Revenue Equivalence, Profit Maximization, and Transparency in Dynamic Mechanisms*

Ilya Segal Juuso Toikka

December 5, 2007

Abstract

We consider the problem of designing incentive-compatible mechanisms in a general dynamic environment in which agents receive serially correlated private information over time and decisions may be made over time. The private information is assumed independent across agents, and its distribution is allowed to depend on past allocations. We derive a dynamic revenue equivalence result showing that any two mechanisms that implement the same allocation rule must yield the same expected payoffs to the agents and hence the same expected revenue regardless of the transfer scheme and of the information disclosed by the mechanism to the agents. We then use the result as a tool for designing profit maximizing mechanisms. As an example of the applications we analyze the problem of designing a profit-maximizing sequence of auctions when the bidders' types follow a linear AR(1)-process.

PRELIMINARY AND INCOMPLETE.

1 Introduction

We consider the problem of designing incentive-compatible mechanisms in a dynamic environment in which agents receive private information over time and decisions may be made over time. The model allows for serial correlation of agents' private information as well as its dependence on past allocations. For example, it includes such problems as the allocation of resources to randomly arriving agents whose valuations follow a stochastic process, or the design of multiperiod procurement auctions for bidders whose cost parameters evolve over time and may exhibit learning-by-doing effects. We assume that private information is independently distributed across agents. For much of the paper we concentrate on settings where the agents' utilities are quasilinear.

Unlike many of the recent papers on dynamic mechanism design we do not restrict attention to surplus-maximizing mechanisms.[1] Instead, we derive a dynamic payoff formula that gives the derivative of each agent's expected payoff conditional on his first

*Department of Economics, Stanford University, Stanford, CA 94305, ilya.segal@stanford.edu, juuso.toikka@stanford.edu

[1]See Athey and Segal (2007) for a discussion of the literature on efficient dynamic mechanisms.

period private information for any implementable choice rule. This expression summarizes local incentive compatibility constraints in a manner analogous to the "Mirrlees's trick" (Mirrlees, 1971) familiar from static mechanism design. For the special case of quasilinear environments the formula implies a dynamic "revenue equivalence" result: any two Bayesian incentive-compatible mechanisms that implement the same allocation rule must yield the same expected payoff to each agent and therefore yield the same expected revenue regardless of the transfer rule or the information revealed by the mechanism. Just like in the static setting, this result requires assuming that agents' private information in each period lies in a connected interval, and that utility depends smoothly on the private information. In addition, however, it requires that the agents' private information evolve in a "smooth" way.

The dynamic payoff formula allows us to express the expected profits in the mechanism as the expected "virtual surplus," appropriately defined for the dynamic setting. This expression is derived using only the agents' local incentive constraints, and not all allocation rules that maximize it can be implemented in an incentive-compatible mechanism. However, we show that when the agent's types follow a Markov process with transitions that are increasing in the first order stochastic dominance sense, under the usual single-crossing assumption on the agents' utilities, any allocation rule in which each agent's consumption in each period is nondecreasing in the history of his own reports is implementable in an incentive-compatible mechanism. Furthermore, any such allocation rule can be implemented in a mechanism that is "transparent," i.e., in which all agents' reports are made public. (In general, implementing a given allocation rule may require using a mechanism in which agents' reports are kept confidential.)

In the special case where the agents' types follow an AR(1) process, with the standard monotone hazard rate assumption on the agents' initial type distribution and the standard third-derivative assumption on their utility functions, virtual-surplus maximization indeed yields an allocation rule that is monotone in the desired sense, and therefore implementable in an incentive-compatible transparent mechanism. This mechanism is therefore profit-maximizing among all incentive-compatible mechanisms. The optimal mechanism exhibits some interesting properties: for example, an agent's consumption in a given period depends only on his initial report and his current report, but not on any intermediate reports. This can be interpreted as a scheme where the agents make up-front payments that reduce their future distortions. The model can also be viewed as an approximation of a continuous-time model in which agents' types follow a continuous-time stochastic process.

Many of the existing papers on optimal mechanisms in dynamic environments consider special cases of the general model for which we establish revenue equivalence. Courty and Li (2000) consider a two-period model of sequential screening of a single agent who in the first period receives a private signal of his private valuation for consumption in the second period. We generalize their model in three respects. We allow for multiple periods, multiple agents, and for more general specifications for the payoffs and the stochastic structure. Nevertheless, some of the ideas or results of the present paper, such as sufficient conditions for incentive compatibility, appear in a specialized form in Courty and Li (2000). Baron and Besanko (1984) is an earlier contribution that contains special cases of our results, even though they too consider a single agent and state most of their results for two periods. Theirs is the insight about the informativeness of the first period type about the period t type determining the period t distortions.

Battaglini (2005) considers a model with one agent and two types, and derives an optimal selling strategy for a monopolist facing a Markovian consumer. Our results for a model with a continuous type show that many of his predictions seem special to the setting with only two types. We discuss some of the differences between our results and his in section 5.1.

Gershkov and Moldovanu (2007) consider both efficient and profit maximizing ways to allocate a fixed set of objects to randomly arriving buyers. While the model does have multiple agents, they assume that each agent only lives instantaneously. Hence the problem of each agent is a static one. They show a payoff equivalence result for the environment, but the result is essentially the static payoff equivalence result applied separately to each short lived agent. In contrast, we allow the agents to be long lived.

Dynamic mechanism design is inherently related to the literature on multidimensional screening—see Rochet and Stole (2003) for a discussion. In fact, there is a sense in which the problem of dimensionality is particularly severe in dynamic mechanism design. With serial correlation the agent's current type corresponds to a distribution over future types, and with a continuum of possible future types such distributions are infinite dimensional objects. However, it is possible to impose structure on the payoff functions and distributions that ensure that a single crossing property is satisfied in each period. This makes it possible to have relatively simple characterizations of incentive compatibility. In part this is due to the fact that, in a dynamic environment where the private information is observed sequentially, there are less deviations available to the agent compared to the situation where the same information is available to him at once as would be the case in a static multi-dimensional environment.

We touch here upon the issue of transparency. Pancs (2007) studies its role in trading mechanisms in environments where agents take nonenforceable actions such as investment or information acquisition.

We set up the model in the next section, and derive the payoff equivalence result in section 3. We establish sufficient conditions for incentive compatibility in section 4.

We then consider the problem of designing optimal mechanisms in section 5. Section 6 is devoted to an example.

2 The Model

We derive the results about payoff equivalence and incentive compatibility in a single-agent model where, in addition to the agent's privately known type, there is a stochastically evolving state of the world. The state is unobservable to the agent, but his payoff may depend on it, and mechanisms can make decisions contingent on it. Our leading interpretation for the state is to think of it as private information held by other (unmodeled) agents (i.e., their types). With this interpretation the model corresponds to the problem faced by an individual agent in a multi-agent environment when other agents are playing their equilibrium strategies. More generally, the state can capture any other causes of the agent's uncertainty about the environment or account for deliberate randomizations by the mechanism.

2.1 The Environment

For expositional ease we consider a Markov environment with one agent and finitely many periods, indexed by $t = 1, 2, \ldots, T$. We remark below how the results extend to non-Markov and infinite horizon models.

In each period t there is a (contractible) *decision* $y_t \in Y_t = Y_{ut} \times Y_{ot}$. The decision $y_t = (y_{ut}, y_{ot})$ consists of two parts. The second component y_{ot} is observable to the agent, whereas the first component y_{ut} is not. The set of all period t decision histories is denoted $Y^t = \prod_{\tau=0}^{t} Y_\tau$. The sets Y_u^t and Y_o^t are defined analogously. An element of $Y = Y^T$ is denoted $y = (y_u, y_o)$. Each Y_{ut} and Y_{ot} is assumed to be a standard Borel (also known as "nice measurable") space as to guarantee the existence of regular conditional probability distributions (see, e.g., Durrett, 2004). All product sets are endowed with the product sigma-algebra.

The observability of the decision should be thought of as a technological constraint. Below we allow mechanisms to reveal y_{ut}, but not to conceal y_{ot}. Hence y_{ot} is a lower bound on the observability of the decision y_t. E.g., if the decision includes allocating private goods, then y_{ot} would include the agent's own consumption but not consumption of the other (unmodeled) agents, as this could be concealed from him. On the other hand, if the decision is about a public good, y_{ot} may include everyone's consumption.

Before the period t decision, the agent privately observes his *current type* $\theta_t \in \Theta_t \subset \mathbb{R}$. The set of possible type histories at period t is denoted $\Theta^t = \prod_{\tau=1}^{t} \Theta_\tau$. An element θ of $\Theta = \Theta^T$ is referred to as the agent's *type*.

In addition to the private type, there is an auxiliary state of the world, which is unobservable to the agent.[2] We let ω_t denote the *current state*, which is an element of a measurable space Ω_t. We let $\Omega^t = \prod_{\tau=1}^{t} \Omega_\tau$ denote the set of possible period t state histories, which is endowed with the product sigma-algebra. The *state* ω is an element of $\Omega = \Omega^T$.

The distribution of the next period's type θ_{t+1} and state ω_{t+1} may depend on the current type and state as well as on the entire decision history. In particular, we assume that the joint distribution of $(\theta_{t+1}, \omega_{t+1})$ on $\Theta_{t+1} \times \Omega_{t+1}$ is governed by a history dependent probability measure $\Lambda_{t+1}\left(\cdot|\theta_t, \omega_t, y^t\right)$ such that $\Lambda_{t+1}\left(A|\cdot\right) : \Theta_t \times \Omega_t \times Y^t \to \mathbb{R}$ is measurable for all measurable $A \subset \Theta_{t+1} \times \Omega_{t+1}$. We assume that the type and the state are independent given past decisions in the following sense. For all t, and all histories $\left(\theta_t, \omega_t, y^t\right)$, the probability measure takes the form

$$\Lambda_{t+1}\left(\cdot|\theta_t, \omega_t, y^t\right) = F_{t+1}\left(\cdot|\theta_t, y_o^t\right) \times G_{t+1}\left(\cdot|\omega_t, y^t\right),$$

where $F_{t+1}\left(\cdot|\theta_t, y_o^t\right)$ and $G_{t+1}\left(\cdot|\omega_t, y^t\right)$ are probability measures on Θ_{t+1} and Ω_{t+1}, respectively. This implies that the current type and current state evolve according to a Markov process, which is controlled by the decision history. Note that the distribution of the type depends only on objects observable to the agent. We abuse notation by using $F_{t+1}(\cdot|\theta_t, y_o^t)$ to denote also the cumulative distribution function (cdf) corresponding to the measure $F_{t+1}(\cdot|\theta_t, y_o^t)$.

