

An Efficient Dynamic Mechanism*

Susan Athey[†]

Ilya Segal[‡]

First Draft: November, 2004 This Draft: February, 2007

Abstract

This paper constructs an efficient, budget-balanced, Bayesian incentive-compatible mechanism for a general dynamic environment with private information. As an intermediate result, we construct an efficient, ex post incentive-compatible mechanism, which is not budget balanced. We also provide conditions under which participation constraints can be satisfied in each period, so that the mechanism can be made self-enforcing if the horizon is infinite and players are sufficiently patient.

In our dynamic environment, agents observe a sequence of private signals over a number of periods (either finite or countable). In each period, the agents report their private signals, and make public (contractible) and private decisions based on the reports. The probability distribution over future signals may depend on both past signals and past decisions. The construction of an efficient mechanism hinges on the assumption of “private values” (each agent’s payoff is determined by his own observations). Balancing the budget relies on the assumption of “independent types” (the distribution of each agent’s private signals does not depend on the other agents’ private information, except through public decisions).

*We are grateful to a co-editor, three anonymous referees, Eric Maskin, Andy Postlewaite, and participants at the fall 2004 Stanford-Berkeley theory fest, the 2007 winter meetings of the AEA, and IDEI Toulouse for helpful comments, and to the National Science Foundation for generous financial support through grants SES-0351500 (Athey) and SES-0318447 (Segal).

[†]Professor of Economics, Harvard University, Littauer Center, 1875 Cambridge St., Cambridge, MA 02138. <http://kuznets.fas.harvard.edu/~athey/>

[‡]Roy and Betty Anderson Professor in the Humanities and Sciences, Stanford University, Department of Economics, 579 Serra Mall, Stanford, CA 94305-6072. <http://www.stanford.edu/~isegal/>

1. Introduction

Multi-agent mechanism design has been used to model many important economic settings. However, most of the existing literature on multi-agent mechanism design studies a one-time decision. In reality, a sequence of decisions often needs to be made, but a static mechanism cannot be employed if the parties receive information over time that should affect the decisions, and the agents' preferences and/or technology are not time-separable. For example, parties in a long-term relationship may need to make a sequence of trading and investment decisions in a changing environment. A procurement authority may wish to conduct a sequence of auctions, where bidders have serially correlated values or capacity constraints or learning-by-doing. Ongoing utilization decisions may have to be made for a renewable resource that confers private benefits to individuals. An electricity production schedule may need to be allocated among generating plants with privately known costs that include start-up and shut-down costs. Computational capacity or bandwidth may need to be allocated in a network whose users have private values for various time blocks. When different decisions could be made either simultaneously or sequentially, sequential timing is typically more efficient since later decisions can be made contingent on the information revealed over time. Even in settings in which the resource allocation decision is a one-time decision, the agents' costs of computation (which may be viewed as information acquisition) or communication may be reduced by making the computation and communication decisions sequentially, which helps explain the prevalence of dynamic allocation mechanisms such as iterative auctions.

This paper considers the problem of sustaining incentives, balancing the budget, and satisfying participation constraints in an efficient dynamic mechanism. The additional problem arising in a dynamic mechanism is that it faces more stringent incentive constraints. In a static mechanism, it suffices to prevent deviations in which an agent pretends that he has a different type. In a dynamic mechanism, however, an agent can make his reporting strategy at a point in time contingent on the information he has gleaned about the other agents' types from past interaction, and so each possible contingency has associated incentive constraints.

We construct incentive-compatible dynamic mechanisms for a general infinite-horizon dynamic model in which each agent observes a sequence of private signals over time, and public and private decisions are made over time. The distribution of signals at any point in time may depend on the previously observed signals (e.g., allowing serial correlation) and/or previously made decisions (e.g., allowing investments that stochastically affect payoffs or affect information about payoffs). We assume that agents' payoffs depend on the signals and decisions and are quasilinear in monetary transfers. Also, we make use of two assumptions, that extend "standard" assumptions in static mechanism design to the general dynamic setting: (1) "Private Values," which means that each agent's payoff does not depend on the other agents' private information, and (2) "Independent Values," which means that the distribution of each

agent’s private signals is independent of the other agents’ private information, conditional on past public decisions.

The main result of this paper is that under these two assumptions, the additional incentive constraints do not rule out the implementation of efficient decision plans using budget-balanced transfers. To show this, we construct budget-balanced mechanisms that induce “truthfulness and obedience”: they provide agents with the incentives to report private signals truthfully and to obey the mechanism’s prescriptions for private decisions at all times and for all possible histories. Budget balance is desirable because there may not be an external source of funds, and burning money is inefficient for the agents. Thus, our result extends to the dynamic setting the static results of Arrow (1979) and d’Aspremont and Gerard-Varet (1979) (AGV), as well as Rogerson’s (1992) results incorporating private actions. Under more restrictive conditions, we show that participation constraints can be satisfied as well.

We begin by ignoring both participation constraints and budget balance. We observe that with private values, it is possible to induce truthfulness and obedience throughout the dynamic mechanism using a “Team Mechanism,” where transfers give each agent the sum of the other agents’ utilities in each period. Such transfers make each agent the residual claimant for total surplus and provide him with the incentive to be truthful and obedient as long as the mechanism prescribes an efficient decision rule. This mechanism could be viewed as an extension of the famous Vickrey-Groves-Clarke mechanism to a fully dynamic setting.¹

The problem with the Team Mechanism is that it is not budget-balanced. As we illustrate using a simple example in Section 2, a naive attempt to balance the budget using the idea of the AGV mechanism runs into the following difficulty. In a static setting, the AGV mechanism supports truthtelling in a Bayesian-Nash equilibrium: an agent has an incentive to report truthfully given his beliefs about opponents’ types, because the expected value of his transfer is equal to the “expected externality” his report imposes on the other agents through its effect on decisions. Thus, an agent’s current beliefs about opponent types play an important role in determining his transfer. However, in a dynamic setting these beliefs evolve over time as a function of opponent reports and the decisions those reports induce. If the transfers are constructed using the agents’ prior beliefs at the beginning of the game, the transfers will no longer induce truthful reporting after agents have gleaned some information about each other’s types. If, instead, the transfer to one agent is conditioned on information that is revealed by a second agent earlier in the game, the second agent’s own incentives to report truthfully will be undermined at that earlier stage.

Despite these difficulties, we show that dynamic efficiency can be implemented with balanced budget in the case of independent private values. We demonstrate a mechanism that achieves

¹For some special dynamic settings, similar non-budget-balanced efficient mechanisms are proposed by Friedman and Parkes (2003), and subsequently to the present paper, by Bapna and Weber (2005) and Bergemann and Valimaki (2006).

this, a mechanism that we call “the Balanced Mechanism.” We construct transfers that sustain an equilibrium in truthful and obedient strategies by giving each agent in each period an incentive payment equal to the *change* in the present value of the expected utilities of the other agents that is induced by his current report. We show that on the one hand, these incentive payments give incentives for truthful reporting and obedience, by letting each agent internalize the expected externality imposed on the other agents by his reports. On the other hand, these incentive payments have the property that the expected incentive payment to opposing agents is zero when he is truthful and obedient, no matter what the others report. Hence, the expected incentive payment to one agent cannot be manipulated by the other agents. This allows us to balance the budget by letting the incentive payment of a given agent be paid by the other agents without affecting those agents’ reporting incentives.

A well-known shortcoming of the AGV mechanism is that it need not satisfy the agents’ participation constraints. In fact, in some some well-known static cases, there exists no efficient balanced-budget mechanism satisfying participation constraints (e.g., Myerson and Satterwaite (1983), Mailath and Postlewaite (1990)). The same will obviously remain true for dynamic settings with a finite horizon or low patience. For this reason, we focus on an infinite-horizon model, and provide conditions under which participation constraints in each period can be satisfied when agents are sufficiently patient. Intuitively, the conditions ensure that agents’ private information in a given period is not “too” persistent and so has relatively little effect on the continuation payoffs of patient agents. Under these conditions, the payments in the Balanced Mechanism are bounded even as the continuation payoffs grow with the agents’ patience, hence the mechanism can be made self-enforcing, with agents implementing all decisions and payments without an external enforcer, and punishing any detected deviation with a breakdown in cooperation.

Most of the existing literature on dynamic mechanism design has avoided dealing with the problem of contingent deviations, by focusing on one of the following simple cases: (i) a single agent with private information (e.g., Courty and Li (2000), Battaglini (2005)), (ii) a continuum of agents with i.i.d. private information whose aggregate is predictable (e.g., Atkeson and Lucas (1993)), or (iii) information is independent across periods and preferences and technology are time-separable (Wang (1995), Athey and Bagwell (2001), Athey, Bagwell, and Sanchirico (2004), Levin (2003), Rayo (2003), Miller (2004), Athey and Miller (2004), but see Athey and Bagwell (2004) for an exception).² In each of these cases, an individual agent learns nothing in the course of the mechanism about the others’ types that is relevant for the future, hence there is no need to consider contingent deviations. In more general settings, however, even if the mechanism hides the agents’ reports from each other, an agent would typically be able to infer

²Part of the literature on dynamic contracting considers the case of imperfect commitment (e.g. Bester and Strausz (2001), Battaglini (2003), and Krishna and Morgan (2004)). We sidestep this issue by allowing the agents to commit to the mechanism in advance.

something about the other agents' types from the prescribed decisions, and try to exploit this information in contingent deviations. Thus, a dynamic mechanism has to satisfy more incentive constraints than a corresponding static one, in which no information leaks out. Satisfying these additional incentive constraints constitutes the main contribution of this paper.

2. Some motivating examples

We consider a series of examples with alternative assumptions.

Example 1. (*Two-Period Trading Game*) A seller (agent 1) and a buyer (agent 2) engage in a two-period relationship. In each period $t = 1; 2$, they can trade a contractible quantity $x_t \in [0; 1]$. Before the first period, the seller observes a signal $\tilde{c}_1 \in [0; 1]$, which determines his cost function $\frac{1}{2\tilde{c}_1} (x_t)^2$ in each period $t = 1; 2$. The buyer's value per unit of the good in period 1 is equal to 1, and in period 2 it is given by a signal $\tilde{v}_2 \in [0; 2]$ that she observes between the periods.

An efficient (surplus-maximizing) mechanism must have trading decisions x_1 and x_2 determined by the allocation rules $x_1(\tilde{c}_1) = \tilde{c}_1$, $x_2(\tilde{c}_1; \tilde{v}_2) = \tilde{c}_1 \tilde{v}_2$ respectively. Note in particular that the first-period trade will reveal the seller's type \tilde{c}_1 to the buyer.

The problem of designing an efficient mechanism comes down to designing transfers to each agent as a function of their reports. In this simple setting, each agent makes only one report. Let us first consider the AGV mechanism for this problem, for the case where the buyer does not learn the seller's type before making his announcement. To give the buyer an incentive for truthful revelation, she is charged an "incentive payment" equal to the expected externality he imposes on the seller, i.e., the seller's expected cost:

$$t_2(\tilde{v}_2) = -\mathbb{E}_{\tilde{c}_1} \left[\frac{1}{2\tilde{c}_1} \left(\tilde{c}_1(\tilde{c}_1) \right)^2 + \frac{1}{2\tilde{c}_1} \left(\tilde{c}_1(\tilde{c}_1; \tilde{v}_2) \right)^2 \right] = -\frac{1}{2} \mathbb{E}_{\tilde{c}_1} [\tilde{c}_1] \cdot (1 + (\tilde{v}_2)^2):$$

Similarly, the seller's incentives are provided by paying him an "incentive payment" equal to the expectation of the buyer's utility:

$$t_1(\tilde{c}_1) = \mathbb{E}_{\tilde{v}_2} \left[\tilde{c}_1(\tilde{c}_1) + \tilde{v}_2 \cdot \tilde{c}_1(\tilde{c}_1; \tilde{v}_2) \right] = \tilde{c}_1 \left(1 + \mathbb{E}_{\tilde{v}_2} [(\tilde{v}_2)^2] \right):$$

Now, since each party's incentive payment does not depend on the other's report, we can balance the budget simply by charging each party's incentive payment to the other party, i.e., letting the total transfer to each agent i be $t_i(\tilde{c}_i; \tilde{v}_i) = t_i(\tilde{c}_i) - t_{-i}(\tilde{v}_i)$.

Now we turn to the case of interest, where the buyer makes his announcement after the seller's type \tilde{c}_1 is revealed. If we use the AGV transfers described above, the buyer anticipates that the second-period trade will be determined by $x_2(\tilde{c}_1; \cdot)$. However, the buyer must pay (through $t_2(\tilde{v}_2)$) the expectation (over \tilde{c}_1) of the cost of the seller. Then, if the seller's cost \tilde{c}_1 is known to be higher than average, the buyer "over-reports" his value to induce inefficiently

high trade, since the buyer will not have to pay the actual cost of this trade. Similarly, the buyer does not internalize the benefit of an unexpectedly low cost, and in that case he “under-reports” his value to induce less-than-efficient trade.

