau=t+1}^{\infty} \delta^{\tau-t} \left(\mathbb{E}_{t-1} \left[\tilde{\theta}_{s,\tau} \chi_b^*(\tilde{\theta}_\tau) \middle| \theta_{b,t-1} \right] - \mathbb{E}_t \left[\tilde{\theta}_{s,\tau} \chi_b^*(\tilde{\theta}_\tau) \middle| \theta_{b,t} \right] \right) \right). \end{aligned}$$

We simply compare these expressions to the outside option, where the seller keeps the object in the outside option. This condition combines the implied bounds on K for the worst-case scenario for the each player:

$$\begin{aligned} & \sum_{\tau=t+1}^{\infty} \delta^{\tau-t} \mathbb{E}_t \left[\tilde{\theta}_{b,\tau} \chi_b^*(\tilde{\theta}_\tau) \middle| \underline{\theta}_b \right] + \\ & \underline{\varphi} + \sum_{\tau=t+1}^{\infty} \delta^{\tau-t} \left(\mathbb{E}_t \left[\tilde{\theta}_{s,\tau} \chi_s^*(\tilde{\theta}_\tau) \middle| \bar{\theta}_b \right] - \mathbb{E}_{t-1} \left[\tilde{\theta}_{s,\tau} \chi_s^*(\tilde{\theta}_\tau) \middle| \underline{\theta}_b \right] \right) \\ & \geq \sum_{\tau=t+1}^{\infty} \delta^{\tau-t} \mathbb{E}_t \left[\tilde{\theta}_{s,\tau} \chi_b^*(\tilde{\theta}_\tau) \middle| \bar{\theta}_b \right] + \bar{\varphi} + \bar{\theta}_s \\ & + \sum_{\tau=t+1}^{\infty} \delta^{\tau-t} \left(\mathbb{E}_t \left[\tilde{\theta}_{s,\tau} \chi_s^*(\tilde{\theta}_\tau) \middle| \underline{\theta}_b \right] - \mathbb{E}_{t-1} \left[\tilde{\theta}_{s,\tau} \chi_s^*(\tilde{\theta}_\tau) \middle| \bar{\theta}_b \right] \right). \end{aligned}$$

The condition given in the proposition is sufficient for this. ■

Proof of Proposition 2: Payoffs to player i are

$$\begin{aligned} & \theta_{i,t} \chi_i^*(\theta_t; z_t) + \sum_{\tau=t+1}^{\infty} \delta^{\tau-t} \mathbb{E}_{\tilde{\theta}} \left[\tilde{\theta}_{i,\tau}^* \chi_i^*(\tilde{\theta}_\tau^*; \tilde{z}_\tau) \middle| \sigma(z_t, \chi_b^*(\theta_t; z_t)) \right] \\ & + (1\{i = s\} - 1\{i = b\}) \cdot \\ & \left(\begin{aligned} & -\varphi(\theta_t, z_{t-1}) + K - \left(\Pr \left(\chi_b^*(\theta_{b,t}, \tilde{\theta}_{s,t}; z_t) = 1 | z_t \right) - \Pr \left(\chi_b^*(\tilde{\theta}_t; z_t) = 1 | z_t \right) \right) \cdot \\ & \sum_{\tau=t}^{\infty} \delta^{\tau-t} \left(\mathbb{E}_{\tilde{\theta}} \left[\tilde{\theta}_{s,\tau}^* \chi_s^*(\tilde{\theta}_\tau^*; \tilde{z}_\tau) | \sigma(z_t, 1) \right] - \mathbb{E}_{\tilde{\theta}} \left[\tilde{\theta}_{s,\tau}^* \chi_s^*(\tilde{\theta}_\tau^*; \tilde{z}_\tau) | \sigma(z_t, 0) \right] \right) \\ & + \left(\Pr \left(\chi_b^*(\tilde{\theta}_{b,t}, \theta_{s,t}; z_t) = 1 | z_t \right) - \Pr \left(\chi_b^*(\tilde{\theta}_t; z_t) = 1 | z_t \right) \right) \cdot \\ & \sum_{\tau=t+1}^{\infty} \delta^{\tau-t} \left(\mathbb{E}_{\tilde{\theta}} \left[\tilde{\theta}_{b,\tau}^* \chi_b^*(\tilde{\theta}_\tau^*; \tilde{z}_\tau) | \sigma(z_t, 1) \right] - \mathbb{E}_{\tilde{\theta}} \left[\tilde{\theta}_{b,\tau}^* \chi_b^*(\tilde{\theta}_\tau^*; \tilde{z}_\tau) | \sigma(z_t, 0) \right] \right) \end{aligned} \right). \end{aligned}$$

The expression given in the proposition relies on finding a K such that this expression exceeds the outside options in the worst-case scenarios. ■