We note for future reference that the history dependent measures $\Lambda_{t+1}\left(\cdot|\theta_t, \omega_t, y^t\right)$ induce a unique probability distribution on $\Theta \times \Omega$ when accompanied with any measurable function that maps type and state histories to current decisions. More precisely, suppose that $h : \Theta \times \Omega \to Y$ is a measurable function that is adapted such that $h_t(\theta, \omega)$ depends only on (θ^t, ω^t). Then h induces probability kernels $\Lambda_{t+1}\left(\cdot|\theta_t, \omega_t, h^t(\theta^t, \omega^t)\right)$ that uniquely define a probability measure μ_h on $\Theta \times \Omega$.

The agent's payoff from a sequence of decisions y given type θ and state ω is

$$\sum_{t=1}^{T} \delta^t u_t\left(y^t, \theta_t, \omega_t\right),$$

where $\delta \in (0,1]$ is the discount factor and the functions $u_t : Y^t \times \Theta_t \times \Omega_t \to \mathbb{R}$ are assumed measurable. In general the agent does not observe his realized utility, but only the observable decisions.[3] We note that the current payoff $u_t\left(y^t, \theta_t, \omega_t\right)$ may depend on

[2] A complete description of the state of the world at any period t would include the agent's type history θ^t as well as the history of the state ω^t.

[3] Alternatively, if the agent does observe his utility, then it is independent of the unobservable decisions y_u^t and the state ω_t.

the entire history of decisions. This allows for effects such as learning-by-doing or habit-formation. It also includes situations where the agent wants to consume only once but the time of consumption is a choice variable. Yet another example is an intertemporal capacity constraint that could be present, e.g., in repeated procurement.

For much of the paper we restrict attention to quasilinear environments defined as follows.

Definition 1. The environment is *quasilinear* if for all periods t all of the following hold:

1. The decision is of the form $y_t = (y_{ut}, y_{ot}) = (x_{ut}, (x_{ot}, p_t))$, where $x_t = (x_{ut}, x_{ot}) \in X_{ut} \times X_{ot} = X_t$ is an *allocation*, and $p_t \in \mathbb{R}$ is a *transfer*. The set of possible decisions is $Y_t = X_t \times \mathbb{R}$. Only x_{ot} and p_t are observable to the agent.[4]

2. The distributions are independent of the transfers: $F_{t+1}\left(\cdot|\theta_t, y_o^t\right) = F_{t+1}\left(\cdot|\theta_t, x_o^t\right)$ and $G_{t+1}\left(\cdot|\omega_t, y^t\right) = G_{t+1}\left(\cdot|\omega_t, x^t\right)$.

3. The periodic utility function is quasilinear: $u_t(y^t, \theta_t, \omega_t) = v_t(x^t, \theta_t, \omega_t) + p_t$ for some measurable $v_t : X^t \times \Theta_t \times \Omega_t \to \mathbb{R}$.

2.2 Mechanisms

By the revelation principle for dynamic games of Myerson (1986) it is without loss to restrict attention to a subclass of all possible mechanisms. However, before describing the class of canonical mechanisms, we first present general definitions of mechanisms, strategies, and implementation in order to illustrate the scope of the results, and to facilitate discussion of information disclosure policies, which are a feature of dynamic mechanism not present in static models.

A mechanism in the above environment assigns a set of possible messages to the agent in each period. The agent sends a message from this set to the mechanism and an outcome function maps the history of messages into a current decision. We allow the mechanism to make decisions contingent on the state so that the outcome function generally depends also on the history of the state.

Given that the environment is dynamic and the mechanism designer has superior information about the state and the part of the decision unobservable to the agent, the mechanism also includes an information disclosure policy. Even if the mechanism does not directly reveal information about the state, the agent can in general infer something about it based on the knowledge of the mechanism and the observable decisions. An

[4] The assumption that the agent can observe his current transfer is without loss of generality since in a quasilinear model all transfers can be done in the last period.

information disclosure policy can reveal more information about the state either directly or by revealing information about the unobservable decisions. Since the agent's messages affect the decisions, they in part contribute to what the agent learns about the state. We assume throughout that the designer can commit to the mechanism.

Formally, the period t message m_t is an element of a measurable space M_t. We denote the set of period t message histories by $M^t = \prod_{\tau=1}^{t} M_\tau$ with $M = M^T$ referred to as the *message space*. Each M^t is endowed with the product sigma-algebra.

An *outcome function* is a measurable mapping

$$g : M \times \Omega \to Y,$$

which determines the decisions as a function of the agent's messages and the state. In order to respect observability assumptions we assume that the outcome function is adapted[5]: $g_t(m, \omega)$ depends only on (m^t, ω^t). Where convenient we write $g_t(m^t, \omega^t) = g_t(m, \omega)$. We use the notation g_o and g_u to denote the component mappings of g corresponding to the observable and unobservable parts of the decisions. Assuming that the outcome function is deterministic is without loss as randomizations can be introduced through its dependence on ω.

We model information disclosure as follows. Assume first that, before the period t decision, the agent has simply seen the past observable decision y_{ot-1}. An outcome function g, a message history m^{t-1}, and a state history ω^{t-1} then induce observations $g_{ot-1}(m^{t-1}, \omega^{t-1})$ in the natural way. Notice that fixing the outcome function, the observations are a function only of message and state histories (m^{t-1}, ω^{t-1}). A general information disclosure policy generalizes this by assuming that in each period t, before reporting, the agent observes the value $I_t(m^{t-1}, \omega^t)$ of some measurable function I_t.[6] This allows the mechanism to communicate more information about the state and decision histories than is included in $g_{ot-1}(m^{t-1}, \omega^{t-1})$. Furthermore, $I_t(m^{t-1}, \omega^t)$ is a function of the current state as this information is available to the mechanism designer. For example, the mechanism could simply reveal the current state to the agent.[7] Given perfect recall, the agent's information about the state in period t before sending his message is captured by the information history $I^t(m^{t-1}, \omega^t)$. The collection of mappings $I := (I_t)_{t=1}^T$ is referred to as the *information (disclosure) policy.*

We are now ready to state the formal definition.

Definition 2. A *mechanism* is a triple $\langle M, g, I \rangle$, where M is a message space, g is an

[5] Athey and Segal (2007) use the term *observationally measurable.*

[6] The function I_t is a mapping from $M^{t-1} \times \Omega^t$ to some set $range(I_t)$. The range is of little consequence to the present analysis and hence left implicit.

[7] In a multi-agent environment, where the state is the profile of other agents' types, learning the current state before sending the message corresponds to a periodic ex post equilibrium.

outcome function, and I is an information policy.

Consider a mechanism $\langle M, g, I\rangle$. In period t, before sending his message, the agent's information consists of his message history m^{t-1}, type history θ^t, and information history $I^t(m^{t-1}, \omega^t)$. (Note that I^t contains the realized observed decisions y_o^{t-1}.) A strategy for the agent in the mechanism $\langle M, g, I\rangle$ is thus a mapping $\sigma : M \times \Theta \times range(I) \to M$.[8] We assume that σ is adapted in the natural sense: for all t and all (m, θ, I), $\sigma_t(m, \theta, I) = \sigma_t(m^{t-1}, \theta^t, I^t)$.

A strategy σ induces a *strategic plan* $\bar{\sigma} : \Theta \times \Omega \to M$ as follows. Let $\bar{\sigma}_1(\theta^1, \omega^1) = \sigma_1(\theta^1, I^1(\emptyset, \omega^1))$. For $t = 2, \ldots, T$, let

$$\bar{\sigma}_t(\theta^t, \omega^t) = \sigma_t(\bar{\sigma}^{t-1}(\theta^{t-1}, \omega^{t-1}), \theta^t, I^t(\bar{\sigma}^{t-1}(\theta^{t-1}, \omega^{t-1}), \omega^t).$$

Then $\bar{\sigma}(\theta, \omega)$ gives the agent's messages given type θ and state ω. Note that the function $\bar{\sigma}$ is measurable and adapted by construction. Hence the composite mapping $(\theta, \omega) \longmapsto g(\bar{\sigma}(\theta, \omega), \omega)$ is measurable and adapted. Thus, by the above observation, it induces a unique probability distribution on $\Theta \times \Omega$. We denote this measure by μ_σ leaving its dependence on the mechanism implicit. Similarly, given a first period type θ_1, there exists a unique probability measure $\mu_\sigma|\theta_1$ on $\Theta \times \Omega$ such that the first period type equals θ_1 with probability 1.

Definition 3. The strategy σ is *optimal* in the mechanism $\langle M, g, I\rangle$ if for all $\theta_1 \in \Theta_1$, and all strategies ρ,

$$\mathbb{E}^{\mu_\sigma|\theta_1}\Big[\sum_{t=1}^{T} \delta^t u_t\left(g^t\left(\bar{\sigma}^t(\theta, \omega), \omega^t\right), \theta_t, \omega_t\right)\Big] \geq \mathbb{E}^{\mu_\pi|\theta_1}\Big[\sum_{t=1}^{T} \delta^t u_t\left(g^t\left(\bar{\rho}^t(\theta, \omega), \omega^t\right), \theta_t, \omega_t\right)\Big].$$

That is, the optimality of a strategy is required at the stage where the agent knows his first period type, but is yet to receive any information from the mechanism. With slight abuse of terminology, we refer to the inequalities embedded in this definition as the ex ante incentive compatibility constraints.