To fix this problem, we could instead give the buyer an incentive transfer based on the actual externality he imposes given the seller’s report: $\tilde{t}_2(c_1; c_2) = -\frac{1}{2} c_1 \cdot (1 + (c_2)^2)$. This will give the buyer the incentive to report c_2 truthfully, no matter what c_1 the seller reports. However, this transfer depends on the seller’s report c_1 . Thus, if we attempted to balance the budget by having it be paid by the seller, making his total transfer $\tilde{t}_1(c_1; c_2) = c_1(c_1) - \tilde{t}_2(c_1; c_2)$, the seller would want to reduce $\tilde{t}_2(c_1; c_2)$ by overreporting his cost c_1 in period 1, thus exaggerating the externality imposed on him by the buyer.

The problem of contingent deviations illustrated here arises not only when types are persistent as in the above example, but in any dynamic setting in which the agents’ preferences and/or technology are not separable across periods.

In this paper we propose a different way to construct transfers that resolves the problem. Similarly to the AGV mechanism, our construction proceeds in two steps: (1) Construct incentive transfers $t_1(c)$, $t_2(c)$ to make each agent report truthfully if he expects the other to do so, and (2) charge each agent’s incentive transfer to the other agent, making the total transfer to agent i equal $t_i(c) = t_i(c) - t_{-i}(c)$. However, in contrast to AGV transfers, the incentive transfer $t_i(c)$ to agent i will now depend not just on agent i ’s announcements c_i but on those of the other agents. How then do we then ensure that step 2 does not destroy incentives? For this purpose, we ensure that even though agent $-i$ can affect the other’s incentive payment $t_{-i}(c; c_{-i})$, he cannot manipulate the *expectation* of that payment given that agent $-i$ reports truthfully. We achieve this by letting $t_{-i}(c; c_{-i})$ be the *change* in the expectation of agent $-i$ ’s final utility conditional on all the previous announcements that is brought about by the report of agent i . (In the general model in which an agent reports in many periods, these incentive transfers would be calculated in each period for the latest report). No matter what announcement strategy agent $-i$ adopts, if agent i announces truthfully, the expectation of agent $-i$ ’s utility follows a martingale with respect to agent i ’s announcements — that is, the changes have zero expectation. Hence agent $-i$ can be charged $t_{-i}(c; c_{-i})$ without affecting his incentives.

In Example 1, our construction entails giving the buyer an incentive transfer of

$$t_2(c_1; c_2) = -\frac{1}{2} c_1 \cdot \left((c_2)^2 - \mathbb{E} \left[(\tilde{c}_2)^2 \right] \right);$$

which on the one hand gives him correct incentives by letting him internalize the seller’s expected cost (since $\mathbb{E}_{-1} \left[t_2(\tilde{c}_1; c_2) \right] = -\frac{1}{2} \mathbb{E}_{-1} \left[\tilde{c}_1 \right] \cdot (c_2)^2 + \text{const}$), and on the other hand ensures that the expectation of this transfer cannot be manipulated by seller: $\mathbb{E}_{-2} \left[t_2(c_1; \tilde{c}_2) \right] = 0$ for any c_1 . Therefore, we can now charge this incentive transfer to the seller — i.e., let $t_1(c_1; c_2) = c_1(c_1) - t_2(c_1; c_2)$ — without undermining the seller’s incentives for truthful reporting. Also, letting then $t_2(c_1; c_2) = -t_1(c_1; c_2) = t_2(c_1; c_2) - c_1(c_1)$ balances the budget and provides

incentives for the buyer to report truthfully.

In the rest of this paper, we show that the idea of charging an agent for the change in opponent utilities induced by his report can be applied to yield an efficient mechanism in a general dynamic setting. The idea that agent i 's incentive payments cannot be manipulated by others remains the same in the more general model, but it becomes more subtle to show that the anticipation by agent i of his own future incentive payments provides the correct incentives.

The general setting we study includes, for example, the following variations of Example 1.

Example 2. (*Serially Correlated Valuations*) Suppose that the buyer privately observes his first-period valuation ($v_{i,1}$) before the first period, and his second-period valuation ($v_{i,2}$) before the second period; and that these two valuations are correlated.³

Example 3. (*Decisions affect Valuations*) In addition to “exogenous” correlation between the buyer’s first and second period valuations, the buyer’s consumption level in the first period (x_1) affects the probability distribution over his second-period valuation. This might arise due to the buyer’s “habit formation.”

Example 4. (*Decisions Affect Information*) The buyer never observes his valuation $v_{i,2}$, but after the first period, she observes a signal S that is correlated with $v_{i,2}$ and whose probability distribution (e.g., informativeness) depends on her first-period consumption x_1 . This could describe the buyer’s learning about the product by consuming it.

Example 5. (*Hidden Actions affect Valuations*) The buyer makes a private investment a in the first period that affects the probability distribution over her second-period valuation. This investment is not observable by the other agents and is not verifiable.

Example 6. (*Hidden Actions affect Information*) The buyer never observes his valuation $v_{i,2}$, but after the first period, she observes a signal S that is correlated with $v_{i,2}$ and whose probability distribution (e.g., informativeness) depends on a private investment a she makes in the first period. For example, the investment could model the buyer’s gathering of information about or computation of his preferences.

3. The Setup

3.1. Technology and Payoffs

We consider a model with I agents and a countable number of periods, indexed by $t = 1; 2; \dots$. (A finite-horizon model with T periods is a special case.) In each period t , each agent $i \in \{1; \dots; I\}$

³A model of dynamic procurement auctions with serially correlated values has been considered by Athey and Bagwell (2004).

observes a private signal $\omega_{i,t} \in \Theta_{i,t}$, where $\Theta_{i,t}$ is a measurable space.⁴ We use $\Theta_i^t = \prod_{s=1}^t \Theta_{i,s}$ to denote the space of possible signal histories of agent i at time t , and let $\Theta_t = \prod_{i=0}^I \Theta_{i,t}$ and $\Theta^t = \prod_{i=0}^I \Theta_i^t$. We interpret $\Theta_i = \Theta_i^\infty$ as the agent's type space, even though the agent only observes his type over time, and $\Theta = \prod_{i=0}^I \Theta_i$ as the state space. All these sets will be treated as measurable spaces with the product sigma-algebra.

In each period $t = 1, \dots$, after this period's signals are observed, each agent $i \in \{1, \dots, I\}$ makes a private decision $x_{i,t} \in X_{i,t}$, and there is a public (contractible) decision $x_{0,t} \in X_{0,t}$, where $X_{0,t}$ and $X_{i,t}$ for each i are measurable spaces. We let $X_t = \prod_{i=0}^I X_{i,t}$, $X^t = \prod_{s=1}^t X_s$, and $X = X^\infty$.

The distribution of signals at any moment may depend both on the history of observed information $\omega^{t-1} \in \Theta^{t-1}$ and the history of decisions $x^{t-1} \in X^{t-1}$. Thus, the evolution of uncertainty at time t is described by history-contingent probability measures over the period- t signal space Θ_t ; denoted $\mu_t(\cdot | x^{t-1}; \omega^{t-1})$. For any measurable set $\mathcal{F} \subset \Theta_t$; $\mu_t(\mathcal{F} | \cdot; \cdot)$ is a measurable function of the history on $X^{t-1} \times \Theta^{t-1}$.^{5,6} (Note that μ_1 is a constant since there is no history in period 1.)

A *decision plan* is a measurable function $\mu : \Theta \rightarrow X$, where each $\mu_t(\cdot)$ represents the decision made at time t . We say that decision plan μ is *observationally measurable* if it is measurable and for each t , $\mu_t(\cdot)$ depends only on ω^t – the information observed by time t . We will sometimes write $\mu_t(\omega^t)$ rather than $\mu_t(\cdot)$ to emphasize observational measurability. The sequence of decisions up to time t given type vector θ is denoted $\mu^t(\theta) = (\mu_1(\theta); \dots; \mu_t(\theta))$.

Any decision plan μ uniquely determines a probability measure over Θ ; which we denote μ . The existence and uniqueness of this measure follow from the Tulcea product theorem.⁷

⁴We could also allow for publicly observed (contractible) signals. Formally, they could be incorporated, e.g., by introducing an additional “agent 0” who observes signals $\theta_{0,t}$ in each period t and always reports them truthfully.

⁵We can view these history-contingent distributions as describing Nature's behavioral strategy in the dynamic game in which Nature chooses types and the mechanism chooses decisions. Alternatively, Nature's mixed strategy could be described in a different way: e.g., in Example 4, we could be given a distribution of the buyer's value θ_2 and a distribution of the second-period signal s conditional on θ_2 for any given first-period decision x_1 . From this description, we could deduce an equivalent behavioral strategy, which gives a distribution of the current signal conditional on the history (e.g., of $s|x_1$ and of $\theta|x_1, s$). However, we prefer to start directly with behavioral strategies: The problem with alternative approaches is that conditional distributions are not well-defined on zero-probability events (e.g., when signal spaces are continuous, almost all points would have zero probabilities), while it will be important for our analysis to have well-defined forward expectations at all times and events.

⁶One may want to allow the set X_t of possible decisions at time t and the set Θ_t of possible observations at time t to depend on the history (θ^{t-1}, x^{t-1}) . (The latter dependence, for example, may occur in situations in which decisions determine which signal is observed, as in Example 4 above). However, these situations could be incorporated in our model simply by letting X_t describe the set of all potential decisions that could be made at time t at all histories, and similarly letting Θ_t be the set of all signals that could be potentially observed at time t at all histories. The actual history then determines which observations from Θ_t and which decisions from X_t are relevant for the agents' payoffs.

⁷Formally, the distributions $\nu_t |_{\mathcal{X}_{t-1}}(\theta^{t-1}), \theta^{t-1}$ are probability kernels that define a consistent family

Our results on the existence of a budget-balanced mechanism rely on the following assumption

Definition 1. We have *independent types* if for each t , \mathbf{x}^{t-1} ; and θ^{t-1} , the probability measure $\mu_t | \mathbf{x}^{t-1}; \theta^{t-1}$ over Θ_t can be written in the form

$$\mu_t | \mathbf{x}^{t-1}; \theta^{t-1} = \prod_{i=0}^I \mu_{i;t} | \mathbf{x}_0^{t-1}; \mathbf{x}_i^{t-1}; \theta_i^{t-1};$$

for some probability measures $\mu_{i;t}$ over $\Theta_{i;t}$ that depend only on the public decision history \mathbf{x}_0^{t-1} and on agent i 's private history $(\theta_i^{t-1}; \mathbf{x}_i^{t-1})$.

This definition means that an agent's private information does not have any stochastic effects on the other agents' private signals, conditional on the public decisions. Note that this assumption still allows one agent's private signals or decisions to affect another agent's future private signals through the implemented public decisions. Thus, the probability measure μ_t over Θ induced by a given decision rule need not be independent across agents.

Now we describe the agents' payoffs. The payoff of each agent i is given as a function of the stream \mathbf{x} of decisions, state θ , and stream of monetary transfers $\mathbf{y}_i = (y_{i;t})_{t=1}^{\infty}$; where $y_{i;t} \in \mathbb{R}$:

$$\sum_{t=1}^{\infty} \beta^t (u_{i;t}(\mathbf{x}^t; \theta^t) + y_{i;t});$$

where $\beta \in (0; 1)$ is a discount factor, and the functions $u_{i;t} : \mathbf{X}^t \times \Theta^t \rightarrow \mathbb{R}$ are assumed to be uniformly bounded and measurable.

Since each function $u_{i;t}$ in each period t can depend on the entire history of signals and decisions, writing the payoff as a present discounted value is not too restrictive. For example, in a setting with a finite horizon T , any utility functional $U_i(\mathbf{x}; \theta)$ can be represented as

$$U_i(\mathbf{x}; \theta) = \sum_{t=1}^{\infty} \beta^t u_{i;t}(\mathbf{x}^t; \theta^t); \tag{3.1}$$

simply by letting $u_{i;t} \equiv 0$ for $t \neq T$ and $u_{i;t}(\mathbf{x}^T; \theta^T) \equiv \beta^{-T} U_i(\mathbf{x}; \theta)$. As for the general infinite-horizon case, in the Appendix we show that a functional $U_i : \mathbf{X} \times \Theta \rightarrow \mathbb{R}$ can be written in the form (3.1) for some uniformly bounded sequence $(u_{i;t})_{t=1}^{\infty}$ if and only if U_i has the following property, which ensures that decisions and observations in the distant future have a vanishing effect on the payoffs, and which will be important for our analysis:⁸

$\{\mu^T\}_{T=1}^{\infty}$ of probability measures on Θ^T for all finite horizons T . The Tulcea product theorem (see, e.g., Pollard (2002), Chapter 4, Theorem 49) establishes that there exists a unique stochastic process on Θ whose marginals for any finite horizon T coincide with μ^T .

⁸This is a mild strengthening of the usual "continuity at infinity" assumption in infinite-horizon games (e.g., Fudenberg and Tirole (1991a), p. 110).

