

Dynamic Mechanism Design:

Revenue Equivalence, Profit Maximization and Information Disclosure*

PRELIMINARY AND INCOMPLETE.

Alessandro Pavan

Northwestern University

Ilya Segal

Stanford University

Juuso Toikka

Stanford University

June 2, 2008

Abstract

This paper examines the problem of how to design incentive-compatible mechanisms in environments in which the agents' private information evolves stochastically over time and in which decisions have to be made in each period. The environments we consider are fairly general in that the agents' types are allowed to evolve in a non-Markov way, decisions are allowed to affect the type distributions and payoffs are not restricted to be separable over time. Our first result is the characterization of a dynamic formula for (the derivative of) the agents' equilibrium payoffs in an incentive-compatible mechanism. The formula summarizes all local first-order conditions taking into account how current types affect the dynamics of expected payoffs. The formula generalizes the familiar envelope condition from static mechanism design: the key difference is that a variation in the current types now impacts payoffs in all subsequent periods both directly and through the effect on the distributions of future types. We first identify assumptions on the primitive environment that guarantee that our dynamic payoff formula is a necessary condition for incentive compatibility. Next, we specialize this formula to quasi-linear environments and use it to establish a dynamic revenue-equivalence result. Lastly, we turn to the characterization of sufficient conditions for incentive compatibility. We then apply the results to study the properties of revenue-maximizing mechanisms in a variety of applications that include dynamic auctions with AR(k) values and the provision of experience goods.

JEL Classification Numbers: D82, C73, L1.

Keywords: dynamic mechanisms, asymmetric information, stochastic processes, long-term contracting, incentives

*This paper supercedes previous working papers "Revenue Equivalence, Profit Maximization, and Transparency in Dynamic Mechanisms" by Segal and Toikka and "Long-Term Contracting in a Changing World" by Pavan. Acknowledgements to be added.

1 Introduction

We consider the problem of how to design 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 the agents' private information as well as the dependence of information on past allocations. For example, it covers as special cases such problems as the allocation of resources to agents whose valuations follow a stochastic process, the procedures for selling new experience goods whose value is refined by the buyers upon consumption, or the design of multiperiod procurement auctions for bidders whose cost parameters evolve stochastically over time and may exhibit learning-by-doing effects.

The fundamental difference between dynamic and static mechanism design is that in the former, an agent has access to a lot more potential deviations. Namely, instead of a simple misrepresentation of his true type, the agent can make this representation conditional on the information he has observed in the mechanism, in particular on his past types, his past reports (which need not have been truthful), and what he inferred about the other agents' types in the course of the mechanism. Despite the resulting complications, we deliver some general necessary conditions for incentive compatibility and some sufficient conditions, and use them to characterize profit-maximizing mechanisms in several applications.

The cornerstone of our analysis is the derivation of a formula for the derivative of an agent's expected payoff in an incentive-compatible mechanism with respect to his private information. Similarly to Mirrlees's first-order approach for static environments (Mirrlees, 1971), our formula (hereafter referred to as *dynamic payoff formula*) provides an envelope-theorem condition summarizing local incentive compatibility constraints. In contrast to the static model, however, the derivation of this formula relies on incentive compatibility in all the future periods, not just in one given period. Furthermore, unlike some of the earlier papers about dynamic mechanism design, we identify conditions on the primitive environment for which the dynamic payoff formula is a necessary condition for *any* incentive-compatible mechanism (not just for "well-behaved" ones). In addition to carrying over the usual static assumptions of "smoothness" of the agent's payoff function in his type and connectedness of the type space (see, e.g., Milgrom and Segal, 2002), the dynamic setting requires additional assumptions on the stochastic process governing the evolution of each agent's information. Intuitively, our dynamic payoff formula represents the impact of an (infinitesimal) change in the agent's current type on his equilibrium expected payoff. This change can be decomposed into two parts. The first one is the familiar effect of the current type on the agent's expected utility, as in static mechanism design. The second part captures the indirect effect of the current type on the expected utility through its impact on the type distributions in each of the subsequent periods. Note that in general the current type may affect the future type distributions

directly as well as indirectly through its impact on the type distributions in intermediate periods. All changes in the type distributions are then evaluated by looking at their ultimate impact on the agent’s utility, holding constant the agent’s messages to the mechanism (by the usual envelope theorem logic).

The dynamic payoff formula can be established either by iterating backward the local incentive-compatibility conditions or by using the *quantile function theorem* (see, e.g., Angus, 1994) to represent the agents’ types as the result of independent innovations (shocks). While the two approaches lead to the same formula, the conditions on the primitive environment that validate this formula a necessary condition for incentive compatibility are somewhat different. In this sense the two approaches are complementary (see also Eso and Szentes, 2007, for a similar approach in a two-period-one-decision model).