A *(social) choice rule* is an adapted measurable mapping $\varphi : \Theta \times \Omega \to Y$ that assigns a sequence of decisions to each type-state-pair. If the environment is quasilinear, then $Y = X \times \mathbb{R}^T$ and we write $\varphi = (\chi, \psi)$. The component mapping χ is an *allocation rule* and ψ is a *transfer rule.*

Definition 4. The mechanism $\langle M, g, I\rangle$ *implements the choice rule* φ if there exists an

[8] We require a strategy specify messages after own deviations in anticipation of the application to multi-agent environments.

optimal strategy σ such that for all (θ, ω),

$$\varphi(\theta, \omega) = g\left(\bar{\sigma}(\theta, \omega), \omega\right).$$

The choice rule φ is *implementable* if there exists a mechanism $\langle M, g, I\rangle$ that implements it. In the quasilinear environment, the mechanism $\langle M, g, I\rangle$ *implements the allocation rule* χ if $\langle M, g, I\rangle$ implements the choice rule (χ, ψ) for some transfer rule ψ. The allocation rule χ is *implementable* if there exists a mechanism $\langle M, g, I\rangle$ and a transfer rule ψ such that $\langle M, g, I\rangle$ implements (χ, ψ).

2.3 Revelation Principle

Myerson (1986) argues that in a general dynamic Bayesian game, any outcome that can be achieved by adding mediated communication between the players can be achieved with a mechanism where in each period (1) each agent confidentially reports any new private information to the mechanism, (2) the mechanism confidentially recommends an action to each player, and (3) each agent is honest and obedient. Adapted to our setting this dynamic revelation principle implies that it is without loss to consider mechanisms where in each period t:

(1) the agent privately[9] reports a current type to the mechanism,

(2) no information about the state or the unobservable decisions is revealed to the agent (beyond what can be inferred from the observable decisions), and

(3) the agent is honest.

We refer to mechanisms satisfying condition (1) as direct mechanisms, and to those satisfying conditions (1) and (2) as canonical mechanisms. Formally, we have the following definitions.

Definition 5. The mechanism $\langle M, g, I\rangle$ is a *direct mechanism for the choice rule* φ if $M = \Theta$ and $g = \varphi$.

We identify a direct mechanism $\langle \Theta, g, I\rangle$ with the tuple $\langle g, I\rangle$. Note that the outcome function g in a direct mechanism is a choice rule.

Definition 6. A direct mechanism $\langle g, I\rangle$ is *canonical* if $I = g_o$.

That is, the information policy in a canonical mechanism is minimal in that it only reveals the observable decision (i.e., $I_t(m^{t-1}, \omega^{t-1}) = g_{ot}(m^{t-1}, \omega^{t-1})$ for all t). When

[9] In the model with one agent there is no difference between private or public reporting, but we add the qualifier in anticipation of the application to multi-agent environments.

this causes no confusion, we identify a canonical mechanism with the outcome function g.

Definition 7. The choice rule φ is *truthfully implementable* if honesty is optimal in the canonical mechanism $\langle \varphi, \varphi_o \rangle$. That is, if there exists an optimal strategy σ such that for all (θ, ω), the strategic plan $\bar{\sigma}$ induced by σ satisfies $\bar{\sigma}(\theta, \omega) = \theta$.

Dynamic Revelation Principle (Myerson, 1986). *A choice rule φ is implementable if and only if it is truthfully implementable.*

3 The Dynamic Payoff Formula and Revenue Equivalence

We now set out to derive an explicit expression for the derivative of the agent's expected payoff from implementing a choice rule φ. The resulting dynamic payoff formula summarizes local incentive compatibility constraints in a manner analogous to the "Mirrless trick" (Mirrlees, 1971) familiar from static mechanism design. If the environment is quasilinear, the formula implies that the agent's payoff is pinned down by the allocation rule and hence is independent of the transfer rule (up to a constant). The formula also proves useful in designing optimal dynamic mechanisms.

In deriving the dynamic payoff formula we use the following key assumptions.

Assumption 1. *For all t, $\Theta_t = (\underline{\theta}_t, \overline{\theta}_t) \subset \mathbb{R}$ for some $-\infty \leq \underline{\theta}_t \leq \overline{\theta}_t \leq +\infty$.*

Assumption 2. *For all t, and all $(y^t, \theta_t, \omega_t)$, the partial derivative $\frac{\partial u_t(y^t, \theta_t, \omega_t)}{\partial \theta_t}$ exists and is bounded uniformly in $(y^t, \theta_t, \omega_t)$.*

Assumption 3. *For all t, and all (θ_t, y_o^t), the Lebesgue measure on Θ_{t+1} is absolutely continuous with respect to $F_{t+1}(\cdot|\theta_t, y_o^t)$.*

Assumption 4. *For all t, and all (θ_t, y_o^t), $\int \theta_{t+1} dF_{t+1}(\theta_{t+1}|\theta_t, y_o^t) < +\infty$.*

Assumption 5. *For all t, and all $(\theta_{t+1}, \theta_t, y_o^t)$, the partial derivative $\frac{\partial F_{t+1}(\theta_{t+1}|\theta_t, y_o^t)}{\partial \theta_t}$ exists. Furthermore, there exists an integrable function $b_{t+1} : \Theta_{t+1} \to \mathbb{R}$ such that for all $(\theta_{t+1}, \theta_t, y_o^t)$,*

$$\left| \frac{\partial F_{t+1}(\theta_{t+1}|\theta_t, y_o^t)}{\partial \theta_t} \right| \leq b_{t+1}(\theta_{t+1}).$$

Assumptions 1 and 2 are familiar from static settings (see, e.g., Milgrom and Segal, 2002). Note, however, that we do not require that the set of types be bounded. Assumptions 3 and 4 are also typically made in static models. Assumption 3 is a full support assumption which implies that the cdf $F_{t+1}(\cdot|\theta_t, y_o^t)$ is strictly increasing over any nondegenerate subinterval of Θ_{t+1}, but it can have jumps. Assumption 4 about the existence of the expectation is trivial if Θ_{t+1} is bounded.

Assumption 5 requires that the distribution of the current type depend sufficiently smoothly on the past type. The motivation for it is essentially the same as for requiring that, even in static settings, utilities depend smoothly on types. However, since this assumption does not have an immediate counterpart in the static model, it is perhaps instructive to consider what it restrictions it imposes on the stochastic process for θ_t. In particular, it implies that derivative of the expected future type with respect to the current, $\frac{\partial}{\partial\theta_t}\mathbb{E}[\theta_{t+1}|\theta_t, y_o^t]$, exists and is bounded uniformly in (θ_t, y_o^t)—see Lemma B.1 in the Appendix B.

We are now ready to state our main result.

Theorem 1. *Assume that the environment satisfies Assumptions 1–5. Let φ be an implementable choice rule. Then there exists a Lipschitz-continuous function $U_1 : \Theta_1 \to \mathbb{R}$ such that in any mechanism $\langle M, g, I\rangle$ that implements φ, the agent's expected payoff given first period type θ_1 equals $U_1(\theta_1)$. Furthermore, for a.e. θ_1,*

$$U_1'(\theta_1) = \sum_{t=1}^{T} \delta^t \int_{\Omega^t \times \Theta^t} \frac{\partial u_t(\varphi^t(\theta^t, \omega^t), \theta_t, \omega_t)}{\partial \theta_t} \prod_{\tau=2}^{t} \left(-\frac{\partial F_\tau(\theta_\tau|\theta_{\tau-1}, \varphi_o^{\tau-1}(\theta^{\tau-1}, \omega^{\tau-1}))}{\partial \theta_{\tau-1}} \right)$$
$$\times\, dG_t(\omega_t|\omega_{t-1}, \varphi^{t-1}(\theta^{t-1}, \omega^{t-1})) \cdots dG_1(\omega_1) d\theta_t \cdots d\theta_2 d\gamma_{\theta_1},$$

where $\prod_{\tau=2}^{1}\left(-\frac{\partial F_\tau(\theta_\tau|\theta_{\tau-1}, \varphi_o^{\tau-1}(\theta^{\tau-1}, \omega^{\tau-1}))}{\partial\theta_{\tau-1}}\right) \equiv 1$ and γ_{θ_1} is the Dirac measure at θ_1.

The proof of this result can be found in the Appendix A. As the dynamic payoff formula is relatively complicated, it may be useful to note that if the cdf's F_{t+1} have strictly positive densities f_{t+1}, then the formula simplifies to

$$U_1'(\theta_1) = \mathbb{E}^{\mu_\varphi|\theta_1} \left[\sum_{t=1}^{T} \delta^t \frac{\partial u_t(\varphi^t(\theta^t, \omega^t), \theta_t, \omega_t)}{\partial \theta_t} \prod_{\tau=2}^{t} \left(-\frac{\partial F_\tau(\theta_\tau|\theta_{\tau-1}, \varphi_o^{\tau-1}(\theta^{\tau-1}, \omega^{\tau-1}))/\partial\theta_{\tau-1}}{f_\tau(\theta_\tau|\theta_{\tau-1}, \varphi_o^{\tau-1}(\theta^{\tau-1}, \omega^{\tau-1}))} \right) \right],$$

where the expectation is taken over $\Theta \times \Omega$ with respect to the measure $\mu_\varphi|\theta_1$ induced by the choice rule φ and the first period type θ_1.

If the environment is quasilinear, then the partial derivative of the current payoff with respect to the current type (holding the messages fixed) is independent of the current transfer. Furthermore, the distributions F_t and G_t are then independent of the transfers by assumption. By inspection of the dynamic payoff formula of Theorem 1 this implies that U_1' is determined by the allocation rule alone. Thus we have the following corollary.

Corollary 1 (Payoff Equivalence). *Assume that the environment is quasilinear and satisfies Assumptions 1–4. Let χ be an implementable allocation rule. If the mechanisms $\langle M, g, I\rangle$ and $\left\langle \tilde{M}, \tilde{g}, \tilde{I}\right\rangle$ implement χ, then there exists $a \in \mathbb{R}$ such that the agent's* *expected payoffs in the two mechanisms satisfy* $U_1(\theta_1) = \tilde{U}_1(\theta_1) + a$ *for all* θ_1*. That is, the agent's expected payoff is independent of the transfer scheme and of the information policy up to a constant.*

The corollary is a dynamic version of the well-known static payoff equivalence theorems. The fact that the agent's payoff here is independent of the information policy is reminiscent of the fact that, in static settings, Bayesian and dominant strategy implementation of a given allocation rule give rise to the same interim payoffs to the agents. It should be noted that information policy is irrelevant for the payoff only as long as the policy is such that the allocation rule remains implementable. Just like it is possible to implement a larger set of choice rules with Bayesian rather than dominant strategy implementation in static settings, the revelation principle implies that it is in general possible to implement a larger set of dynamic choice rules by using mechanisms that reveal less information.