Definition 2. A functional $F : \prod_{t=1}^{\infty} Z_t \rightarrow \mathbb{R}$ is *Lipshitz Continuous at Infinity* with discount factor $\delta \in (0; 1)$ (denoted *LCI*) if it is Lipshitz continuous in the metric $d(z; z') = \inf\{\delta^t : z_t \neq z'_t\}$, that is, there exists $C > 0$ such that $|F(z) - F(z')| \leq C \delta^t$ for all $z, z' \in Z$:

Our main results will depend on the following property of payoffs:

Definition 3. We have *private values* if each agent i 's utility $U_i(x; \cdot)$ depends only on the public decisions x_0 and the agent's private history $(x_i; \cdot)$

This definition simply means that each agent observes his payoff, i.e., can calculate it as a function of his own information. It is less restrictive than appears at first because it allows an agent's *expected* utility conditional on time- t history $(\cdot^t; x^t)$ to depend on the other agents' information, as long his *final* utility is fully determined by his observations (as, e.g., in Mezzetti (2004)).⁹ We could also allow some cases in which agent i may not observe his final utility (as, e.g., in Example 4), by letting $U_i(x; \cdot)$ represent the agent's expectation of his utility conditional on the state ω , provided that this expectation does not depend on the other agents' private signals or decisions.

3.2. Mechanisms and Strategies

We consider mechanisms in which in each period, agents report their private information, and based on the reports, decisions are implemented, and transfers are made.¹⁰ Formally, a *mechanism* is described with two observationally measurable functions: a decision plan $\sigma : \Theta \rightarrow X$ and a uniformly bounded transfer plan $\tau : \Theta \rightarrow (\mathbb{R}^I)^\infty$, which prescribe the current decisions and transfers to the agents as a function of the reporting history. The public decisions $\sigma_0(\cdot)$ prescribed by the plan are implemented directly, while the prescribed private decisions $\sigma_i(\cdot)$ for agents $i \geq 1$ are nonbinding recommendations that the agents are free to disobey. The transfer plan is implemented directly. Its observational measurability means that the transfer $\tau_{i;t}(\cdot)$ made to agent i in period t depends only on the current history \cdot^t . The total discounted payments in the mechanism with a discount factor $\delta \in (0; 1)$ can be calculated as

$$\Psi_i(\cdot) = \sum_{t=0}^{\infty} \delta^t \tau_{i;t}(\cdot^t); \tag{3.2}$$

⁹The property of Private Values is violated in some well-known settings, such as auctions for mineral rights, For such settings, however, it is known that efficiency is in general not implementable even in a static problem, as in Akerlof's (1970) classic example of trading with adverse selection. See also Jehiel and Moldovanu (2001).

¹⁰Thus, we focus on direct revelation mechanisms in which agents report truthfully in equilibrium. This focus is innocuous since our goal is to propose particular mechanisms rather than characterize everything that is implementable. However, if we were interested in such characterization, we could restrict attention to direct revelation mechanisms with truthful reporting by making use of a revelation principle proposed by Myerson (1986) for general dynamic games.

where uniform convergence of the series is ensured by the uniform boundedness of the transfer plan. Note that the agents' reporting incentives in a mechanism depend on the transfer plan only through the total discounted payments $\Psi(\cdot)$ and not on how they are spread over time. (However, if agents have opportunities to exit the mechanism during the course of the game, the satisfaction of their participation constraints will depend on the spread of payments - see Section 6 below.)

We say that the mechanism is *budget balanced* if $\sum_i v_{i,t}(\cdot) \equiv 0$ for all t .

In principle, the mechanism also determines what information agents observe about the history of reports. In general, the less information is revealed to an agent, the smaller is the set of contingent deviations available to him, and the easier it is to satisfy his incentive constraints (this idea underlies Myerson's (1986) "communication equilibrium"). In the extreme, if nothing is revealed to the agent, the problem is essentially static. In practice, the agent will observe at least some components of the decisions—e.g., the history of his own consumption—and so he will have access to at least some contingent deviations. We will propose mechanisms in which truthful-obedient play is incentive-compatible even if the agents observe the complete reporting history, and therefore will remain incentive-compatible no matter what the agents observe about the history (so long as agents observe enough information to implement their private decision plans).

The mechanism induces a multiperiod game, in which each period t consists of three stages:

Stage $t:0$: Each agent i privately observes his private signal $v_{i,t}$.

Stage $t:1$: Each agent i makes a public report $\hat{v}_{i,t}$.

Stage $t:2$: Each agent i makes a private decision $x_{i,t}$. The mechanism implements the public decision $x_{0,t} = x_{0,t}(\hat{v}^t)$ and the payments $y_{i,t} = y_{i,t}(\hat{v}^t)$ to each agent i .

Now we turn to describing the agents' strategies in the mechanism. By stage $t:1$ each agent i observes the reporting history $\hat{v}^{t-1} \in \Theta^{t-1}$ and the history of own private signals $v_i^t \in \Theta_i^t$ and private decisions $x_i^{t-1} \in X_i^{t-1}$. Thus, the agent's reporting strategy can be represented by a function $\hat{v}_i : \Theta \times \Theta_i \times X_i \rightarrow \Theta_i$, where the report of agent i in period t , $\hat{v}_{i,t}(\hat{v}^{t-1}; v_i^t; x_i) = \left(\hat{v}_i(\hat{v}^{t-1}; v_i^t; x_i) \right)_t$, is "observationally measurable" in that it depends only on the history of reports \hat{v}^{t-1} and the agent's private history $v_i^t; x_i^{t-1}$. (For our results about the "Team Mechanism" below, we will allow deviations to more general strategies \hat{v}_i in which $\hat{v}_{i,t}$ can depend on the reporting history \hat{v}_i^t that includes the concurrent reports of the other agents, which we call "measurable with respect to *current* observations.")

In stage $t:2$, each agent i observes in addition the reports from stage $t:1$. Thus, his private decision strategy can be represented by $x_i : \Theta \times \Theta_i \times X_i \rightarrow X_i$, where in each period t , $x_{i,t}(\hat{v}^{t-1}; v_i^t; x_i) = \left(x_i(\hat{v}^{t-1}; v_i^t; x_i) \right)_t$, is "observationally measurable" in that it depends only on

the reporting history $\hat{\tau}^t$ and the agent's private history $\tau_i^t; x_i^{t-1}$. The complete strategy of agent i is given by $(\tau_i; \sigma_i)$.

A given observationally measurable strategy $(\tau_i; \sigma_i)$ of agent i induces a *strategic plan* $(\bar{\tau}_i; \bar{\sigma}_i)$, where $\bar{\tau}_i(\cdot)$ denotes the agent's reports and $\bar{\sigma}_i(\cdot)$ his decisions from following the strategy given that his type is τ_i and the opponents report τ_{-i} . This plan is constructed recursively, by letting for each $t = 1; \dots$, $\bar{\tau}_{i;t}(\cdot) = \tau_{i;t}((\bar{\tau}_{-i}^{t-1}(\tau_{-i}^{t-1}); \tau_{-i}^{t-1}); \tau_i^{t-1}; \bar{\tau}_{-i}^{t-1}(\tau_{-i}^{t-1}))$, and $\bar{\sigma}_{i;t}(\cdot) = \sigma_{i;t}((\bar{\tau}_{-i}^t(\tau_{-i}^t); \tau_{-i}^t); \tau_i^{t-1}; \bar{\tau}_{-i}^{t-1}(\tau_{-i}^{t-1}))$. Observational measurability means that $\bar{\tau}_{i;t}(\cdot)$ depends only on $(\tau_i^{t-1}; \tau_{-i}^{t-1})$ and $\bar{\sigma}_{i;t}(\cdot)$ depends only on τ_{-i}^t . Observe that agent i 's expected payoff in the mechanism from using strategy $(\tau_i; \sigma_i)$ depends only on the induced strategic plan $(\bar{\tau}_i; \bar{\sigma}_i)$, for any strategies chosen by the other agents. Thus, all the strategies giving rise to the same strategic plan are strategically equivalent in the normal-form representation of the mechanism. (However, they may have different implications for extensive-form refinements, as discussed in Section 3.3 below).

Definition 4. For a given decision plan σ , agent i 's strategy $(\tau_i; \sigma_i)$ is a *truthful-obedient strategy* if the corresponding strategic plan has $\bar{\tau}_i(\cdot) = \tau_i$ and $\bar{\sigma}_i(\cdot) = \sigma_i(\cdot)$ for all $\tau_{-i} \in \Theta$:

This definition means that if an agent has been truthful and obedient so far he will continue to be truthful and obedient. It does not specify anything about reporting after a lie (i.e., $\hat{\tau}_i^{t-1} \neq \tau_i^{t-1}$) or a disobedience (i.e., $x_i^{t-1} \neq \tau_i^{t-1}(\hat{\tau}_i^{t-1})$), and so there are many truthful-obedient strategies.

Our proofs will make use of the following simple observation: given a decision plan σ and transfer function τ ; when opposing agents use truthful-obedient strategies, any strategy $(\tau_i; \sigma_i)$ for agent i induces the same decisions and transfers in every state as those that arise when agent i uses a truthful-obedient strategy but the decision plan and transfers are defined as follows:

$$\hat{\tau}(\cdot) \equiv (\tau_{-i}(\bar{\tau}_i(\cdot)); \tau_{-i}); \bar{\sigma}_i(\cdot) \quad (3.3)$$

$$\hat{\sigma}(\cdot) \equiv (\bar{\tau}_i(\cdot); \tau_{-i}); \text{ and } \hat{\Psi}(\cdot) \equiv \sum_{t=0}^{\infty} \tau_{i;t}^{\hat{\tau}}(\cdot) \quad (3.4)$$

Note that if decision plan σ and payment plan τ are observationally measurable, then the decision plan $\hat{\tau}$ and payment plan $\hat{\sigma}$ given by (3.3-3.4) are also observationally measurable.

When agent i uses reporting strategy τ_i with the associated $\bar{\tau}_i$, $\hat{\tau}$; and $\hat{\Psi}$; while the other agents are truthful, agent i 's payoff in the mechanism is

$$\mathbb{E}_{\tau_{-i}}^{[\hat{\tau}]} \left[U_i \left(\hat{\tau}(\tau_{-i}); \hat{\sigma} \right) + \hat{\Psi}_i(\tau_{-i}) \right] \quad (3.5)$$

We use $\mathbb{E}_{\tau_{-i}}$ to denote the expectation of random variable $\hat{\tau}$ whose probability distribution is τ_{-i} . Note that different reporting strategies by the agent induce different distributions over types, and this is captured by the fact that we take expectations using probability distribution $\tau_{-i}^{[\hat{\tau}]}$ where $\hat{\tau}$ varies with the reporting strategy.

3.3. Solution Concept

In this paper, we focus on the solution concept of Bayesian Nash Equilibrium (BNE), so that each agent’s ex ante choice of a history-contingent reporting strategy must be a best response to the history-contingent reporting strategies of opponents. Thus, the existence of a truthful-obedient Bayesian-Nash equilibrium means that expression (3.5) is maximized when agent i uses a truthful reporting strategy, and so $\hat{\sigma}_i = \sigma_i$ and $\hat{\Psi}_i = \Psi_i$.

In the dynamic setting, one may be interested in dynamic equilibrium refinements. Here we focus on Perfect Bayesian Equilibrium (PBE) as defined in Fudenberg and Tirole (1991) and modified to fit our model (which has types learned over time rather than at the beginning of the game, as in their model). We argue that under some conditions, if a truthful-obedient BNE exists, we can also construct a truthful-obedient PBE.

To simplify the discussion, we focus on the special case where the following hold: independent types; decisions and signals are drawn from finite sets in each period; and all signals have positive probabilities for any history. Then, the requirements of PBE are: (i) Player i ’s belief about the types of other agents is the product of agent i ’s beliefs about the types of individual agents. These beliefs do not depend upon agent i ’s type. (ii) Players i and j always have the same beliefs about the type of any third agent k , and agent i ’s belief about agent j ’s type is consistent with the understanding that the behavior of any other agent cannot signal this information. (iii) Beliefs are determined by Bayes’ rule whenever possible. This includes circumstances in which observed play in a previous period is incompatible with the beliefs and equilibrium strategies for that period. That is, after a zero-probability event, agents form beliefs, but then they continue to use Bayes’ rule when possible from there. (iv) Given any history and beliefs, the agents’ strategies from that point forward are a Bayesian-Nash equilibrium of the continuation game.

Note that when all signals have positive probabilities, opponents cannot report anything “unexpected,” and so all zero-probability information sets of an agent arise following his own deviations. Since the agent cannot signal something to himself that he does not know, his beliefs about the others are unchanged following such a deviation. Then, since the definition of a truthful-obedient strategy does not restrict the agent’s behavior following own deviations, we can modify his strategy to be a best-response to such beliefs at this information set, while still being truthful-obedient.¹¹ Doing this for all information sets and for all agents, we obtain a truthful-obedient PBE.

¹¹To see that a best response is guaranteed to exist, note that since U_i is LCI $_{\delta}$, the agent’s objective is continuous in his strategy $\sigma_i = (\alpha_i, \beta_i)$ in the metric $\rho_{\delta}(\sigma_i, \sigma'_i) = \delta^{\inf\{t:\sigma_{i,t} \neq \sigma'_{i,t}\}}$, and his strategy space is compact in the metric when the action set is finite in each period. To establish compactness, take any sequence of strategies, find a set of first-period strategies that occurs infinitely often, then find a strategy with the same first-period strategy and that has a second-period strategy that is used infinitely often, etc., and this will be a converging subsequence in the metric.