To ease the exposition, in the first part of the paper (Section 3) we consider an environment with a single agent who observes all the relevant history of the mechanism. There we derive the envelope formula that determines the agent’s equilibrium payoff in a incentive-compatible mechanism. In Section 4 we then show how to adapt the envelope formula to a multi-agent environment. The key difference between the two settings is that in the latter an agent observes only a part the entire history generated by the mechanism: an agent must thus form beliefs about the unobserved types of the other agents as well as the decisions that the mechanism has induces with these agents. We show that the derivation for the single-agent case extends to multi-agent mechanisms provided that the stochastic processes governing the evolution of the agents’ types are independent of one another, except through their effect on the decisions that are observed by the agents. In other words, we show how the familiar “Independent Types” assumption for static mechanism design should be properly adjusted to a dynamic setting to guarantee that the agents’ equilibrium payoffs can still be pinned down by an envelope formula.

For the special case of quasilinear environments, we first use the dynamic envelope formula to establish a dynamic “*revenue equivalence theorem*” that links the payment rules in any two Bayesian incentive-compatible mechanisms that implement the same allocation rule. In particular, if we have a single agent who participates in a deterministic mechanism, this theorem pins down, in each state, the total payment that is necessary to implement a given allocation rule, up to a scalar that does not depend on the state. With many agents, or with a stochastic mechanism, the theorem pins down the *expected* payments as function of each agent’s type history, where the expectation is with respect to the other agents’ types and/or the stochastic decisions taken by the mechanism. However, if one requires a strong form of “robustness”—according to which the mechanism must remain incentive-compatible even if an agent is shown at the very beginning of the game all the other agents’ (future) types—then the theorem again pins down the total payments for *each state*.

Next, we use the dynamic envelope formula to express the expected profits in an incentive-compatible and individually rational mechanism as the expected “*virtual surplus*,” appropriately defined for the dynamic setting. This derivation uses only the agents’ local incentive constraints, and only the participation constraints of the lowest-types in the initial period. Ignoring all the other incentive and participation constraints yields a dynamic “*Relaxed Program*,” which is in general a dynamic programming problem. In particular, the Relaxed Program gives us a simple intuition for the optimal distortions introduced by a profit-maximizing principal: Since only the first-period participation constraints bind (this is due to the unlimited bonding possibilities in the quasilinear environment with unbounded transfers), the distortions are created to balance the rent-extraction versus efficiency trade-off, as perceived from the perspective of period one. However, due to informational linkages in the stochastic type process, the principal will not only distort the agent’s consumption in period one but also in any subsequent period whenever his type in period t is informative about the first-period type. The informativeness is here measured by an “*information index*” that captures all the direct and indirect effects of the first-period type on the type distributions in all subsequent periods.

It turns out that when an agent’s type in period $t > 1$ hits its highest or lowest possible value, the informational linkage disappears and the principal implements the efficient level of consumption in that period (provided that payoffs are additively time-separable). However, for intermediate types in period t , the optimal mechanism entails distortions (for example, when types are positively correlated over time in the sense of First-Order Stochastic Dominance, and the agent’s payoffs satisfy the single-crossing property, the optimal mechanism entails downward distortions). Thus, in contrast to the static model, with a continuous but bounded type space, distortions in each period $t > 1$ are never monotonic in the agent’s type. This is also in contrast with the results of Battaglini (2005) for the case of a Markov process with only two types in each period.

Studying the Relaxed Program is not fully satisfactory unless one also provides sufficient conditions for its solution to satisfy all of the remaining incentive and participation constraints. We are indeed able to provide some such conditions. In particular, we show that in the case where the agents’ types follow a Markov process and their payoffs are Markovian in their types (so that it is enough to check one-stage deviations from truth-telling), a sufficient condition for an allocation rule to be implementable is that the partial derivative of the agent’s expected utility with respect to his current type when he misreports be nondecreasing in the report. One can then use the dynamic payoff formula to calculate this partial derivative—the condition is fairly easy to check. (Unfortunately, this condition is not necessary for incentive-compatibility—a tight characterization is evasive because of the multidimensional decision space of the problem.) This sufficient condition also turns useful when checking incentive compatibility in some non-Markov settings that are

sufficiently “separable.”