Corollary 1 generalizes the revenue equivalence theorem of Myerson (1981) to dynamic settings. This result is obtained by interpreting the state ω to consist of the other bidders' types. Then applying the payoff equivalence result to each individual bidder, we see that the bidders' expected payoff is the same in any two auction formats that induce the same allocation and where the bidder with the lowest possible first period type has the same payoff. Since the allocation is the same in the two mechanisms, also the aggregate surpluses must coincide. Revenue equivalence then follows by noting that the seller's revenue consists of the aggregate surplus less the bidders' expected payoffs.

Another corollary concerns uniqueness of the incentive payments needed to implement efficient allocations. By an argument similar to the one behind revenue equivalence, we can establish that the Team Mechanism of Athey and Segal (2007) is essentially the unique mechanism that implements efficient allocations in Markov environments. Indeed, other variants of efficient dynamic mechanisms such as the Dynamic Marginal Contribution Mechanism of Bergemann and Välimäki (2007) differ from the Team Mechanism only in payments that are not contingent on an agent's own report. Their relationship is analogous to that among different versions of static VCG mechanisms.

Remark 1. *Theorem 1 and Corollary 1 extend to an infinite horizon model with discounting under additional smoothness assumptions on the stochastic process. In particular, suppose that the periodic utility functions* u_t *are bounded, and that each distribution* $F_t(\cdot|\theta_t, y_o^t)$ *has a density* $f_t(\cdot|\theta_t, y_o^t)$ *that satisfies a condition analogous to Assumption 5. Suppose further, that*

$$\delta \int_{\Theta_{t+1}} \left| \frac{\partial F_{t+1}(\theta_{t+1}|\theta_t, y_o^t)}{\partial \theta_t} \right| d\theta_{t+1}$$

is bounded away from 1. Then it is possible to take $T = +\infty$ in the dynamic payoff formula of Theorem 1. Corollary 1 follows just like in the finite horizon case.

Remark 2. *The results also extend to a finite horizon non-Markov environment where the distributions satisfy assumptions analogous to those required in the Markov model. Although the analog of Theorem 1 is substantially more complicated, Corollary 1 can nevertheless be recovered. The results for the non-Markov model also allow us to extend the payoff equivalence result of Corollary 1 to Markov environments where private information in each given period can be multi-dimensional. This is done by transforming the multi-dimensional Markov model into a one-dimensional non-Markov model where instead of observing his type θ_t at once, the agent observes it component by component in some prespecified order. This only reduces the deviations available to the agent; if truthtelling is optimal when θ_t is observed at once, it remains optimal when observations are sequential. This is enough for payoff equivalence.*

4 Sufficient Conditions for Implementability

The dynamic payoff formula is derived using only local incentive compatibility (IC) constraints. That is, we are considering only "small" deviations from truthtelling. However, just like in the static case, an allocation rule that satisfies the local constraints need not be implementable as it may fail to satisfy global IC constraints. I.e., an agent may find telling a "'big lie" to be a profitable deviation.

It is well known that in the static case, if the agent's payoff function satisfies a single-crossing condition, then monotonicity of the allocation rule closes the gap between local and global IC constraints. It is then possible to obtain a characterization of all implementable choice rules that facilitates finding optimal ones. In contrast, it is difficult to characterize implementability in a general dynamic environment, because the agents' private information is effectively multi-dimensional. Even if the current type is a real number, it generally corresponds to a distribution over future types. And with a continuum of possible types, distributions are infinite dimensional objects. Put another way, the literature on multi-dimensional screening has made it clear that it is meaningful to speak of the private information as having a particular dimensionality only if it is with respect to a single-crossing condition.

Given the "curse of dimensionality," an analytically tractable model has to place restrictions on the stochastic process. Assume for simplicity that the distribution of the current type θ_{t+1} is independent of the past decisions. Then a natural candidate is to assume that the distributions $F_{t+1}(\cdot|\theta_t)$ are ordered by θ_t in the sense of first order stochastic dominance (FOSD). Coupled with a single-crossing condition on the periodic payoffs, FOSD gives rise to a natural ordering of current types. Heuristically, single-crossing of the periodic payoffs implies that high values of θ_t are the "high" types from the perspective of the current period. FOSD implies that high values of θ_t also correspond to high values of θ_{t+1} (in a stochastic sense). But high values of θ_{t+1} are again "high" types from the period $t+1$ perspective and so on. This complementarity, or positive feedback loop, guarantees that high current types are the ones who should be receiving (in expectation) high present and future allocations.

For the linear model studied in section 6 it is actually possible to arrive at a characterization of all incentive compatible allocation rules. In general we have the following sufficient conditions that extend a result of Courty and Li (2000) to multiple periods and many agents. The conditions are sufficient for implementation not only in the canonical mechanism, but with full transparency as defined below.

Definition 8. The choice rule φ is *truthfully implementable with full transparency* if honesty is optimal in the direct mechanism $\langle \varphi, I^* \rangle$, where

$$I_t^*(m^{t-1}, \omega^t) = (\varphi_{t-1}(m^{t-1}, \omega^{t-1}), \omega_t).$$

That is, the information policy I^* is such that in each period the agent observes the entire past decision and the current state before sending his message.

For the statement of the result, we need to make precise what me mean by local IC constraints.

Definition 9. For all t and all (θ^t, ω^t), let $U_t(\theta^t, \omega^t)$ denote the agent's expected continuation utility from truthtelling given history (θ^t, ω^t) in period t in the fully transparent mechanism $\langle (\chi, \psi), I^* \rangle$. The choice rule (χ, ψ) *satisfies local IC constraints (ICFOC)* if for all t, all (θ^{t-1}, ω^t), and almost every θ_t,

$$\frac{\partial U_t(\theta^t, \omega^t)}{\partial \theta_t} = \frac{\partial v_t(\chi_t(\theta^t, \omega^t), \theta_t, \omega_t)}{\partial \theta_t} + \delta \int_{\Theta_{t+1} \times \Omega_{t+1}} \frac{\partial U_{t+1}(\theta^{t+1}, \omega^{t+1})}{\partial \theta_{t+1}} \left(-\frac{\partial F_{t+1}(\theta_{t+1}|\theta_t)}{\partial \theta_t} \right) d\theta_{t+1} dG_t(\omega_{t+1}|\omega_t).$$

We simplify by assuming through the rest of this subsection that the distributions of types and states are independent of the past decisions, and that periodic payoffs depend only on current allocations. Both of these assumptions can be relaxed by assuming sufficient complementarity at the cost of obfuscating the argument.

Assumption 6. *The distributions and current payoffs are independent of past decisions:* $F_{t+1}\left(\cdot|\theta_t, y_o^t\right) = F_{t+1}\left(\cdot|\theta_t\right)$, $G_{t+1}\left(\cdot|\omega_t, y^t\right) = G_{t+1}\left(\cdot|\omega_t\right)$, *and* $u_t(y^t, \theta_t, \omega_t) = u_t(y_t, \theta_t, \omega_t)$ *for all* $(y^t, \theta_t, \omega_t)$ *and all* t.

Lemma 1. *Assume that the environment is quasilinear and satisfies Assumptions 1–6. Assume further that $X_t \subset \mathbb{R}$ for all t. Let (χ, ψ) be a choice rule that satisfies ICFOC. Then (χ, ψ) is truthfully implementable with full transparency if all of the following conditions hold:*

1. *SCP: for all t, and all $(x_t, \theta_t, \omega_t)$, $\frac{\partial v_t(x_t,\theta_t,\omega_t)}{\partial \theta_t}$ is strictly increasing in x_t;*
2. *FOSD: for all t, and all $(\theta_{t+1}, \theta_t, x_o^t)$, $\frac{\partial F_{t+1}(\theta_{t+1}|\theta_t)}{\partial \theta_t} \leq 0$;*
3. *Strong-MON: $\chi_t(\theta^t, \omega^t)$ is nondecreasing in θ_τ for all τ, all ω^t, and all t.*

Proof. The continuation utilities $U_t(\theta^t, \omega^t)$ satisfy the recursion

$$U_t(\theta^t, \omega^t) = \max_{m_t \in \Theta_t} \Big\{ v_t(\chi_t(\theta^{t-1}, m_t, \omega^t), \theta_t, \omega_t) + \psi_t(\theta^{t-1}, m_t, \omega^t) \\ + \delta \int_{\Theta_{t+1} \times \Omega_{t+1}} U_{t+1}(\theta_{-t}^{t+1}, m_t, \omega^{t+1}) d\Lambda_{t+1}(\theta_{t+1}, \omega_{t+1} | \theta_t, \omega_t) \Big\}.$$

Claim: $\frac{\partial U_t(\theta^t,\omega^t)}{\partial \theta_t}$ *is nondecreasing in θ_τ for $\tau < t$.*

Proof of Claim. This can be seen by backward induction on t. For $t = T+1$, this holds since $U_{T+1} \equiv 0$. Furthermore, by ICFOC,

$$\frac{\partial U_t(\theta^t, \omega^t)}{\partial \theta_t} = \frac{\partial v_t(\chi_t(\theta^t, \omega^t), \theta_t, \omega_t)}{\partial \theta_t} \\ + \delta \int_{\Theta_{t+1} \times \Omega_{t+1}} \frac{\partial U_{t+1}(\theta^{t+1}, \omega^{t+1})}{\partial \theta_{t+1}} \left(-\frac{\partial F_{t+1}(\theta_{t+1}|\theta_t)}{\partial \theta_t} \right) d\theta_{t+1} dG_t(\omega_{t+1}|\omega_t).$$

The first term is nondecreasing in θ_τ with $\tau < t$ by Strong-MON and SCP, and the second term is nondecreasing in θ_τ by FOSD and the inductive hypothesis for $t+1$. □