4. The Team Mechanism

We want to implement an efficient decision plan δ^* , that is, one that maximizes the total expected surplus

$$\mathbb{E}_{\tilde{\omega}} \left[\sum_i U_i \left(\delta^*(\tilde{\omega}); \tilde{\omega} \right) \right]; \quad (4.1)$$

among all observationally measurable decision plans. The problem of calculating δ^* is a stochastic dynamic programming problem, but for the purpose of this paper we simply assume that a solution exists.¹²

The Team Mechanism consists of the efficient decision plan δ^* together with the transfer functions

$$M_{i;t}^M(\omega) = \sum_{j \neq i} U_{j;t}(\delta^*(\omega); \omega);$$

which give each agent the sum of other agents' utilities. (Note that the per-period transfers are uniformly bounded since per-period utilities are uniformly bounded.) The resulting total discounted transfers are

$$\Psi_i^M(\omega) = \sum_{j \neq i} U_j(\delta^*(\omega); \omega);$$

We say that a strategy profile is a *within-period ex post equilibrium* if each agent's reporting strategy is a best response even among strategies that are measurable with respect to *current* observations (recall that in general, we restrict reporting strategies to be measurable only with respect to *past* opponents' reports). This is a refinement of Bayesian-Nash equilibrium, which does not consider deviations based on the opponents' current types. Note that the usual concept of ex post equilibrium considered in the static mechanism design literature allows deviations based on all the information eventually obtained. This would be too strong for the dynamic setting, where, even with one agent, this agent might want to report differently and make different decisions if he could foresee the future signals.¹³ A notable exception is given by settings in which there exists an efficient decision rule that is "distribution-free," i.e., maximizes the total surplus for all states of the world, rather than just its expectation. An example is a case where $U_{i;t}(x^t; \omega_i^t)$ and $\omega_i^t | x^{t-1}; \omega_i^{t-1}$ do not depend on the past decisions x^{t-1} , and the efficient

¹²For example, it is guaranteed to exist when all the signal spaces and decision spaces are finite, by the same argument as in footnote 11 above.

¹³The refinement of ex post Nash equilibrium has been considered mainly in static mechanism design, where it requires that an agent's announcement must be a best response to the realized values of opponent announcements, when opponents follow the proposed equilibrium strategies. In a static setting, private-values context ex post equilibrium implementation is equivalent to implementation in dominant strategies, but this equivalence breaks down in a dynamic model where strategies might specify that a player responds to an opponent's announcement in pathological ways. See Miller (2004), who motivates and applies "within-period" ex post equilibrium to dynamic games.

policy $\bar{t}^*(t)$ maximizes period-by-period surplus $\sum_i u_{i,t}(\bar{t}^*(t); t_i)$. In such settings, the Team Mechanism we propose sustains truthful reporting as a “true” ex post equilibrium.

Proposition 1. *Suppose that we have private values. Then for any mechanism with an efficient decision plan \bar{t}^* and total discounted transfers Ψ that equal the total discounted transfers $\Psi_i^M(\cdot)$ in the Team Mechanism, any truthful-obedient strategy profile is a Bayesian-Nash equilibrium and a within-period ex post equilibrium.*

Proof. Consider a mechanism satisfying the conditions of the proposition. Suppose that agent i deviates to induce a strategic plan $(\bar{t}_{-i}; \bar{t}_i)$ that is measurable with respect to current observations, while the other agents use truthful-obedient strategies. The decision and transfer plans following the deviation are given by \hat{t} and $\hat{\Psi}$ from (3.3) and (3.4), with \hat{t} being observationally measurable. Note that

$$\begin{aligned} \hat{\Psi}_i(\cdot) &= \sum_{j \neq i} U_j(\bar{t}_{-i}(\cdot); \bar{t}_i(\cdot); \bar{t}_i(\cdot); \bar{t}_i(\cdot)) \\ &= \sum_{j \neq i} U_j(\hat{t}_{-i}(\cdot); \bar{t}_i(\cdot); \bar{t}_i(\cdot); \bar{t}_i(\cdot)) \\ &= \sum_{j \neq i} U_j(\hat{t}(\cdot); \cdot); \end{aligned}$$

where the first equality is by construction of the transfers, the second by definition of \hat{t} , and the third by the assumption of private values, which means that U_j does not depend directly on \bar{t}_i or x_i : Thus, agent i 's expected payoff from the deviation (3.5) is given by

$$\mathbb{E}_i^{[\hat{t}]} \left[U_i(\hat{t}(\cdot); \cdot) + \sum_{j \neq i} U_j(\hat{t}(\cdot); \cdot) \right];$$

Since $\hat{t}(\cdot) = \bar{t}^*(\cdot)$ when $(\bar{t}_{-i}; \bar{t}_i)$ is a truthful-obedient strategy, and since by construction \bar{t}^* maximizes expected total surplus among all observationally measurable decision plans, the deviation cannot raise agent i 's expected payoff. ■

The intuition for the proof is that the transfers make each agent into a claimant for the total expected surplus, which is maximized when all agents adhere to truthful reporting strategies. Hence, no agent has an incentive to deviate. This is a straightforward extension of the Vickrey-Groves-Clarke mechanism to our dynamic model.

A simple way to implement the Team Mechanism is by asking the agents to report their realized payoffs (i.e., viewing these payoffs as signals) and giving each agent a transfer equal to the sum of the other agents' reports. One advantage of this implementation is that constructing the transfers does not require knowing any more details about the environment than what is needed to calculate the decision plan. Furthermore, if the designer wants to calculate the decision plan but does not know the agents' utility functions, he can ask each agent to report

his own utility function (which could be modeled as a “rich” signal available in period 1), and by Proposition 1 the agents can be relied upon to report truthfully. Recall, however, that in general the efficient decision rule depends on the agents’ beliefs as well as the their utility functions. If agents’ beliefs are subjective and different agents may hold different beliefs, there is no “detail-free” efficient decision rule, let alone a detail-free mechanism.

5. Balancing

A major problem with the Team Mechanism is that its transfers are not budget balanced. However, we now show that they can be balanced, provided that the types are independent (as defined in Subsection 3.1).

Our construction of the balanced mechanism relies on the following notation: For any period t and history \tilde{t}^{-1} , a given decision plan π induces a probability measure over Θ , which assigns probability 1 to $\tilde{t}^{-1} = \tilde{t}^{-1}$, and which we denote by $\pi_t[\cdot] \mid \tilde{t}^{-1}$: This measure exists and is unique by the Tulcea Product Theorem (see footnote 7). Also, for each agent i , and each \tilde{t}_i , let

$$\pi_t[\cdot] \mid \tilde{t}_i, \tilde{t}_i^{-1} = \mathbb{E}_{\pi_t}^{t \mid \tilde{t}_i^{-1}(\tilde{t}_i^{-1}); \tilde{t}_i^{-1}}[\pi_t[\cdot] \mid \tilde{t}_i; (\tilde{t}_i^{-1}; \Gamma_{-i;t})];$$

This is the measure probability measure over Θ when agent i ’s private signals are given through period t and all the other signals are given through period $t - 1$.¹⁴

Using this notation, for any mechanism that implements decision plan π using transfers τ , we construct a new mechanism with the same decision plan and the following “Balanced Transfers.”

$$\mathbb{B}_{i;t}(\tilde{t}) = \tau_{i;t}(\tilde{t}_i, \tilde{t}_i^{-1}) - \frac{1}{1-\delta} \sum_{j \neq i} \tau_{j;t}(\tilde{t}_j, \tilde{t}_j^{-1}); \text{ where} \quad (5.1)$$

$$\tau_{j;t}(\tilde{t}_j, \tilde{t}_j^{-1}) = -\tau \left(\mathbb{E}_{\pi_t}^{j \mid \tilde{t}_j, \tilde{t}_j^{-1}} \left[\Psi_j(\tilde{\cdot}) \right] - \mathbb{E}_{\pi_t}^{t \mid \tilde{t}_j, \tilde{t}_j^{-1}} \left[\Psi_j(\tilde{\cdot}) \right] \right) \quad (5.2)$$

(where conditioning on $\tilde{t}_i = \tilde{t}_i$ with $t = 0$ is interpreted as vacuous, i.e., no conditioning at all).

By construction, the Balanced Transfers are budget-balanced. Also, note that if the transfers in the original mechanism are uniformly bounded, then by Lemma 2 (in the Appendix) Ψ_j is LCI with some constant C , which implies that, for all t , $\left| \tau_{j;t}(\tilde{t}_j, \tilde{t}_j^{-1}) \right| \leq 2C$, and therefore the balanced transfers are uniformly bounded by $2C$:

To understand the Balanced Transfers, note $\tau_{j;t}$, interpreted as agent j ’s “incentive payment” in period t , gives the change in the expectation of agent j ’s present discounted transfers $-\delta^t \Psi_j(\tilde{\cdot})$ in the original mechanism that results from his report \tilde{t}_j in period t , given the

¹⁴We could interpret measures $\mu_t[\chi] \mid \theta^{t-1}$ and $\mu_t^i[\chi] \mid \theta_i^t, \theta_{-i}^{t-1}$ as conditional measures given θ^{t-1} and $(\theta_i^t, \theta_{-i}^{t-1})$, respectively. However, conditional probability measures are not well-defined at any given zero-probability point (this is the source of the famous Borel-Kolmogorov paradox – see, e.g., Proschan and Presnell (1998)), while these concepts are uniquely defined.

signal history θ^{t-1} .¹⁵ As in the standard AGV mechanism, all the other agents pitch in the same amount $\frac{1}{l-1} \gamma_{j;t}$ to pay agent j 's incentive payment. Subtracting all these payments that a given agent i has to make from his own incentive payment yields the balanced transfers (5.1).

Proposition 2. *Consider any mechanism $(\gamma; \Psi)$ in which there is a truthful-obedient BNE. If the agents' types are independent, then this is also a truthful-obedient BNE in the balanced mechanism $(\gamma; \Psi^B)$.*

Remark 1. *Recall that all that is important for truthful reporting to be a BNE is the present value of the transfers, which is given by*

$$\Psi_i^B(\gamma) = \sum_{t=1}^{\infty} \delta^{t-1} \left[\gamma_{i;t}(\theta_i^t, \theta_{-i}^{t-1}) - \frac{1}{l-1} \sum_{j \neq i} \gamma_{j;t}(\theta_j^t, \theta_{-j}^{t-1}) \right].$$

There are many ways other than (5.1) to spread $\Psi_i^B(\gamma)$ over time while ensuring budget balance in each period, and all these ways will preserve incentive compatibility.

Remark 2. *If decisions are made only up to period T , then the balanced team mechanism depends only on θ^T , and agents need not report θ_t for $t > T$ (they need not be observed).*

We define the *Balanced Team Mechanism* to be a mechanism with an efficient decision plan, together with balanced team transfers γ^{MB} constructed from γ^M using (5.1).

Corollary 1. *With independent types and private values, there is a truthtelling BNE of the Balanced Team Mechanism.*

Before proceeding with a formal proof of Proposition 2, we provide a sketch of the main ideas. We can write each agent j 's incentive payment in period t as

$$\gamma_{j;t}(\theta_j^t, \theta_{-j}^{t-1}) = \gamma_{j;t}^+(\theta_j^t, \theta_{-j}^{t-1}) - \gamma_{j;t}^-(\theta_{-j}^{t-1}); \text{ where}$$

$$\gamma_{j;t}^+(\theta_j^t, \theta_{-j}^{t-1}) = \mathbb{E}_{\theta_{-j}^{t-1}} \left[\Psi_j(\theta_j^t, \theta_{-j}^{t-1}) \right]; \quad \gamma_{j;t}^-(\theta_{-j}^{t-1}) = \mathbb{E}_{\theta_{-j}^{t-1}} \left[\Psi_j(\theta_{-j}^{t-1}) \right]; \quad (5.3)$$

The two terms are expectations of the same function $\Psi_j(\cdot)$, and differ only in the history used to construct beliefs, where $\gamma_{j;t}^-(\theta_{-j}^{t-1})$ uses only period $t-1$ information while $\gamma_{j;t}^+(\theta_j^t, \theta_{-j}^{t-1})$ uses, in addition, agent j 's period t information.

¹⁵In the case where all types in Θ have positive probability, so that the conditional expectations are uniquely defined everywhere, and also using (3.2), agent j 's incentive payment can be written as

$$\gamma_{j,t}(\theta_j^t, \theta_{-j}^{t-1}) = \sum_{\tau=t}^{\infty} \delta^{\tau-t} \left(\mathbb{E}_{\theta}^{\mu[x]} [\psi_{j,\tau}(\tilde{\theta}) | (\tilde{\theta}_j^t, \tilde{\theta}_{-j}^{t-1}) = (\theta_j^t, \theta_{-j}^{t-1})] - \mathbb{E}_{\theta}^{\mu[x]} [\psi_{j,\tau}(\tilde{\theta}) | \tilde{\theta}^{t-1} = \theta^{t-1}] \right).$$

Suppose that agent i deviates to a different strategy, while the other agents are truthful and obedient. The proof shows that, for any such deviation, the present expected value of agent i 's incentive payments $\bar{v}_{i;t}$ in the balanced mechanism equals, up to a constant, to that in the original mechanism (Claim 1). Also, the proof shows that the present expected value of the other agents' incentive payments $\bar{v}_{j;t}$ (for $j \neq i$) is zero regardless of agent i 's deviation (Claim 2). Thus, if the agent had no profitable deviation in the original mechanism, he will have none in the balanced mechanism.