By the one stage deviation principle and the Markov assumption it suffices to rule out single deviations from truthtelling. Write

$$\Phi_t(\theta^{t-1}, m_t, \theta_t, \omega^t) := v_t(\chi_t(\theta^{t-1}, m_t, \omega^t), \theta_t, \omega_t) + \psi_t(\theta^{t-1}, m_t, \omega^t) \\ + \delta \int_{\Theta_{t+1} \times \Omega_{t+1}} U_{t+1}(\theta_{-t}^{t+1}, m_t, \omega^{t+1}) d\Lambda_{t+1}(\theta_{t+1}, \omega_{t+1} | \theta_t, \omega_t),$$

and

$$\phi(\theta^{t-1}, m_t, \theta_t, \omega^t) := \Phi_t(\theta^{t-1}, m_t, \theta_t, \omega^t) - U_t(\theta^t, \omega^t).$$

Then honesty is optimal in period t given reporting history θ^{t-1} and state history ω^t if

$$\max_{\theta_t \in \Theta_t} \phi(\theta^{t-1}, m_t, \theta_t, \omega^t) = \phi_t(\theta^{t-1}, m_t, m_t, \omega^t) = 0.$$

Note that

$$\frac{\partial \phi(\theta^{t-1}, m_t, \theta_t, \omega^t)}{\partial \theta_t} = \frac{\partial v_t(\chi_t(\theta^{t-1}, m_t, \omega^t), \theta_t, \omega_t)}{\partial \theta_t} + \delta \int_{\Theta_{t+1} \times \Omega_{t+1}} \frac{\partial U_{t+1}(\theta_{-t}^{t+1}, m_t, \omega^t)}{\partial \theta_{t+1}} \left(-\frac{\partial F_{t+1}(\theta_{t+1}|\theta_t)}{\partial \theta_t} \right) d\theta_{t+1} dG_t(\omega_{t+1}|\omega_t).$$

The first part is nondecreasing in m_t by SCP and Strong-MON, and the second part by FOSD and the Claim. Thus, $\frac{\partial \phi(\theta^{t-1}, m_t, \theta_t, \omega^t)}{\partial \theta_t}$ is nondecreasing in m_t. Since χ satisfies ICFOC, we know that for almost every θ_t, $\frac{\partial \phi(\theta^{t-1}, m_t, \theta_t, \omega^t)}{\partial \theta_t}$ is zero whenever $m_t = \theta_t$. Thus it must be nonnegative when $\theta_t < m_t$ and nonpositive when $\theta_t > m_t$. Hence $\max_{\theta_t \in \Theta_t} \phi(\theta^{t-1}, m_t, \theta_t, \omega^t) = \phi_t(\theta^{t-1}, m_t, m_t, \omega^t) = 0$ almost surely, implying that (χ, ψ) is truthfully implementable with full transparency. $\square$

5 Towards Optimal Mechanisms

The importance of payoff equivalence in static settings derives largely from it providing a tool for designing optimal mechanisms. We show here that the dynamic payoff equivalence results of the previous section can be used analogously in designing profit maximizing mechanisms in dynamic environments.

5.1 Dynamic Virtual Surplus

Suppose that the environment is quasilinear and the principal who designs the mechanism has a payoff of the form

$$\sum_{t=1}^{T} \delta^t \big(p_t - c_t(x^t, \omega_t)\big).$$

The dynamic payoff formula of Theorem 1 allows us to derive an expression for the principal's expected payoff. In order to get a clean formula we make the following assumption.

Assumption 7. *Each $F_{t+1}(\cdot|\theta_t, y_o^t)$ has a density $f_{t+1}(\cdot|\theta_t, y_o^t)$ which is strictly positive on Θ_t, and there exists a lowest first period type $\underline{\theta}_1 > -\infty$.*

Theorem 2. *Assume that the environment is quasilinear and satisfies Assumptions 1–5, and 7. If the choice rule (χ, ψ) is implementable, then the principal's expected payoff* *is given by the* expected dynamic virtual surplus

$$R = \mathbb{E}^{\mu_\chi}\Big[\sum_{t=1}^{T}\delta^t\Big(v_t(\chi^t(\theta^t,\omega^t),\theta_t,\omega_t) - c_t(\chi^t(\theta^t,\omega^t),\omega_t)$$
$$- \frac{1-F_1(\theta_1)}{f_1(\theta_1)}\prod_{\tau=2}^{t}\Big(-\frac{\partial F_\tau(\theta_\tau|\theta_{\tau-1},\chi_o^{\tau-1}(\theta^{\tau-1},\omega^{\tau-1}))/\partial\theta_{\tau-1}}{f_\tau(\theta_\tau|\theta_{\tau-1},\chi_o^{\tau-1}(\theta^{\tau-1},\omega^{\tau-1}))}\Big)\frac{\partial v_t(\chi^t(\theta^t,\omega^t),\theta_t,\omega_t)}{\partial\theta_t}\Big)\Big]$$
$$- U_1(\underline{\theta}_1).$$

Proof. Given a mechanism that implements the choice rule (χ,ψ), the designer's expected payoff is

$$R = \mathbb{E}^{\mu_\chi}\Big[\sum_{t=1}^{T}\delta^t\big(\psi_t(\theta^t,\omega^t) - c_t(\chi^t(\theta^t,\omega^t),\omega_t)\big)\Big].$$

As the environment is quasilinear, we can express the designer's payoff as the difference between the aggregate surplus and the agent's information rent. This gives

$$\begin{aligned}
R &= \mathbb{E}^{\mu_\chi}\Big[\sum_{t=1}^{T}\delta^t\big(v_t(\chi^t(\theta^t,\omega^t),\theta_t,\omega_t) - c_t(\chi^t(\theta^t,\omega^t),\omega_t)\big)\Big] - \mathbb{E}^{\mu_\chi}\Big[U_1(\theta_1)\Big]\\
&= \mathbb{E}^{\mu_\chi}\Big[\sum_{t=1}^{T}\delta^t\big(v_t(\chi^t(\theta^t,\omega^t),\theta_t,\omega_t) - c_t(\chi^t(\theta^t,\omega^t),\omega_t)\big)\Big] - \mathbb{E}^{\mu_\chi}\Big[\int_{\underline{\theta}_1}^{\theta_1}U_1'(q)dq\Big] - U_1(\underline{\theta}_1)\\
&= \mathbb{E}^{\mu_\chi}\Big[\sum_{t=1}^{T}\delta^t\big(v_t(\chi^t(\theta^t,\omega^t),\theta_t,\omega_t) - c_t(\chi^t(\theta^t,\omega^t),\omega_t)\big) - \frac{1-F_1(\theta_1)}{f_1(\theta_1)}U_1'(\theta_1)\Big] - U_1(\underline{\theta}_1).
\end{aligned}$$

Recalling the formula for $U_1'(\theta_1)$ from Theorem 1 and using iterated expectations we arrive at the dynamic virtual surplus. □

The dynamic virtual surplus is the difference between aggregate surplus and the agent's information rent. The first period type plays a special role in contracting as its hazard rate plays a role in determining the information rent in all periods. Heuristically, this follows since the informational asymmetry at the time of contracting is about the first period type, whereas once the agent has revealed his type (and before the second period type is realized) there is only uncertainty about the future types with the principal and the agent sharing a common expectation. Notice that the static virtual surplus obtains as a special case if $T = 1$.

The effect of the intermediate types on the information rent in period t is to determine how informative the first period type is about the current type in period t. Baron and Besanko (1984) appear to have been the first to make this observation. In the special case where current types are independent, the product term in the dynamic virtual surplus vanishes and hence the principal extracts all the surplus in periods $t > 1$. This is to be expected as without linkages between periods, contracting over decisions in periods $t > 1$ is under symmetric information with the principal having all the bargaining power.

An interesting observation concerns the dynamics of distortions. In particular, with serial correlation there are information rents in all periods with probability one. This can be seen by noting that the information rent term vanishes only if the current type equals the supremum or the infimum of its support. If this is the case, then there are no information rents from that period onwards. (A special case of this obtains if already the first period type is the highest possible.) However, even with bounded supports this is a probability zero event. Notice that the lowest and the highest type are in this respect symmetric except for the first period. The intuition for this comes from the above observation that the product term measures the informativeness of the first period type about the period t current type. When the current type hits either boundary the signal is compressed resulting in the first period type losing its informativeness about the future types.

These observations paint a quite different picture of the determinants of distortions than the intuition based on the two-type model of Battaglini (2005), where efficiency always obtains, and the low and the high type have asymmetric roles. For example, if the current type in our model follows a stationary linear first order autoregressive process—a case which we analyze in some detail in sections 5.2 and 6—it is true that distortions vanish asymptotically even though exact efficiency never obtains. However, this is due to the fact that then the correlation of the first period type with the period t type vanishes as t goes to infinity. Indeed, if the AR(1) process is nonstationary, then distortions actually increase over time. This is consistent with the observation that what really matters for distortions is the informativeness of the first period type about future types.

It is straightforward to derive the dynamic virtual surplus for multi-agent environments. Suppose that there are N agents with independent types. As noted above, all the results from the single-agent environment carry over to this case if we denote the agent by i and take $\omega = \theta_{-i}$. This follows since in equilibrium each agent is taking the behavior of the other agents as given so that it is as if he was participating in a single-agent mechanism with randomizations (which in the multi-agent case are due to the uncertainty about the other players' types). In particular, if we assume that values are private and distributions are independent of decisions, then the multi-agent version of the dynamic virtual surplus formula is

$$R_N = \mathbb{E}^{\mu}\Big[\sum_{t=1}^{T}\delta^t\sum_{i=1}^{N}\Big(v_{it}(\chi^t(\theta^t),\theta_{it}) - \frac{c_t(\chi^t(\theta^t))}{N} - \frac{1-F_{i1}(\theta_{i1})}{f_{i1}(\theta_{i1})}\prod_{\tau=2}^{t}\Big(-\frac{\partial F_{i\tau}(\theta_{i\tau}|\theta_{i\tau-1})/\partial\theta_{i\tau-1}}{f_{i\tau}(\theta_{i\tau}|\theta_{i\tau-1})}\Big)\frac{\partial v_{it}(\chi^t(\theta^t),\theta_{it})}{\partial\theta_{it}}\Big)\Big] - \sum_{i=1}^{N}U_{i1}(\underline{\theta}_{i1}). \quad (5.1)$$

5.2 Profit Maximizing Mechanisms

Maximizing the principal's profit amounts to maximizing dynamic virtual surplus over all implementable choice rules. By the revelation principle it suffices to consider choice rules that are truthfully implementable. Until now the only constraints on truthful implementation have been the incentive compatibility constraints which require honest reporting be at least as good as any other strategy. However, in order to analyze profit maximization it is necessary to introduce some form of participation constraints. Since the only binding participation constraint in a dynamic model is the first period one, we incorporate the participation constraints by requiring simply that[10]

$$U_1(\theta_1) \geq \bar{u} \quad \text{for all } \theta_1, \quad (5.2)$$

where $\bar{u}$ is the reservation utility, normalized to be type independent. A particularly simple case obtains if $\frac{\partial v_t(x^t,\theta_t,\omega_t)}{\partial\theta_t} \geq 0$ as then U_1' is non-negative by Theorem 1 so that the participation constraint (5.2) binds only at the lowest type.

As in the static model, maximizing virtual surplus is straightforward if pointwise maximization yields an implementable allocation rule. By Lemma 1 this obtains in a model with SCP and FOSD if the pointwise maximizer happens to be monotone in the sense of Strong-MON. However, as Strong-MON is a condition on an endogenous object, it is natural to look for assumptions on the primitives that ensure that it is satisfied by some pointwise maximizer of the virtual surplus. To this end, we consider here an environment with just one agent. The example in the following section considers a particular multi-agent environment.

The current type is said to evolve according to a linear first order autoregressive process (AR(1)) if for $t \geq 2$,

$$\theta_t = a\theta_{t-1} + \varepsilon_t, \quad (5.3)$$

[10] Participation constraints in periods $t > 1$ can be relaxed without affecting the incentive compatibility constraints or the payoff to the principal by requiring the agent to post bonds.

where $a \geq 0$ and the innovations ε_t are iid with a differentiable cdf $H : \mathbb{R} \to \mathbb{R}$ with finite mean and full support on $\mathbb{R}$. Note that the distribution of the first period type is allowed to be different from H.

Theorem 3. *Assume that the environment satisfies Assumptions 1–7 with the current type evolving according to an AR(1) process.*[11] *Assume further that (i) $X_t \subset \mathbb{R}$ for all t, (ii) SCP, (iii) $\frac{\partial v_t(x_t,\theta_t,\omega_t)}{\partial \theta_t}$ is non-negative and supermodular in (x_t, θ_t) for all ω_t, and (iv) $f_1(\theta_1)/(1 - F_1(\theta_1))$ is strictly increasing. Let $\langle(\chi, \psi), I^*\rangle$ be a fully transparent mechanism that satisfies the following:*

1. *For all t, and for almost every (θ^t, ω^t),*

$$\chi_t(\theta^t, \omega^t) \in \arg\max_{x_t \in X_t} \left\{ v_t(x_t, \theta_t, \omega_t) - c_t(x_t, \omega_t) - \frac{1 - F_1(\theta_1)}{f_1(\theta_1)} a^{t-1} \frac{\partial v_t(x_t, \theta_t, \omega_t)}{\partial \theta_t} \right\},$$

2. *(χ, ψ) satisfies ICFOC with $U_1(\underline{\theta}_1) = \bar{u}$.*

Then $\langle(\chi, \psi), I^*\rangle$ is a profit maximizing mechanism. Furthermore, any other profit maximizing allocation rule coincides with χ with probability 1.

Proof. Recall the formula for dynamic virtual surplus from Theorem 2. The virtual surplus in period t given history (θ^t, ω^t) takes the form

$$v_t(x_t, \theta_t, \omega_t) - c_t(x_t, \omega_t) - \frac{1 - F_1(\theta_1)}{f_1(\theta_1)} \prod_{\tau=2}^{t} \Big(- \frac{\partial F_\tau(\theta_\tau|\theta_{\tau-1})/\partial \theta_{\tau-1}}{f_\tau(\theta_\tau|\theta_{\tau-1})} \Big) \frac{\partial v_t(x_t, \theta_t, \omega_t)}{\partial \theta_t},$$

where we have used the simplification of Assumption 6 that distributions and current payoffs are independent of past decisions. With an AR(1) type we have $F(\theta_t|\theta_{t-1}) = H(\theta_t - \theta_{t-1})$ and

$$- \frac{\partial F(\theta_t|\theta_{t-1})/\partial \theta_{t-1}}{f_t(\theta_t|\theta_{t-1})} = a.$$

Thus the virtual surplus in period t given history (θ^t, ω^t) simplifies to the objective function in condition 1. Hence, if (χ, ψ) satisfies condition 1, it maximizes the expected virtual surplus pointwise almost everywhere. It suffices to prove that it is implementable and satisfies the participation constraint with no slack. Notice that the objective function in 1 is strictly supermodular in (x_t, θ_t) and (x_t, θ_1) by conditions (i)–(iv). By the monotone selection theorem of Milgrom and Shannon (1994), any selection of the

[11] It can be shown that the result holds for first order autoregressive processes of the form

$$b(\theta_t) = a(\theta_{t-1}) + \varepsilon_t,$$

where $a : \mathbb{R} \to \mathbb{R}$ is strictly increasing, weakly concave and differentiable, and $b : \mathbb{R} \to \mathbb{R}$ is strictly increasing, weakly convex and differentiable.

maximizers to this problem is nondecreasing in θ_1 and θ_t. Thus condition 1 implies that the allocation rule χ satisfies Strong-MON. Given that (χ, ψ) satisfies ICFOC, this implies that (χ, ψ) is implementable with full transparency. Note then that Theorem 1 and assumption (iii) imply that U_1 is nondecreasing. Thus condition 2 implies that the mechanism satisfies the participation constraint with no slack. So (χ, ψ) is profit maximizing. Finally, if there exists an implementable mechanism that maximizes virtual surplus pointwise with probability 1, any profit maximizing mechanism has to do so. □

A curious feature of the model with AR(1)-type is that the period t virtual surplus given history (θ^t, ω^t) depends only on the current type and state and the first period type. Since here pointwise maximization of the dynamic virtual surplus yields a profit maximizing allocation rule, the profit maximizing mechanism inherits this feature. Furthermore, the period t virtual surplus, and hence the profit maximizing mechanism, is independent of the distribution of the innovations ε_t. By inspection of the virtual surplus formula in condition 1 we also see that if the process is stationary so that $a < 1$, then distortions are gradually reduced over time; if $a = 1$, they remain constant; and if $a > 1$, the distortions actually increase over time. Finally, conditions 1–3 give a method of constructing the profit maximizing mechanism: First, maximize virtual surplus pointwise. Second, choose a nondecreasing selection from the maximizer correspondence. Third, calculate the transfers using the dynamic payoff formula.

6 Optimal Auctions with AR(1)-types

6.1 Setup

We consider an environment with N agents, the *bidders*, and T periods. For $t = 2, \ldots, T$ the current type of bidder i evolves according to

$$\theta_{it} = a_i \theta_{it-1} + \varepsilon_{it}, \tag{6.1}$$

where ε_{it} are iid with a differentiable cdf H_i with full support on the real line. We let F_{i1} denote the cdf of the first period type, which is assumed to have strictly positive density f_{i1} on some interval $\Theta_{i1} = (\underline{\theta}_i, \overline{\theta}_i)$ with $\underline{\theta}_i > -\infty$. We assume that the hazard rate of the first period type, $f_{i1}(\theta_{i1})/(1 - F_{i1}(\theta_{i1}))$, is nondecreasing.

Types are independent across agents: The first period type distributions are independent across agents, and the period t innovations are drawn independently across agents.

In each period there is a single object for sale. Bidder i's current payoff is $x_{it}\theta_{it} - y_{it}$ so that his current type θ_{it} is simply his valuation for the object.[12] The *seller* who designs the mechanism could have a stochastic reservation value or cost, but for simplicity we take her costs to be given by some deterministic sequence (c_t). The seller's allocation is denoted x_{0t}. Feasibility requires that in each period t, $\sum_{i=0}^{N} x_{it} = 1$ with $x_{it} \geq 0$ for all i.

6.2 Optimal Auctions

An optimal auction for this linear AR(1) environment can be found by simply maximizing the dynamic virtual surplus pointwise for all periods t, and all histories of valuation profiles θ^t. From (5.1) we obtain the period t history θ^t virtual surplus

$$-x_{0t}c_t + \sum_{i=1}^{N} x_{it}\Big(\theta_{it} - \frac{1 - F_{i1}(\theta_{i1})}{f_{i1}(\theta_{i1})} a_i^{t-1}\Big),$$

which is independent of the bidders' valuations in periods $t = 2, \ldots, t-1$ due to the types following linear AR(1) processes. Note that the hazard rate is that of the first period type, not that of the current type.

If there is only one bidder, the allocation rule that maximizes the dynamic virtual surplus pointwise allocates the good to the bidder in period t iff

$$\theta_{it} - \frac{1 - F_{i1}(\theta_{i1})}{f_{i1}(\theta_{i1})} a_i^{t-1} \geq c_t.$$

This is a special case of the one agent environment considered in Theorem 3, and hence the discussion about distortions after Theorem 3 applies here as well. In particular, the distortions may be decreasing, constant, or increasing in t depending on whether $a_i < 1$, $a_i = 1$, or $a_i > 1$.

In general the optimal auctions take the following form.