The proofs of both Claim 1 and Claim 2 hinge on the assumption of independent types, which implies that for any agent $j \neq i$, the mechanism designer forms correct beliefs over the type $\tilde{v}_{j;t}$ of agent j in period t based on the public history in that period, \hat{h}^{t-1} , even when agent i has deviated from truthful-obedient behavior. Thus, for any public history \hat{h}^{t-1} and private history \hat{h}_i^t , agent i 's expectation of $\bar{v}_{j;t}^+(\tilde{v}_{j;t}; \hat{h}^{t-1})$ equals $\bar{v}_{j;t}^-(\hat{h}^{t-1})$. By the Law of Iterated Expectations, this implies that the ex ante expectation of all the other agents' incentive payments equals zero for any possible deviation of agent i , which establishes Claim 2.

To show Claim 1, we examine the present expected value of agent i 's own incentive payments following his deviation (which need not be zero, since the mechanism designer may have incorrect beliefs over agent i 's future signals following his deviation). For this purpose, we compare the terms $\bar{v}_{i;t}^+(\hat{h}_i^t; \hat{h}^{t-1})$ and $\bar{v}_{i;t+1}^-(\hat{h}^t)$, which differ only in that $\bar{v}_{i;t+1}^-$ incorporates period t reports from i 's opponents, $\hat{h}_{-i;t}$: Recall again that by independence of types, that the mechanism designer has correct beliefs over the opponents' reports $\tilde{v}_{-i;t}$ in period t based on the public history \hat{h}^{t-1} , despite agent i 's deviation. Then, the expectation of $\bar{v}_{i;t+1}^-(\hat{h}_i^t; \tilde{v}_{-i;t}; \hat{h}^{t-1})$ in period t is equal to $\bar{v}_{i;t}^+(\hat{h}_i^t; \hat{h}^{t-1})$, irrespective of agent i 's private history. By the Law of Iterated Expectations, the same two terms have the same ex ante expectations as well, regardless of agent i 's deviation. Using this, and for simplicity considering a finite horizon T , the present expected incentive payment of agent i is equal to the expectation of

$$\sum_{t=1}^T \left(\bar{v}_{i;t+1}^-(\tilde{v}^t) - \bar{v}_{i;t}^-(\tilde{v}^{t-1}) \right) = \bar{v}_{i;T+1}^-(\tilde{v}^T) - \bar{v}_{i;1}^- = \Psi_i(\tilde{v}^T) - \mathbb{E}_\sim [\Psi_i(\tilde{v}^T)];$$

where $\tilde{v} = \left(\tilde{v}_{-i}(\tilde{v}); \tilde{v}_{-i} \right)$ represents the agents' random reports following the deviation (with \tilde{v} distributed according to $\mathbb{P}[\tilde{v}]$, with \hat{v} defined in (3.3)). Thus, for any possible deviation of agent i , the present expected value of the agent's incentive payments in the balanced mechanism equals, up to a constant, the present expected payment in the original mechanism, which establishes Claim 1. The formal proof makes these arguments more precise and extends them to the general infinite-horizon case.

Proof of Proposition 2: First note that by the construction of the induced stochastic process $\tilde{v}^t[\tilde{v}] | \tilde{v}^{t-1}$ and by independence, the probability distribution of $\tilde{v}_{i;t}$ given history \tilde{v}^{t-1} can be written as $\tilde{v}_{i;t} | \tilde{v}_0^{t-1}(\tilde{v}^{t-1}); \tilde{v}_i^{t-1}(\tilde{v}^{t-1}); \tilde{v}_i^{t-1}$. Similarly, by independence of types the

probability distribution of $\tilde{\omega}_{-i;t}$ given history ω^{t-1} can be written as $\omega_{-i;t}^{t-1}(\omega^{t-1})$; $\omega_{-i}^{t-1} \equiv \prod_{j \neq i} \omega_{j;t} \left(\omega_0^{t-1}(\omega^{t-1}); \omega_j^{t-1}(\omega^{t-1}); \omega_j^{t-1} \right)$. To clarify observational measurability, we sometimes write $\omega_{-i}^{t-1}(\omega_{i;t}; \omega^{t-1})$ rather than $\omega_{-i}^{t-1}(\cdot)$:

Suppose that agent i deviates to induce an observationally measurable strategic plan $(\tilde{\omega}_{-i}; \tilde{\omega}_i)$, while the other agents use truthful-obedient strategies. The decision and transfer plans following the deviation are given by $\hat{\omega}$ and $\hat{\Psi}$ from (3.3) and (3.4), respectively.

Claim 1: Agent i 's expected discounted sum of incentive payments $\omega_{i;t}$ in the Balanced Mechanism equals his present expected transfer in the original mechanism from the same deviation plus a constant:

$$\mathbb{E}_{\omega_{-i}^{t-1}}^{[\hat{\omega}]} \left[\sum_{t=1}^{\infty} \omega_{i;t} \left(\omega_{-i}^{t-1}(\tilde{\omega}); \tilde{\omega}_{-i}^{t-1} \right) \right] = \mathbb{E}_{\omega_{-i}^{t-1}}^{[\hat{\omega}]} \left[\Psi_i \left(\omega_{-i}(\tilde{\omega}); \tilde{\omega}_{-i} \right) \right] - \mathbb{E}_{\omega_{-i}^{t-1}}^{[\omega]} \left[\Psi_i(\tilde{\omega}) \right]:$$

Proof of Claim 1: First, applying LIE, it follows that

$$\begin{aligned} \omega_{i;t} \left(\omega_{-i}^{t-1}(\omega_{i;t}; \omega^{t-1}); \omega_{-i}^{t-1} \right) &= \mathbb{E}_{\omega_{-i,t}^{t-1}}^{\omega_{-i,t}^{t-1}(\omega_{-i}^{t-1}(\omega_{i;t}; \omega^{t-1})); \omega_{-i}^{t-1}} \left[\omega_{i;t+1} \left(\omega_{-i}^{t-1}(\omega_{i;t}; \omega^{t-1}); \omega_{-i}^{t-1}(\omega_{i;t}; \omega^{t-1}) \right) \right] \\ &= \mathbb{E}_{\omega_{-i,t}^{t-1}}^{\omega_{-i,t}^{t-1}(\omega_{-i}^{t-1}(\omega_{i;t}; \omega^{t-1})); \omega_{-i}^{t-1}} \left[\omega_{i;t+1} \left(\omega_{-i}^{t-1}(\omega_{i;t}; \omega^{t-1}); \omega_{-i}^{t-1}(\omega_{i;t}; \omega^{t-1}) \right) \right]; \end{aligned} \quad (5.4)$$

where the second inequality follows by the definition of $\hat{\omega}$:

Now, taking expectations over $\tilde{\omega}_{i;t}$ with respect to probability distribution induced by $\omega_{-i}^{t-1}[\hat{\omega}]$, we can write

$$\begin{aligned} &\mathbb{E}_{\omega_{-i,t}^{t-1}}^{\omega_{-i,t}^{t-1}[\hat{\omega}]} \left[\omega_{i;t} \left(\omega_{-i}^{t-1}(\tilde{\omega}_{i;t}; \omega^{t-1}); \omega_{-i}^{t-1} \right) \right] \\ &= \mathbb{E}_{\omega_{-i,t}^{t-1}}^{\omega_{-i,t}^{t-1}[\hat{\omega}]} \mathbb{E}_{\omega_{-i,t}^{t-1}}^{\omega_{-i,t}^{t-1}(\omega_{-i}^{t-1}(\tilde{\omega}_{i;t}; \omega^{t-1})); \omega_{-i}^{t-1}} \left[\omega_{i;t+1} \left(\omega_{-i}^{t-1}(\tilde{\omega}_{i;t}; \omega^{t-1}); \omega_{-i}^{t-1}(\omega_{i;t}; \omega^{t-1}) \right) \right] \\ &= \mathbb{E}_{\omega_{-i,t}^{t-1}}^{\omega_{-i,t}^{t-1}[\hat{\omega}]} \left[\omega_{i;t+1} \left(\omega_{-i}^{t-1}(\tilde{\omega}_{i;t}; \omega^{t-1}); \omega_{-i}^{t-1}(\omega_{i;t}; \omega^{t-1}) \right) \right]: \end{aligned}$$

Thus,

$$\begin{aligned} &\mathbb{E}_{\omega_{-i,t}^{t-1}}^{\omega_{-i,t}^{t-1}[\hat{\omega}]} \left[\omega_{i;t} \left(\omega_{-i}^{t-1}(\tilde{\omega}_{i;t}; \omega^{t-1}); \omega_{-i}^{t-1} \right) \right] \\ &= \mathbb{E}_{\omega_{-i,t}^{t-1}}^{\omega_{-i,t}^{t-1}[\hat{\omega}]} \left[\omega_{i;t+1} \left(\omega_{-i}^{t-1}(\tilde{\omega}_{i;t}; \omega^{t-1}); \omega_{-i}^{t-1}(\omega_{i;t}; \omega^{t-1}) \right) \right] - \omega_{i;t} \left(\omega_{-i}^{t-1}(\omega_{i;t-1}; \omega^{t-2}); \omega_{-i}^{t-1} \right): \end{aligned}$$

Taking expectations over $\tilde{\omega}^{t-1}$ with respect to the probability distribution $\omega_{-i}^{t-1}[\hat{\omega}]$ and using LIE yields

$$\begin{aligned} &\mathbb{E}_{\omega_{-i}^{t-1}}^{[\hat{\omega}]} \left[\omega_{i;t} \left(\omega_{-i}^{t-1}(\tilde{\omega}_{i;t}; \tilde{\omega}_{-i}^{t-1}); \tilde{\omega}_{-i}^{t-1} \right) \right] = \Gamma_{t+1} - \Gamma_t, \\ &\text{where } \Gamma_t = \mathbb{E}_{\omega_{-i}^{t-1}}^{[\hat{\omega}]} \left[\omega_{i;t} \left(\omega_{-i}^{t-1}(\tilde{\omega}_{i;t-1}; \tilde{\omega}_{-i}^{t-1}); \tilde{\omega}_{-i}^{t-1} \right) \right]: \end{aligned}$$

Thus, we can write

$$\begin{aligned} &\mathbb{E}_{\omega_{-i}^{t-1}}^{[\hat{\omega}]} \left[\sum_{t=1}^{\infty} \omega_{i;t} \left(\omega_{-i}^{t-1}(\tilde{\omega}_{i;t}; \tilde{\omega}_{-i}^{t-1}); \tilde{\omega}_{-i}^{t-1} \right) \right] = \lim_{T \rightarrow \infty} \sum_{t=1}^T (\Gamma_{t+1} - \Gamma_t) \\ &= \lim_{T \rightarrow \infty} \Gamma_{T+1} - \Gamma_1 = \mathbb{E}_{\omega_{-i}^{t-1}}^{[\hat{\omega}]} \left[\Psi_i \left(\omega_{-i}(\tilde{\omega}); \tilde{\omega}_{-i} \right) \right] - \mathbb{E}_{\omega_{-i}^{t-1}}^{[\omega]} \left[\Psi_i(\tilde{\omega}) \right]; \end{aligned}$$

where the first equality obtains from the previous derivation using the Lebesgue Dominated Convergence Theorem and the LCI of Ψ_i , the second by canceling intermediate terms in the partial sum, and the third by the Dominated Convergence Theorem, the LCI of Ψ_i , and the fact that distribution $\tau_{t+1}[\cdot|\cdot]\left(\bar{\tau}_i\left(\bar{\tau}_i; \bar{\tau}_i^{-1}\right); \bar{\tau}_i\right)$ puts probability 1 on the event $\tilde{\tau}^T = \left(\bar{\tau}_i\left(\bar{\tau}_i; \bar{\tau}_i^{-1}\right); \bar{\tau}_i\right)$.

Now consider a second claim:

Claim 2: The present expected value of agent j 's incentive payments $(\tau_{j;t})$ is equal to zero when agent j follows a truthtelling strategy no matter what reporting strategies the other agents use:

$$\mathbb{E}_{\tilde{\tau}}^{[\wedge]} \left[\sum_{t=1}^{\infty} \tau_{j;t} \left(\tilde{\tau}_{j;t}^{t-1, -t-1} \left(\tilde{\tau}^{t-1} \right) \right) \right] = 0;$$

Proof of Claim 2: For any $j \neq i$; let $\bar{\tau}_{-j}(\cdot) = (\bar{\tau}_i(\cdot); \bar{\tau}_{-i-j})$ denote the reporting strategy of agents other than j . We write $\bar{\tau}_{-j}^t(\cdot)$ for the reporting strategy up to period t . Using (5.2), we can write for any t , and any $\bar{\tau}^{t-1}$,

$$\begin{aligned} & \mathbb{E}_{\tilde{\tau}, t}^{[\wedge]} \left[\tau_{j;t} \left(\left(\tilde{\tau}_{j;t}^{t-1, -t-1} \left(\bar{\tau}^{t-1} \right) \right) \right) \right] \tag{5.5} \\ &= \mathbb{E}_{\tilde{\tau}, t}^{[\wedge]} \left[\tau_{j;t}^{t-1} \mathbb{E}_{\bar{\tau}_{-j,t}^{t-1}} \left[\tau_{j;t}^{t-1, -t-1} \left(\bar{\tau}^{t-1} \right); \bar{\tau}_{-j}^{t-1} \left(\bar{\tau}^{t-1} \right) \right] \mathbb{E}_{\tilde{\tau}}^{t+1} \left[\tau_{j;t}^{t-1, -t-1} \left(\bar{\tau}^{t-1} \right); \tilde{\tau}_{j,t}; \bar{\tau}_{-j,t} \right] \left[\Psi_j \left(\tilde{\tau} \right) \right] \right] \\ &= \mathbb{E}_{\bar{\tau}_{-j,t}^{t-1}} \left[\tau_{j;t}^{t-1} \left(\bar{\tau}_{-j}^{t-1, -t-1} \left(\bar{\tau}^{t-1} \right) \right); \bar{\tau}_{-j}^{t-1} \left(\bar{\tau}_{-j}^{t-1, -t-1} \left(\bar{\tau}^{t-1} \right) \right); \bar{\tau}_{-j}^{t-1} \left[\begin{array}{l} \mathbb{E}_{\bar{\tau}_{-j,t}^{t-1}} \left[\tau_{j;t}^{t-1, -t-1} \left(\bar{\tau}^{t-1} \right); \bar{\tau}_{-j}^{t-1} \left(\bar{\tau}^{t-1} \right) \right] \\ \mathbb{E}_{\tilde{\tau}}^{t+1} \left[\tau_{j;t}^{t-1, -t-1} \left(\bar{\tau}^{t-1} \right); \bar{\tau}_{-j,t} \right] \left[\Psi_j \left(\tilde{\tau} \right) \right] \end{array} \right] \right] \\ &= \mathbb{E}_{\tilde{\tau}}^{t-1} \left[\tau_{j;t}^{t-1, -t-1} \left(\bar{\tau}^{t-1} \right) \right] \mathbb{E}_{\tilde{\tau}}^{t+1} \left[\tau_{j;t}^{t-1, -t-1} \left(\bar{\tau}^{t-1} \right); \tilde{\tau}_{j,t} \right] \left[\Psi_j \left(\tilde{\tau} \right) \right] \\ &= \mathbb{E}_{\tilde{\tau}}^{t-1} \left[\tau_{j;t}^{t-1, -t-1} \left(\bar{\tau}^{t-1} \right) \right] \left[\Psi_j \left(\tilde{\tau} \right) \right] = \tau_{j;t} \left(\tilde{\tau}_{j;t}^{t-1, -t-1} \left(\bar{\tau}^{t-1} \right) \right); \end{aligned}$$

The first equality involves substitutions and definitions, and the LIE. The second is by definition of $\tau_{j;t}^{[\wedge]}|\bar{\tau}^{t-1}$ and $\tilde{\tau}$ and independence of types (which implies that the distribution of $\tilde{\tau}_{j;t}$ in $\tau_{j;t}^{[\wedge]}|\bar{\tau}^{t-1}$ depends only on the public decision history and on the private history of agent j). The third follows by definition of $\tau_{j;t}|\bar{\tau}^{t-1}$, LIE and independent types. The fourth uses LIE and definitions, and the fifth uses definitions. By LIE and the fact that $\tau_{j;t} = \tau_{j;t}^+ - \tau_{j;t}^-$, this implies that

$$\mathbb{E}_{\tilde{\tau}}^{[\wedge]} \left[\tau_{j;t} \left(\tilde{\tau}_{j;t}^{t-1, -t-1} \left(\tilde{\tau}^{t-1} \right) \right) \right] = \mathbb{E}_{\tilde{\tau}}^{[\wedge]} \left[\mathbb{E}_{\tilde{\tau}}^{t[\wedge]|\tilde{\tau}^{t-1}} \left[\tau_{j;t} \left(\tilde{\tau}_{j;t}^{t-1, -t-1} \left(\tilde{\tau}^{t-1} \right) \right) \right] \right] = 0;$$

Then, the Lebesgue Dominated Convergence Theorem and LCI of Ψ_j imply that

$$\mathbb{E}_{\tilde{\tau}}^{[\wedge]} \left[\sum_{t=1}^{\infty} \tau_{j;t} \left(\tilde{\tau}_{j;t}^{t-1, -t-1} \left(\tilde{\tau}^{t-1} \right) \right) \right] = \sum_{t=0}^{\infty} \mathbb{E}_{\tilde{\tau}}^{[\wedge]} \left[\tau_{j;t} \left(\tilde{\tau}_{j;t}^{t-1, -t-1} \left(\tilde{\tau}^{t-1} \right) \right) \right] = 0;$$

completing the proof of Claim 2.

So, adding up, we have

$$\mathbb{E}_{\sim}^{[\wedge]} \left[\Psi_i^B \left(\sigma_i(\tilde{\cdot}); \tilde{\sigma}_{-i} \right) \right] = \mathbb{E}_{\sim}^{[\wedge]} \left[\Psi_i \left(\sigma_i(\tilde{\cdot}); \tilde{\sigma}_{-i} \right) \right] - \mathbb{E}_{\sim}^{[1]} \left[\Psi_i(\tilde{\cdot}) \right]: \quad (5.6)$$

Since the second term does not depend on agent i 's strategy $(\sigma_i; \sigma_{-i})$, the expected transfer to agent i for any strategy σ_i chosen by him is equal, up to a constant, to the expected transfer under the original mechanism, assuming that other agents use truthful-obedient strategies. Thus, if the original mechanism has a truthtelling BNE, it remains a BNE in the balanced mechanism. ■

6. Efficient Equilibria in Markov Games with Private States

In this section, we apply the results to show a Folk-Theorem like result for Markov games with private states — namely, that efficiency can be sustained in equilibrium when the discount factor is close enough to 1. In the games, there is no external enforcer, and so the decisions and payments must be self-enforcing. For simplicity we restrict attention to games with transfers so that our results can be applied directly.¹⁶ The requirement of decisions and transfers to be self-enforcing means that we need to consider deviations in which they not only misreport, but also choose off-equilibrium decisions and transfers. Such deviations can trigger a permanent “punishment phase” of the game. Making sure that such deviations are unprofitable is similar to ensuring participation constraints in the mechanism where agents can choose to exit in each period. In a static context or in a model with very impatient agents, it is well known that whether or not participation constraints can be satisfied depends on the structure of the game. Here, we focus instead on the case where agents are patient and the horizon is infinite, and identify conditions under which participation constraints could be satisfied in this model. In a follow-on paper (Athey and Segal (2007)), we explore a special case of our model in greater depth, providing weaker sufficient conditions for participation constraints to hold.

To see the problem with guaranteeing participation constraints even with patient agents, note that the balanced transfers in each period are a function of discounted sums of future payoffs, and so they could potentially grow with the discount factor at a rate proportional to future surplus. Thus, in general the future surplus may not be enough to guarantee that all agents will wish to make the necessary transfers. In this section, we provide conditions under which the balanced transfers remain bounded as patience increases, and so the growing surplus from future cooperation can be used to make these transfers self-enforcing.

Formally, we consider an infinite-horizon game where each period t consists of three stages:

Stage $t:0$: Each agent i simultaneously observes a private signal $\theta_{i;t} \in \Theta_{i;t}$,

Stage $t:1$: Each agent i sends a report $\hat{\theta}_{i;t} \in \Theta_{i;t}$.

¹⁶In games without transfers, one could in principle extend the approach of Fudenberg, Levine, and Maskin (1994), using small changes in continuation play mimic the role of transfers.

Stage $t:2$: Each agent i chooses a publicly observable action $x_{0;i;t} \in X_{0;i;t}$; a privately observed action $x_{i;t} \in A_{i;t}$, and a publicly observable payment $z_{i;j;t} \geq 0$ to each agent j . (Note that the payments are not wasted by construction).

Letting $x_{0;t} = (x_{0;i;t} ::: x_{0;l;t})$, the collective decision at time t is $x_t \equiv (x_{0;t} x_{1;t} ::: x_{l;t}) \in \prod_{i=0}^l X_{i;t} \equiv X_t$, and the total payment received by agent i in period t is $y_{i;t} = \sum_{j=1}^l (z_{j;i;t} - z_{i;j;t})$.

The payoff of agent i in this game takes the same form as before,

$$\sum_{t=1}^{\infty} \delta^t (u_{i;t}(x^t; \tau^t) + y_{i;t}) :$$

We want to construct an efficient equilibrium for this game, i.e., one that implements an observationally measurable decision plan maximizing the expected joint surplus (4.1). We construct an equilibrium in which agents use truthful-obedient reporting strategies. Let τ denote the transfer plan sustained in the equilibrium.

One class of deviations is that an agent uses a different reporting strategy and/or private decision policy, but uses the same public decision and payment strategies. These deviations are unprofitable as long as the reporting strategies and private decision strategies form a BNE of the mechanism with enforcement. However, with self-enforcement, we also need to deter “off-schedule deviations,” which involve, in addition to possible misreporting and deviation in private decisions, choosing public decisions or payments different from those required by the equilibrium strategies given past reports. Off-schedule deviations are publicly observable as deviations, and as soon as they are observed they can be punished by the agents switching to a “punishment” decision plan.

To deter off-schedule deviations, we need to make sure that the benefits of changing a decision or withholding a payment in one period are outweighed by the future gain from cooperation. As mentioned above, the balanced transfers could potentially grow large as the discount factor increases, since they depend on the expected discounted value of future utility, and so in general the payments are not guaranteed to be smaller than the future gains from cooperation. However, the balanced transfers are defined as functions of the *change* in future expected discounted utility for a particular agent induced by the arrival of a single agent’s signal in a particular period. We now specialize the model in a way that guarantees that a particular period’s signal and decisions have vanishing effect on the continuation expected utility. This will guarantee that the transfers stay within a fixed bound even as the discount factor grows, while the present value of cooperation grows without a bound, hence off-schedule deviations are deterred.

Assumptions

- (i) $\Theta_t = \bar{\Theta}$ and $X_t = \bar{X}$ for all t , and all the sets are finite.

- (ii) Per-period utilities $u_{i,t}(x^t; \theta^t)$ depend only on period- t variables and do not depend on t directly, and so can be written as $u_i(x_t; \theta_t)$.
- (iii) The contingent probability measures $\pi_t|x^{t-1}; \theta^{t-1}$ depend only on period $t - 1$ variables and do not depend on t directly, and so can be written as $\pi|x_{t-1}; \theta_{t-1}$ (where for $t = 1$ we take some fixed $x_0; \theta_0$).
- (iv) Types are independent (Definition 1) and values are private (Definition 3).
- (v) There exists $\bar{x}_t \in X_t$ such that $u_i(x_{i,t}; \bar{x}_{-i,t}; \theta^t) = 0$ for all i , all $x_{i,t} \in \bar{X}$, and all $\theta^t \in \bar{\Theta}$.

Assumptions (i)-(iii) imply that we have a Markov Decision Process (MDP), as studied, e.g., in Puterman (1994).¹⁷ In such problems, it is natural to focus on Markov policies. Furthermore, it turns out that there exists a single Markov policy that is optimal for β close enough to 1:

Definition 5. *Decision policy σ is a Markov policy if each $\sigma_t(\theta^t)$ depends only on the last state θ^t and does not depend directly on t , and so can be written as $\sigma(\theta^t)$. Decision policy σ^* is a Blackwell policy if it is a Markov policy and there exists $\bar{\beta} < 1$ such that for any $\beta \in (\bar{\beta}; 1)$, σ^* achieves a weakly greater present expected total surplus starting at any state than any other observationally measurable decision policy σ .*

Puterman (1994) shows that a Blackwell policy exists (Theorem 10.1.4) and describes an algorithm that constructs it with a finite number of steps (Corollary 10.3.7, p.518).

Assumption (v) can be understood by interpreting decision $\bar{x}_{i,t} \in X_{i,t}$ as “non-participation” by agent i and saying that an agent’s utility is fixed when no other agent participates. For simplicity, the fixed utility is normalized to zero. This ensures that all agents choosing non-participation and making zero payments is a BNE of the stage game regardless of agents’ types and beliefs about others’ types. This allows us to use such BNE as a punishment for the observed off-schedule deviations.¹⁸

Observe that in the Markov game, the team transfers can be written as $t_{i,t}(\theta^t) = \sum_{j \neq i} u_j(\sigma^*(\theta^t); \theta^t)$, and the balanced team transfers can be written as functions of the current and most recent

¹⁷Note that dependence on $k > 1$ periods of history could be incorporated by expanding the sets Θ_i and X to remember k periods of past types and decisions.