Proposition 1. *Consider the linear environment with AR(1)-types. Assume that the hazard rate of the first period type, $f_{i1}(\theta_{i1})/(1 - F_{i1}(\theta_{i1}))$, is non-decreasing for all i. Then an optimal sequence of auctions allocates the good to bidder i in period t iff*

$$i \in \arg\max_j \left\{\theta_{jt} - \frac{1 - F_{j1}(\theta_{j1})}{f_{j1}(\theta_{j1})} a_j^{t-1}\right\} \quad \text{and} \quad \theta_{it} - \frac{1 - F_{i1}(\theta_{i1})}{f_{i1}(\theta_{i1})} a_i^{t-1} \geq c_t. \tag{6.2}$$

Proof. This allocation rule is obtained by pointwise maximization of the dynamic virtual surplus. Hence it is optimal as long as it is implementable. But by inspection bidder

[12] This assumes that the valuation can be negative. This is of no consequence to the present analysis, but a reader who finds this disturbing can think of θ_{it} as bidder i's personal state, and define his current valuation to be $\max\{0, \theta_{it}\}$. The results are unchanged. If the first period types are restricted to be positive, then the optimal mechanism does not even require negative values of θ_{it} be ever reported.

i's probability of receiving the good in period t under the allocation rule is nondecreasing in θ_{it} and θ_{i1}. Thus the allocation rule satisfies Strong-MON, and it is truthfully implementable (with full transparency) by Lemma 1. □

Lemma 1 actually implies that the allocation rule (6.2) can be implemented with full transparency. Translated to the present environment (where $\omega = \theta_{-i}$) this means that it is possible to implement the allocation rule in "periodic ex post equilibrium" as defined in Athey and Segal (2007). That is, the allocation rule is implementable even if the bidders are able to observe each others' current types before sending their messages.

In the special case where types are independent across periods (i.e., $a_i = 0$), the optimal sequence of auctions consists of a static optimal auction in the first period followed by a sequence of efficient auctions from which the seller appropriates all the expected surplus since at the time of contracting the bidders do not have private information about their future types.

With persistent types, the optimal sequence of auctions has some interesting features. The marginal revenue from bidder i in period t,

$$\theta_{it} - \frac{1 - F_{i1}(\theta_{i1})}{f_{i1}(\theta_{i1})} a_i^{t-1},$$

depends on the realized first period type. Thus, even if the bidders are symmetric ex ante, they are treated asymmetrically in periods $t > 1$. One of the implications is that the optimal sequence of auctions may generally allocate the good to a bidder whose realized valuation is not the highest among the bidders even with ex ante symmetric bidders. If the bidders are asymmetric ex ante, but the coefficient a_i is the same across bidders, then the ordering of the bidders with respect to the severity of their distortions remains constant over time, even though the distortions may be vanishing (or even increasing if $a_i > 1$). However, if the persistence of valuations (i.e., a_i) varies across bidders, this ordering may vary too, as more persistent is the valuation, the slower is the decay of distortions. It is thus possible that some bidders are "favored" in the early rounds, and others in later rounds.

Suppose that there is no good to be allocated in period 1, but otherwise the model is as before. The optimal sequence of auctions can then be interpreted as consisting of a period of bidding for "terms of trade" followed by a sequence of auctions where bidders' distortions (or their preferential treatment) is determined by their first period bids. In particular, in all periods $t > 1$, the payments can be calculated using a modified Vickrey mechanism where, given distortions determined by the first period types, the winning bidder pays an amount equal to the smallest bid that he could have made and still won the object. Such a pricing mechanism is in fact ex post incentive compatible, since past reports from periods other than the first one do not affect future distortions. The firts period payments, which can be solved for using the dynamic payoff formula, effectively amount to a tariff on reducing future distortions.

The Appendices

A Proof of Theorem 1

Let φ be an implementable choice rule. Let $\langle M, g, I\rangle$ be a mechanism that implements φ when the agent uses strategy σ. Define $U_1 : \Theta_1 \to \mathbb{R}$ by

$$U_1(\theta_1) := \mathbb{E}^{\mu_\sigma|\theta_1}\Big[\sum_{t=1}^{T}\delta^t u_t\left(g^t\left(\bar{\sigma}^t(\theta,\omega),\omega^t\right),\theta_t,\omega_t\right)\Big].$$

Since $\langle M, g, I\rangle$ implements φ we have $\varphi(\theta,\omega) = g\left(\bar{\sigma}(\theta,\omega),\omega\right)$. This implies that $\mu_\sigma|\theta_1 = \mu_\varphi|\theta_1$ and

$$U_1(\theta_1) = \mathbb{E}^{\mu_\varphi|\theta_1}\Big[\sum_{t=1}^{T}\delta^t u_t\left(\varphi^t(\theta,\omega),\theta_t,\omega_t\right)\Big].$$

Thus U_1 is independent of the mechanism used to implement φ. By the revelation principle φ is truthfully implementable. Hence in deriving the properties of U_1 it suffices to consider honest reporting in the canonical mechanism $\langle\varphi,\varphi_o\rangle$.

In the canonical mechanism the agent does not observe the state, but forms expectations about it based on the observable decisions and the history of reports. In order to capture the agent's beliefs we construct a consistent family of regular conditional probability distributions (rcpd) as follows. Suppose that a rcpd $H_\tau(\cdot|\varphi_o^{\tau-1}(m^{\tau-1},\cdot))$ on Ω^τ exists for all message histories $m^{\tau-1}$, and all periods $\tau \le t$. (The conditioning here is on the random variable $\varphi_o^{\tau-1}(m^{\tau-1},\cdot) : \Omega^\tau \to Y^{\tau-1}$.) For $t = 1$ this holds with $H_1 \equiv G_1$. Suppose that it holds for some $t \ge 1$. For any history m^t, the distribution $H_t(\cdot|\varphi_o^{t-1}(m^{t-1},\cdot))$ and the kernel $G_{t+1}(\cdot|\omega_t,\varphi^t(m^t,\omega^t))$ induce a distribution on Ω^{t+1}. The random variable $\varphi_o^t(m^t,\cdot) : \Omega^{t+1} \to Y_o^t$ maps into a standard Borel space. Hence, there exists a rcpd $H_{t+1}(\cdot|\varphi_o^t(m^t,\cdot))$ on Ω^{t+1} (see, e.g., Durrett, 2004). Consistency follows by construction. We postpone an argument showing that the result is independent of the chosen versions of the rcpd's to the end of the proof.

The Markovian assumption means that the agent's true past types are not payoff-relevant even if his past reports may be. So if he finds it optimal to be honest after truthful histories, he will also find it optimal to be honest after misreporting. Hence it suffices to keep track of the history of reports and observable decisions.

Consider period t. Let $U_t(\theta^t, y_o^{t-1})$ denote the continuation expected utility of current type θ_t from optimal reporting given truthful reporting history θ^{t-1} and a history of observable decisions y_o^{t-1}. That is,

$$U_t(\theta^t, y_o^{t-1}) = \sup_{m_t \in \Theta_t} \mathbb{E}^{H_t(\cdot|\varphi_o^{t-1}(\theta^{t-1},\cdot)=y_o^{t-1})}\Big[u_t(\varphi^t(\theta^{t-1}, m_t, \omega^t), \theta_t, \omega_t) \\ + \delta \int_{\Theta_{t+1}} U_{t+1}(\theta_{-t}^{t+1}, m_t, \tilde{y}_o^t) dF_{t+1}(\theta_{t+1}|\theta_t, \tilde{y}_o^t)\Big], \quad \text{(A.1)}$$

where we have defined $\tilde{y}_o^t := (y_o^{t-1}, \varphi_{ot}(\theta^{t-1}, m_t, \omega^t))$ to simplify notation.

Since our definition of an optimal strategy corresponds to a notion of ex ante incentive compatibility, the agent's continuation payoff may in general differ from U_t in any period $t > 1$ in a set of probability zero under the measure μ_φ induced by the choice rule φ. This follows, since although honesty maximizes the expected utility at period 1, it could fail to maximize the agent's (ex post) continuation utility following histories that have probability zero under μ_φ. However, ex ante incentive compatibility implies that for all t, there exists a set A_t of truthful histories $(\theta^{t-1}, y_o^{t-1})$ that has probability 1 under μ_φ such that following any history in A_t honesty is optimal for almost every θ_t. (The histories in the set A_t have to be consistent in the sense that for each $(\theta^{t-1}, y_o^{t-1})$, there exists ω^{t-1} such that $\varphi_o^{t-1}(\theta^{t-1}, \omega^{t-1}) = y_o^{t-1}$.) In what follows we will restrict attention to the set A_t of such histories.

The proof of the theorem essentially consists of establishing the following claim.

Claim: *For all $t = 1, \ldots, T+1$, and all $(\theta^{t-1}, y_o^{t-1}) \in A_t$, $U_t(\theta^t, y_o^{t-1})$ is a Lipschitz-continuous function of θ_t uniformly on A_t, and for almost every θ_t it satisfies the recursion*

$$\frac{\partial U_t(\theta^t, y_o^{t-1})}{\partial \theta_t} = \mathbb{E}^{H_t(\cdot|\varphi_o^{t-1}(\theta^{t-1},\cdot)=y_o^{t-1})}\Big[\frac{\partial u_t(\varphi^t(\theta^t, \omega^t), \theta_t, \omega_t)}{\partial \theta_t} \\ + \int_{\Theta_{t+1}} \frac{\partial U_{t+1}(\theta^{t+1}, \tilde{y}_o^t)}{\partial \theta_{t+1}} \left(-\frac{\partial F_{t+1}(\theta_{t+1}|\theta_t, \tilde{y}_o^t)}{\partial \theta_t}\right) d\theta_{t+1}\Big],$$

where $\tilde{y}_o^t := (y_o^{t-1}, \varphi_{ot}(\theta^t, \omega^t))$.