¹⁸More generally, in particular games without a natural nonparticipation option, there may be punishment equilibria that can be used in place of nonparticipation; however, if the equilibria depend on the initial state of the game, one must guarantee the existence of an equilibrium for each possible state of the game. Equilibria where decisions do not depend on the state (such as the pooling equilibria constructed by Athey and Bagwell (2004) in their study of dynamic collusion with persistent private information) are especially convenient to analyze in this context.

states as follows:

$$\begin{aligned} B_i^B(\tau-1; \tau) &= B_i^B(i; \tau-1) - \frac{1}{I-1} \sum_{j \neq i} B_j^B(j; \tau-1); \text{ where for each } j; \\ B_j^B(j; \tau-1) &= \sum_{k \neq j} \sum_{t=\tau}^{\infty} -t \left\{ \begin{array}{l} \mathbb{E}_{\sim}^{j|t^*} [u_k^*(\sim); \sim] \\ -\mathbb{E}_{\sim}^{t|t^*} [u_k^*(\sim); \sim] \end{array} \right\}; \end{aligned} \quad (6.1)$$

Thus, in the Markov environment, period t transfers have the attractive feature that they depend only on the announcements in periods t and $t-1$; since these announcements fully determine beliefs.

Any Markov decision policy induces a Markov chain on states with transition probabilities described by $\pi^*(\tau-1; \tau-1)$. We assume that the process induced by the Blackwell policy has limited history dependence, in the sense of having a single ergodic set:

Definition 6. *Given a Markov process, a set of states $S \subseteq \bar{\Theta}$ is called ergodic if, starting in S , the system remains there with probability one, and this is not true for any proper subset of S .*

Proposition 3. *Consider a dynamic game satisfying assumptions (i)-(v) above. Let π^* be a Blackwell decision policy for the game, and suppose that the Markov process induced by π^* has a single ergodic set (plus a possibly empty transient set), and that its invariant distribution yields a positive expected total surplus. Then there exists $\delta^* < 1$ such that, for all $\delta \in (\delta^*; 1)$; the game has an efficient PBE that sustains decision plan π^* .*

Proof. We focus on those δ close enough to 1 for which π^* is a Blackwell policy. Define “on-schedule” public histories as those in which all agents in all period have made public decisions and payments consistent with the equilibrium strategies given the public histories, and call any other public histories “off-schedule” histories. We construct an equilibrium in which strategies satisfy the following for each period t and each agent i :

- For an off-schedule history, in Stage $t:1$, the agent reports his true type. In Stage $t:2$, he chooses the nonparticipation action $\bar{x}_{i;t}$, and makes zero payments to all other agents.
- For an on-schedule history, in Stage $t:1$, the agent reports his true signal $\theta_{i;t}$. In Stage $t:2$, given the current reports $\hat{\theta}_t$, he chooses his public action $x_{0;i;t} = x_{0;i}^*(\hat{\theta}_t)$; he chooses his private action $x_{i;t} = x_i^*(\hat{\theta}_t)$ if in Stage $t:1$ he reported truthfully $\hat{\theta}_{i;t} = \theta_{i;t}$, otherwise he chooses a private action $x_{i;t}$ that maximizes the present expected value of his utility given the belief that in the future all agents are truthful and obedient. Also, the agent makes payments determined from the reports in periods $t-1$ and t as follows:

$$z_{i;j;t}(\hat{\theta}_{i;t}; \hat{\theta}_{t-1}) = \frac{1}{I-1} \left[B_j^B(\hat{\theta}_{j;t}; \hat{\theta}_{t-1}) \right] + A_i \text{ to all agents } j \neq i;$$

where the constants $A_i \geq \frac{1}{1-\beta} \max_{i;(t-1:i,t)} |B_i(i;t,t-1)|$ are specified below.

Observe that if all agents follow these strategies, we will only have on-schedule histories, and the agents will always be truthful and obedient, which implements the Blackwell decision plan π^* . Now we show that these strategies are part of a PBE, for some appropriately chosen constants A_i , when β is close enough to 1. The accompanying PBE beliefs of each agent i in each period t about the other agents' current types if there has been no off-schedule deviation are described by $\Pi_{j \neq i} | \pi^* \left(\hat{\theta}_{t-1} \right); \hat{\theta}_{j;t-1}$ (i.e., the agent believes that the other agents were truthful and obedient in the previous period). The beliefs following an off-schedule deviation are described by $\Pi_{j \neq i} | \chi_{0;j;t-1}; \hat{\theta}_{j;t-1}; \hat{\theta}_{j;t-1}$. It is easy to see that these beliefs satisfy PBE conditions (i)-(iii) of Subsection 3.3. It remains to show that the strategies described above form a best response to each other given these beliefs at any information set.

First, it is easy to see that the strategies form a best response to each other for any off-schedule public history, since at this point any possible strategy of agent i will give him a nonpositive expected continuation payoff regardless of his beliefs. Now consider the information sets with on-schedule public histories. First we argue that at such information sets, there is no profitable ‘‘on-schedule’’ deviation for each agent i , i.e., one in which he makes the public decisions and payments prescribed by the equilibrium strategies given the reports. For information sets in the beginning of a period t , this holds because the Blackwell policy by definition must be efficient starting in period t given any initial beliefs, and so together with the total prescribed payments

$$\sum_{j \neq i} \left[z_{j;i;t} \left(\hat{\theta}_{j;t}; \hat{\theta}_{t-1} \right) - z_{i;j;t} \left(\hat{\theta}_{i;t}; \hat{\theta}_{t-1} \right) \right] = B_i \left(\hat{\theta}_{t-1}; \hat{\theta}_t \right) - (1-\beta) A_i + \sum_{j \neq i} A_j$$

the agent faces a Balanced Team Mechanism, in which truthfulness and obedience are optimal for him within the class of on-schedule strategies according to Proposition 2. In the middle of a period t after reports are made but before decisions are chosen, agent i has no profitable on-schedule deviation by the Principle of Dynamic Programming, since we already know that truthfulness and obedience will be optimal for him starting in period $t+1$ no matter what private action $x_{i;t}$ he chooses now, and his choice of $x_{i;t}$ is optimal given truthfulness and obedience in the future.

Now we need to show that, for some appropriately chosen constants A_i , each agent has no profitable ‘‘off-schedule’’ deviations, i.e., those in which he plans a public decision and/or payment that differ from the equilibrium ones prescribed given the reporting histories, at any on-schedule information set. We do this by showing that if at any future information set the agent plans a first off-schedule public action and/or payment, we can modify his strategy to be truthful and obedient starting from this information set without reducing his expected continuation payoff. By doing it for any information set at which a first off-schedule deviation is planned, we would obtain an on-schedule deviation that is at least as profitable as the original off-schedule deviation. Since we already know that there are no profitable on-schedule

deviations, this will complete the proof.

We choose the constants A_i so that the expected surplus in the game is shared equally among the agents, and since the surplus is positive and very large from any starting state when the agents are very patient, this ensures that the agents will want not to “quit” the game by resorting to an off-schedule deviation. For this purpose, recall that since by assumption the Markov process induced by $\tilde{\pi}^*$ has a single ergodic set, it induces a unique invariant distribution over $\tilde{\mathbf{x}}_t$ (Stokey and Lucas, 1989), which we denote by $\tilde{\mu}$. Let \bar{u}_i denote the expected per-period utility of agent i in this distribution:

$$\bar{u}_i \equiv \mathbb{E}_{\tilde{\mu}} \left[u_i \left(\tilde{\pi}^* \left(\tilde{\mathbf{x}}_t \right); \tilde{\mathbf{x}}_t \right) \right]:$$

Since the expectation of \tilde{z}_j^B is zero for all $j \neq i$, we let $A_i = \bar{u}_i + K$, where the constant K should be large enough to ensure that $z_{i,j;t}(\alpha) \geq 0$ for all α . We show that such K can be found independently of α , for which purpose we demonstrate that the transfers $\tilde{z}_j^B(\alpha; i; t-1)$ are bounded uniformly in α . We do it by making use of the following lemma, whose proof (in the Appendix) relies on the structure of Markov chains with a single ergodic subset - see, e.g., Behrens (2000, Chapter 7) and Stokey and Lucas (1989):

Lemma 1. *Take a finite Markov chain with a single ergodic set, let $\mathbf{p}_{\cdot; \cdot}^{(t)}$, $\equiv \Pr\{\tilde{\mathbf{x}}_t = \cdot' | \tilde{\mathbf{x}}_0 = \cdot\}$ (the t -step transition probabilities), and let $\tilde{\mu}$ denote the invariant distribution. The expression $\left| \sum_{t=0}^{\infty} \alpha^t \left(\mathbf{p}_{\cdot; \cdot}^{(t)} - \tilde{\mu} \right) \right|$ is bounded uniformly across all $\alpha \in (0; 1)$ and all $\cdot; \cdot' \in \bar{\Theta}$.*

The Lemma implies that there exists $C < \infty$ such that for all $\alpha \in \bar{\Theta}$, $\alpha \in (0; 1)$:

$$\left| \mathbb{E}_{\tilde{\mu}} \left[\alpha^t \left\{ u_j \left(\tilde{\pi}^* \left(\tilde{\mathbf{x}}_t \right); \tilde{\mathbf{x}}_t \right) - \bar{u}_j \right\} \right] \right| \leq C: \quad (6.2)$$

In particular, (6.2) implies that the transfers are bounded:

$$\begin{aligned} \left| \tilde{z}_j^B(\alpha; i; t-1) \right| &\leq \left| \mathbb{E}_{\tilde{\mu}} \left[\alpha^{j[t-1]} \left\{ u_j \left(\tilde{\pi}^* \left(\tilde{\mathbf{x}}_{j,t-1} \right); \tilde{\mathbf{x}}_{j,t-1} \right) - \bar{u}_j \right\} \right] \right| \\ &\quad + \left| \mathbb{E}_{\tilde{\mu}} \left[\alpha^{t-1} \left\{ u_j \left(\tilde{\pi}^* \left(\tilde{\mathbf{x}}_{j,t} \right); \tilde{\mathbf{x}}_{j,t} \right) - \bar{u}_j \right\} \right] \right| \\ &\leq 2 \sup_t \left| \mathbb{E}_{\tilde{\mu}} \left[\alpha^t \left\{ u_j \left(\tilde{\pi}^* \left(\tilde{\mathbf{x}}_{j,t} \right); \tilde{\mathbf{x}}_{j,t} \right) - \bar{u}_j \right\} \right] \right| \leq 2C; \end{aligned}$$

where the second inequality uses the Law of Iterated Expectations and the fact that $|\mathbb{E}g(\mathbf{X})| \leq \sup |g(\mathbf{X})|$. Therefore, letting $K = 2C = (1 - \alpha) + M = 1$, where $M = \max_{i; \mathbf{x}_t; i, t} |u_i(\mathbf{x}_t; t)|$, ensures that $0 \leq z_{i,j;t}(\alpha; i; t-1) \leq 2K$ for any $\alpha \in (0; 1)$ and any $(j; t-1)$.

Now, given these constants, compare the expected continuation payoff in period t from making a first off-schedule deviation in period t with that from reverting to truthfulness and

obedience in period t . An off-schedule deviation will bring nonpositive payoffs starting in period $t + 1$, while the expected continuation payoff from truthfulness and obedience in period $t + 1$ conditional on any starting state \tilde{x}_t can be bounded below as follows:

$$\begin{aligned}
& \sum_{\tau=t+1}^{\infty} \mathbb{E}_{\tau}^{[*] | t} \left[u_i \left(* \left(\tilde{x}_{\tau} \right); \tilde{x}_{\tau} \right) + \beta_i \left(\tilde{x}_{\tau-1}; \tilde{x}_{\tau} \right) + \frac{1}{I} \sum_j \bar{u}_j - \bar{u}_i \right] \\
&= \sum_{\tau=t+1}^{\infty} \mathbb{E}_{\tau}^{[*] | t} \left[u_i \left(* \left(\tilde{x}_{\tau} \right); \tilde{x}_{\tau} \right) + \frac{1}{I} \sum_j \bar{u}_j - \bar{u}_i \right] \\
&\geq \frac{1}{1-\beta} \frac{1}{I} \sum_j \bar{u}_j - \left| \sum_{\tau=t+1}^{\infty} \left\{ \mathbb{E}_{\tau}^{[*] | t} \left[u_i \left(* \left(\tilde{x}_{\tau} \right); \tilde{x}_{\tau} \right) \right] - \bar{u}_i \right\} \right| \\
&\geq \frac{1}{1-\beta} \left[\frac{1}{I} \sum_i \bar{u}_i - C \right]; \tag{6.3}
\end{aligned}$$

where the equality obtains since

$$\mathbb{E}_{\tau}^{[*] | t} \left[\beta_i \left(\tilde{x}_{\tau-1}; \tilde{x}_{\tau} \right) \right] = 0$$

for $\tau \geq t + 1$ by the Law of Iterated Expectations, and the last inequality is by (6.2). This should be weighed against the possible gain in period t from the off-schedule deviation rather than being truthful and obedient in this period, which is at most

$$\beta^t \left(I \sup_{i,j;t; \tilde{x}_{t-1}; \tilde{x}_t} z_{i,j;t} \left(\tilde{x}_t; \tilde{x}_{t-1} \right) + 2 \max_{i; \tilde{x}_t} |u_i \left(\tilde{x}_t; \tilde{x}_t \right)| \right) = \beta^t (I \cdot 2K + 2M);$$

Thus, we see that for $\beta = (1 - \beta) > I(C + 2IK + 2M) = \left(\sum_i \bar{u}_i \right)$, the off-schedule deviation is dominated by becoming truthful and obedient. Together with the previous argument, this shows that the proposed strategies and beliefs form a PBE for β close enough to 1. ■

We conclude this section with a couple of observations about how the results can be generalized. First, we restricted attention to efficient allocation rules. We could also consider other decision rules that could be implemented using transfers (not necessarily budget-balanced) that are uniformly bounded in β and that have a Markov structure. Then, our proof can be modified to establish that the balanced transfers satisfy participation constraints for sufficiently patient agents.

Second, the Markov structure is more restrictive than necessary. The important feature of the Markov structure is that today's signals and reports have a vanishing impact on the expected value of future payoffs and transfers. This feature could be preserved in models that are nonstationary and have dependence on longer and possibly infinite histories, provided that this dependence is bounded.

7. Conclusions

This paper generalizes to a dynamic setting the results of Arrow (1979) and d’Aspremont and Gerard-Varet (1979) that, given an allocation rule and incentive-compatible transfers, it is possible to construct budget-balanced and incentive-compatible transfers. The result holds for a very general class of dynamic games, with arbitrary intertemporal dependence of agents’ signals, arbitrary effects of the public decision on the distribution of signals and their informativeness, and privately observed decisions (moral hazard problems) so long as the decisions do not have direct externalities on other agents. The main substantive restrictions are additive separability of payoffs in money, and independence of types across agents (conditional on past public decisions). Although the analysis of participation constraints depends on the structure of the specific game, for agents in an infinite-horizon game with a Markov structure and a single ergodic set, participation constraints can be satisfied when agents are sufficiently patient.

In future work we plan to consider models of costly computation and communication. Some problems of computation costs can be considered as special cases of our model, by modeling an agent’s computation of her type as information gathering.¹⁹ We plan to extend our model to include the case where agents send reports from a restricted message space and so cannot fully reveal their information (e.g., because communication is costly), as in Fadel and Segal (2006).

We also intend to study the case of correlated values more closely. In static mechanism design, the case of independent types is viewed as the “hardest” case for mechanism design; with correlated types, efficiency with budget balance is achievable under generic conditions, as shown by d’Aspremont, Cremer, and Gerard-Varet (2003a,b). We expect that these results could be extended to dynamic mechanism design, but this would require different techniques than the ones we have used, and so we leave this extension for future research.²⁰

References

- [1] Akerlof, G. (1970), “The Market for ‘Lemons’: Quality Uncertainty and the Market Mechanism,” *The Quarterly Journal of Economics* 84(3), 488-500.
- [2] Athey, S., and K. Bagwell (2004), “Collusion with Persistent Cost Shocks,” mimeo, Stanford University, 2004.
- [3] Athey, S., and D. Miller (2004), “Efficiency in repeated trade with hidden valuations,” mimeo, Stanford University.

¹⁹See, e.g., Larson and Sandholm (2001) for such a model. Other kinds computational costs fall outside of our model, such as the agents’ costly information acquisition about others’ types, or the cost of computing the decisions and transfers by the mechanism.

²⁰For a two-period setting with a period of private actions followed by a period of private signals, budget-balanced implementation of approximate efficiency that exploits correlation is examined by Obara and Rahman (2006).

- [4] Athey, S. and I. Segal (2007), “Designing Efficient Mechanisms for Dynamic Bilateral Trading Games,” *American Economic Review Papers and Proceedings*, forthcoming.
- [5] Atkeson, A. and R. Lucas (1993), “On Efficient Distribution with Private Information,” *Review of Economic Studies*, 59, 427-453.
- [6] Arrow, K. (1979), “The property rights doctrine and demand revelation under incomplete information.” In M Boskin, ed., *Economics and human welfare*. New York Academic Press.
- [7] d’Aspremont, C., J. Cremer, and Gerard-Varet L. (2003a) “Balanced Bayesian Mechanisms.” *Journal of Economic Theory*, 115, 385-396.
- [8] d’Aspremont, C., J. Cremer, and Gerard-Varet L. (2003b), “Correlation, Independence, and Bayesian Incentives,” *Social Choice and Welfare* 21(2), 281-310.
- [9] d’Aspremont, C. and Gerard-Varet, L. (1979) “Incentives and incomplete information.” *Journal of Public Economics*, 11:25–45.
- [10] Bapna, A., and T. Weber (2005) “Efficient Dynamic Allocation with Uncertain Valuations,” working paper, Stanford University
- [11] Battaglini, M. (2003), “Optimality and Renegotiation in Dynamic Contracting,” Mimeo, Princeton University.
- [12] Battaglini, M. (2005), “Long-Term Contracting with Markovian Consumers,” *American Economic Review* 95(3), pages 637-658.
- [13] Behrends, Ehrhard (2000). *Introduction to Markov Chains, with Special Emphasis on Rapid Mixing*. Vieweg, 2000.
- [14] Bergemann, Dirk, and Juuso Valimaki (2006), “Efficient Dynamic Auctions,” Cowles Foundation Discussion Paper 1584
- [15] Bester, H. and Strausz, R. (2001), “Contracting with Imperfect Commitment and the Revelation Principle: The Single Agent Case,” *Econometrica*, 2001, 69 (4), pp. 1077–1098.
- [16] Courty, P., and H. Li (2000), “Sequential Screening,” *Review of Economic Studies*, 67(4), page 697-717.
- [17] Cremer, J., Y. Spiegel, and C. Zheng (2003), “Optimal Selling Mechanisms with Costly Information Acquisition,” IDEI Working Paper.
- [18] Ericson, R. and A. Pakes (1995), “Markov-Perfect Industry Dynamics: A Framework for Empirical Work,” *Review of Economic Studies* 62 (1), 53-82.

- [19] Fadel, R., and I. Segal (2006), “The Communication Cost of Selfishness,” working paper, Stanford University.
- [20] Friedman, E., and D. C. Parkes (2003), “Pricing WiFi at Starbucks—Issues in Online Mechanism Design,” *Fourth ACM Conference on Electronic Commerce*, pp. 240-241.
- [21] Fudenberg, D., D. Levine, and E. Maskin, “The Folk Theorem with Imperfect Public Information,” *Econometrica*, 62, 997-1039.
- [22] Fudenberg, D. and J. Tirole (1991a), *Game Theory*, MIT Press, Cambridge, MA.
- [23] Fudenberg, D. and J. Tirole (1991b), “Perfect Bayesian Equilibrium and Sequential Equilibrium,” *Journal of Economic Theory*, 53, 236-260.
- [24] Gershkov, A., and B. Szentes (2004), “Optimal Voting Scheme with Costly Information Acquisition,” mimeo, University of Chicago.
- [25] Jehiel, P. and B. Moldovanu (2001), “Efficient Design with Interdependent Valuations,” *Econometrica* 69: 1237-1259.
- [26] Kohlberg, E., and J.-F. Mertens (1986), “On the strategic stability of equilibria,” *Econometrica* 54:1003-1008.
- [27] Krishna, V. and J. Morgan (2004), “Contracting for Information under Imperfect Commitment,” Mimeo, Berkeley.
- [28] Larson, K. and T. Sandholm (2001), “Costly Valuation Computation in Auctions.” In *Proceedings of the Theoretical Aspects of Reasoning about Knowledge (TARK)*, pp. 169-182.
- [29] Levin, J. (2003), “Relational Incentive Contracts,” *American Economic Review*, June.
- [30] Mailath, G. and A. Postlewaite (1990), “Asymmetric Information Bargaining Problems with Many Agents,” *Review of Economic Studies*, 57, 351-360.
- [31] Mezzetti, C. (2004), “Mechanism Design with Interdependent Valuations: Efficiency,” *Econometrica* 72(5): 1617.
- [32] Miller, D. (2004), “The dynamic cost of ex post incentive compatibility in repeated games of private information,” Mimeo, UCSD.
- [33] Myerson, R. (1986), “Multistage Games with Communication,” *Econometrica*, 54, 323-358.
- [34] Myerson, R. (1991), *Game Theory: Analysis of Conflict*, Harvard University Press.
- [35] Obara, I., and D. Rahman (2006), “Approximate Implementability with Ex Post Budget Balance,” working paper, UCLA Department of Economics.

- [36] Pollard, D. (2002), *A User's Guide to Measure-Theoretic Probability*, Cambridge University Press.
- [37] Proschan, M.A., and B. Presnell (1998), "Expect the Unexpected from Conditional Expectation," *The American Statistician*, 52(3), 248-252.
- [38] Puterman, M. (1994), *Markov Decision Processes*. Wiley: Hoboken, NJ.
- [39] Rayo, L. (2003), "Relational Team Incentives and Ownership," Mimeo, Chicago GSB.
- [40] Rogerson, W. (1992), "Contractual Solutions to the Hold-Up Problem," *Review of Economic Studies* 59, 777-793.
- [41] Smorodinsky, R., and M. Tennenholtz (2003), "Overcoming Free Riding in Multi-Party Computations-The Anonymous Case," Mimeo, Technion.
- [42] Stokey, N. and R. Lucas (1989), *Recursive Methods in Economic Dynamics*.
- [43] Wang, Cheng (1995), "Dynamic Insurance with Private Information and Balanced Budgets," *Review of Economic Studies* ; 62.4, October, pp. 577-95.

8. Appendix

Lemma 2. A functional $F : \prod_{t=1}^{\infty} Z_t \rightarrow \mathbb{R}$ is LCI with some $\beta \in (0; 1)$ if and only if $F(z) = \sum_{t=0}^{\infty} \beta^t f_t(z^t)$ for some uniformly bounded sequence of functions $\{f_t : Z^t \rightarrow \mathbb{R}\}_{t=0}^{\infty}$.

Proof of Lemma 2: "If": Letting $M = \sup_{i;t;z^t \in Z^t} |f_i(z^t)| < \infty$, we can write

$$\begin{aligned} |F(z) - F(y)| &\leq \sum_{t=\inf\{i:z_i \neq y_i\}}^{\infty} \beta^t |f_t(z^t) - f_t(y^t)| \leq \sum_{t=\inf\{i:z_i \neq y_i\}}^{\infty} \beta^t \cdot 2M \\ &= \frac{2M}{1-\beta} \beta^{\inf\{i:z_i \neq y_i\}} = \frac{2M}{1-\beta} \beta^{\inf\{i:z_i \neq y_i\}}(z; y) \end{aligned}$$

"Only if": fix $z \in Z$ and take $f_0 = F(z)$ and $f_t(z^t) = \beta^{-t} [F(z^t; z^{-t}) - F(z^{t-1}; z^{-(t-1)})]$ for $t \geq 1$. Then we have $|f_t(z^t)| \leq \beta^{-t} C$. $(z^t; z^{-t}); (z^{t-1}; z^{-(t-1)}) \leq C$ for $t \geq 1$. Furthermore,

$$\left| \sum_{t=0}^T \beta^t f_t(z^t) - F(z) \right| = |F(z^T; z^{-T}) - F(z)| \leq C \beta^T \rightarrow 0 \text{ as } T \rightarrow \infty. \blacksquare$$

Proof of Lemma 1: It suffices to establish a bound for any given pair $(i, j) \in \bar{\Theta}$, since the state space $\bar{\Theta}$ is finite. For the "transient" states $(i, j) \in \bar{\Theta} \setminus S$, the expression is bounded since by definition $\sum_{t=0}^{\infty} \beta^t p_{ij}^{(t)} < \infty$ and $\beta^t = 0$. For the "ergodic" states $(i, j) \in S$, note that all such states

have the same period $n \geq 1$, and let $G \subset S$ denote the cyclically moving subset containing i . Let $t_0 < n$ be the first $t \geq 0$ s.t. $p_{ij}^{(t)} > 0$. Then $p_{ij}^{(t)} > 0$ only for $t = t_0 + rd$, for $r = 0; 1; \dots$. Furthermore, the n -step transition process with the transition matrix $p_{ij}^{(n)}$, behaves like an irreducible aperiodic chain on G , with the invariant probabilities described by d_{ij} on $i \in G$. Applying Lucas-Stokey (1989, Theorem 11.4), the n -step process converges geometrically to its invariant distribution, i.e., there exists $\alpha < 1$ such that $\left| p_{ij}^{(t_0+rd)} - d_{ij} \right| \leq \alpha^r$. Therefore,

$$\begin{aligned} \left| \sum_{t=0}^{\infty} \alpha^t \left(p_{ij}^{(t)} - d_{ij} \right) \right| &\leq \sum_{t=0}^{\infty} \alpha^{t_0+rd} \left| p_{ij}^{(t_0+rd)} - d_{ij} \right| + \alpha^{t_0} \left| n \sum_{t=0}^{\infty} \alpha^{t_0+rd} - \sum_{t=0}^{\infty} \alpha^t \right| \\ &\leq \sum_{r=0}^{\infty} \alpha^{r} + \sum_{r=0}^{\infty} \alpha^{rd} \left| d_{ij} - \sum_{t=0}^{n-1} \alpha^t \right| \\ &\leq \frac{1}{1-\alpha} + \frac{1}{1-\alpha^n} n (1-\alpha^{n-1}) \leq \frac{1}{1-\alpha} + n: \blacksquare \end{aligned}$$