Proof. The proof is by backward induction on t. We assume throughout that $(\theta^{t-1}, y_o^{t-1}) \in A_t$. The claim holds vacuously for $t = T+1$ since $U_{T+1}(\theta^{T+1}, \varphi_o^T) \equiv 0$. Suppose then that it holds for some $t+1$. Notice that although U_{t+1} satisfies the properties only in some set A_{t+1} of probability 1, we may redefine it such that it satisfies them everywhere as it only enters the formulas through an integral. We show first that the objective function in (A.1) is differentiable and Lipschitz in θ_t. Fix θ_t, θ_t' and consider the integral term within the expectation in (A.1), which gives the agent's expected continuation payoff. We suppress the domain of integration Θ_{t+1} in all of the integrals that follow to simplify notation. We have

$$\frac{\int U_{t+1}(\theta_{-t}^{t+1}, m_t, \tilde{y}_o^t) dF_{t+1}(\theta_{t+1}|\theta_t, \tilde{y}_o^t) - \int U_{t+1}(\theta_{-t}^{t+1}, m_t, \tilde{y}_o^t) dF_{t+1}(\theta_{t+1}|\theta_t', \tilde{y}_o^t)}{\theta_t - \theta_t'}$$
$$= \frac{\int U_{t+1}(\theta_{-t}^{t+1}, m_t, \tilde{y}_o^t) d\left(F_{t+1}(\theta_{t+1}|\theta_t, \tilde{y}_o^t) - F_{t+1}(\theta_{t+1}|\theta_t', \tilde{y}_o^t)\right)}{\theta_t - \theta_t'}.$$

Integration by parts gives

$$\frac{1}{\theta_t - \theta_t'}\Bigg[\left[U_{t+1}(\theta_{-t}^{t+1}, m_t, \tilde{y}_o^t)\left(F_{t+1}(\theta_{t+1}|\theta_t, \tilde{y}_o^t) - F_{t+1}(\theta_{t+1}|\theta_t', \tilde{y}_o^t)\right)\right]_{\theta_{t+1}=\inf \Theta_{t+1}}^{\sup \Theta_{t+1}}$$
$$- \int \frac{\partial U_{t+1}(\theta_{-t}^{t+1}, m_t, \tilde{y}_o^t)}{\partial \theta_{t+1}} \left(F_{t+1}(\theta_{t+1}|\theta_t, \tilde{y}_o^t) - F_{t+1}(\theta_{t+1}|\theta_t', \tilde{y}_o^t)\right) d\theta_{t+1}\Bigg].$$

The first term vanishes since U_{t+1} is Lipschitz in θ_{t+1} by the induction hypothesis and the distributions have finite means by Assumption 4.[13] Thus we are left with

$$- \int \frac{\partial U_{t+1}(\theta_{-t}^{t+1}, m_t, \tilde{y}_o^t)}{\partial \theta_{t+1}} \frac{F_{t+1}(\theta_{t+1}|\theta_t, \tilde{y}_o^t) - F_{t+1}(\theta_{t+1}|\theta_t', \tilde{y}_o^t)}{\theta_t - \theta_t'} d\theta_{t+1}.$$

Assumption 5 implies that $F_{t+1}(\theta_{t+1}|\cdot, \tilde{y}_o^t)$ is $b_{t+1}(\theta_{t+1})$-Lipschitz. Since U_{t+1} is B-Lipschitz in θ_{t+1} for some $B < +\infty$, the absolute value of the integrand in the above expression is bounded from above by $Bb_{t+1}(\theta_{t+1})$ for all θ_{t+1}, where b_{t+1} is integrable by assumption. Thus we may use the dominated convergence theorem to pass the limit through the integral to obtain the derivative of the integral term in (A.1)

$$\int \frac{\partial U_{t+1}(\theta_{-t}^{t+1}, m_t, \tilde{y}_o^t)}{\partial \theta_{t+1}} \left(-\frac{\partial F_{t+1}(\theta_{t+1}|\theta_t, \tilde{y}_o^t)}{\partial \theta_t}\right) d\theta_{t+1} \leq B \int b_{t+1}(\theta_{t+1}) d\theta_{t+1}.$$

This implies that the integral term in (A.1) is differentiable and Lipschitz in θ_t uniformly in $(\theta^{t-1}, m_t, \tilde{y}_o^t)$. Note that so is the current payoff function u_t by Assumption 2. Thus the term within the expectation in (A.1) is a differentiable and uniformly Lipschitz function of θ_t. Since H_t is a probability measure and hence finite, we may use the dominated convergence theorem again (with the Lipschitz constant as the dominating function) to conclude that the objective function in (A.1) is differentiable and uniformly Lispchitz in θ_t. This implies that the value function U_t is uniformly Lipschitz in θ_t. Furthermore, we can apply the envelope theorem of Milgrom and Segal (2002) to find its partial derivative. Since U_t is Lipschitz in θ_t, its partial derivative with respect to

[13] If Θ_{t+1} is bounded this follows from U_{t+1} being continuous in θ_{t+1}. Otherwise, bound the first term by using triangle inequality twice. Then take the term with $\theta_{t+1} = \sup \Theta_{t+1}$. U_{t+1} Lipschitz implies $|U_{t+1}| \leq a + b\theta_{t+1}$ with $b > 0$ for large θ_{t+1}. So this term is bounded by $(a + b\theta_{t+1})(1 - F_{t+1}(\theta_{t+1}|\theta_t, \tilde{y}_o^t))$. Since $F_{t+1}(\cdot|\theta_t, \tilde{y}_o^t)$ has a finite mean, this bound vanishes as $\theta_{t+1} \to +\infty$. The other term is analogous.

θ_t exists almost every θ_t. Similarly, honesty is optimal for almost all θ_t. Therefore, for almost every θ_t,

$$\frac{\partial U_t(\theta^t, y_o^{t-1})}{\partial \theta_t} = \mathbb{E}^{H_t(\cdot|\varphi_o^{t-1}(\theta^{t-1},\cdot)=y_o^{t-1})}\Big[\frac{\partial u_t(\varphi^t(\theta^t,\omega^t),\theta_t,\omega_t)}{\partial \theta_t} + \int \frac{\partial U_{t+1}(\theta^{t+1},\tilde{y}_o^t)}{\partial \theta_{t+1}}\left(-\frac{\partial F_{t+1}(\theta_{t+1}|\theta_t,\tilde{y}_o^t)}{\partial \theta_t}\right) d\theta_{t+1}\Big].$$

This proves the claim. □

Now the formula for $U_1'(\theta_1)$ in the statement of the theorem is obtained by iterating the recursive formula in the Claim. Since the conditional probability distributions we constructed are regular, we may use iterated expectations to arrive at the expression which only involves the kernels F_t and G_t. Note that, by construction, the function U_1 defined at the beginning of the proof coincides with the function $U_1(\theta_1, \emptyset)$ obtained by iteration, since ex ante incentive compatibility requires that honesty is optimal in the first period for each type θ_t.

We conclude the proof by returning to the issue of choosing versions of the rcpd's H_t. By inspection the formula for U_1' is independent of the versions. Thus choosing any versions will do.

B Statement and proof of Lemma B.1

Lemma B.1 *Assume that the environment satisfies Assumption 4. Then Assumption 5 implies*

$$\exists B < +\infty : \quad \left|\frac{\partial}{\partial \theta_t}\mathbb{E}[\theta_{t+1}|\theta_t, y_o^t]\right| \leq B \quad \forall(\theta_t, y_o^t).$$

Assume that Assumption 4 is satisfied. Fix θ_t. Then Assumption 5 implies

$$\begin{aligned}\left|\frac{\partial}{\partial \theta_t}\int \theta_{t+1} dF_{t+1}(\theta_{t+1}|\theta_t, y_o^t)\right| &= \left|\lim_{\theta_t' \to \theta_t} \frac{\int \theta_{t+1} d\left(F_{t+1}(\theta_{t+1}|\theta_t', y_o^t) - F_{t+1}(\theta_{t+1}|\theta_t, y_o^t)\right)}{\theta_t' - \theta_t}\right| \\ &= \left|-\lim_{\theta_t' \to \theta_t}\int \frac{F_{t+1}(\theta_{t+1}|\theta_t', y_o^t) - F_{t+1}(\theta_{t+1}|\theta_t, y_o^t)}{\theta_t' - \theta_t} d\theta_{t+1}\right| \\ &= \left|-\int \frac{\partial F_{t+1}(\theta_{t+1}|\theta_t, y_o^t)}{\partial \theta_t} d\theta_{t+1}\right|,\end{aligned}$$

where the second inequality follows by integration by parts and the argument of footnote 13 to show that the substitution term vanishes. The last equality follows by the dominated convergence theorem since the integrand is bounded for all θ_t' by the integrable function b_{t+1}. Thus

$$\left|\frac{\partial}{\partial\theta_t}\int \theta_{t+1} dF_{t+1}(\theta_{t+1}|\theta_t, y_o^t)\right| \leq \int \left|\frac{\partial F_{t+1}(\theta_{t+1}|\theta_t, y_o^t)}{\partial\theta_t}\right| d\theta_{t+1} \leq \int b_{t+1}(\theta_{t+1}) d\theta_{t+1},$$

from which the claim follows by taking $B := \int b_{t+1}(\theta_{t+1}) d\theta_{t+1}$.

References

Athey, S., and I. Segal (2007): "An Efficient Dynamic Mehcanism," Mimeo, Stanford University.

Baron, D. P., and D. Besanko (1984): "Regulation and Information in a Continuing Relationship," *Information Economics and Policy*, 1(3), 267–302.

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

Bergemann, D., and J. Välimäki (2007): "Dynamic Marginal Contribution Mechanism," Discussion Paper, Cowless Foundation.

Courty, P., and H. Li (2000): "Sequential Screening," *Review of Economic Studies*, 67(4), 697–717.

Durrett, R. (2004): *Probability: Theory and Examples*. Brooks/Cole–Thomson Learning, 3rd edn.

Gershkov, A., and B. Moldovanu (2007): "The Dynamic Assignment of Heterogenous Objects: A Mechanism Design Approach," Discussion Paper, University of Bonn.

Milgrom, P., and I. Segal (2002): "Envelope Theorems for Arbitrary Choice Sets," *Econometrica*, 70(2), 583–601.

Milgrom, P., and C. Shannon (1994): "Monotone Comparative Statics," *Econometrica*, 62(1), 157–180.

Mirrlees, J. A. (1971): "An Exploration in the Theory of Optimum Income Taxation," *Review of Economic Studies*, 38(114), 175–208.

Myerson, R. B. (1981): "Optimal Auction Design," *Mathematics of Operation Research*, 6(1), 58–73.

——— (1986): "Multistage Games with Communication," *Econometrica*, 54(2), 323–58.

Pancs, R. (2007): "Optimal Information Disclosure in Auctions and Negotiations: A Mechanism Design Approach," Mimeo, Stanford University.

Rochet, J.-C., and L. Stole (2003): "The Economics of Multidimensional Screening," in *Advances in Economics and Econometrics: Theory and Applications - Eight World Congress*, ed. by M. Dewatripont, L. P. Hansen, and S. J. Turnovsky, vol. 1 of *Econometric Society Monographs*, pp. 150–197. Cambridge University Press.