In some standard settings we can actually state an even simpler sufficient condition for incentive compatibility, which also ensures that incentive compatibility is robust to an agent learning in advance all of the other agents’ types (and therefore to any weaker form of information leakage in the mechanism). This condition is that the transitions that describe the evolution of the agents’ private information are monotone in the sense of First-Order Dominance, the payoffs satisfy a single-crossing property, and the allocation rule is “*strongly monotonic*” in the sense that the consumption of a given agent in any period is nondecreasing in each of the agent’s type reports, for any given profile of reports by the other agents.

In Section 5, we apply the general results to a few simple, yet illuminating, applications. The analysis proves especially simple when the agents’ types follow an autoregressive stochastic process of degree k (AR(k)). If we assume in addition that each agent’s payoff is affine in his types (but not necessarily in his consumption), then the principal’s Relaxed Program turns out to be very similar to the expected social surplus maximization program, the only difference being that the agents’ true values in each period are replaced by their corresponding “virtual values.” In the AR(k) case, the difference between an agent’s true value and his virtual value in period t , which can be called his “handicap” in period t , is determined by the agent’s first-period type, the hazard rate of the first period type’s distribution, and the “impulse response coefficient” of the AR(k) process.¹ Intuitively, the impulse response coefficient determines the informational link between period t and period 1, while the first-period hazard rate captures the importance that the principal assigns to the trade-off between efficiency and rent-extraction as perceived from period one’s perspective (just as in the static model). Importantly, since the handicaps depend only on the first-period type reports, the Relaxed Program at any period $t \geq 2$ can be solved by running an efficient (i.e., expected surplus-maximizing) mechanism on the handicapped values. Thus, while building an efficient mechanism may in general require solving an involved dynamic programming problem (due to possible intertemporal payoff interactions), once a solution is found it can be easily adapted to obtain a solution to the Relaxed Program. We also use the fact that the solution to the Relaxed Program looks “quasi-efficient” from period 2 onward to show that it can be implemented in a mechanism that is incentive compatible from period 2 onward (following truth-telling in period one). This can be done for example using the “Team Mechanism” payments proposed by Athey and Segal (2007) to implement efficient allocation rules. As for verifying incentives in period 1, we have only been able to do it in a few special settings.

We also consider two other applications. The first one is the designing of sequential auctions for environments in which the agents’ payoffs are time-separable while their private types follow an

¹The term “handicapped auction” was first used in Eso and Szentes (2007).

AR(k) process. This setting is particularly simple because the Relaxed Program separates across periods and states and so we do not need to solve a dynamic programming problem. Under the standard monotone hazard rate assumption on the agents' initial type distribution and the standard third-derivative assumption on their utility functions, the Relaxed Program is solved by a Strongly Monotone allocation rule, which then implies that it is implementable in an incentive-compatible mechanism (and one that is robust to information leakage). 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 intermediate reports. This can be interpreted as a scheme where the agents make up-front payments that reduce their future distortions.

The second application is one in which an agent receives a signal about his unknown valuation for a new good each time he consumes it. The agent's expected value for the good then follows a martingale. The solution to the efficient dynamic programming problem in this setting takes the form of a stopping rule. The solution to the profit-maximization problem looks similar, except that the agent again makes a first-period report that determines his up-front payment and his subsequent handicaps. This optimal mechanism achieves a strictly higher expected profit than any pricing policy, even a history-contingent one.

The rest of the paper is organized as follows. Section 2 discusses the related literature. Section 3 presents the results for the single-agent case. Section 4 extends the analysis to quasi-linear settings with multiple agents. Section 5 presents a few applications. Section 5.3 contains proofs omitted in the main text.

2 Related Literature²

The last few years have witnessed a fast-growing literature on dynamic mechanism design. A number of papers propose mechanisms for implementing efficient (profit-maximizing) mechanism that are the dynamic analogues of static VCG and expected-externality mechanisms (see, for example, Athey and Segal (2007) and Bergemann and Valimaki, 2008, and the references therein). These papers do not characterize incentive compatibility, but provide some mechanisms that turn out to be incentive-compatible.

Our analysis is more closely related to the pioneering work of Baron and Besanko (1984) on regulation of a natural monopoly and the more recent paper of Courty and Li (2000) on advance ticket sales. Both papers consider a two-period model with one agent and use the first-order approach to derive optimal mechanisms. The agent's types in the two periods are serially correlated

²This section is still very much incomplete. We apologize to the many authors who feel that their work should have been discussed and that we omitted here.

and this correlation determines the distortions in the optimal mechanism. Courty and Li also provide some sufficient conditions for the allocation rule to be implementable. Our paper builds on the ideas in these papers but extends the approach to allow for multiple periods, multiple agents, and for more general specification of the payoff and information structure. Contrary to these early papers, we also provide conditions on the primitive environment that validate the “first-order approach.”

Related is also a more recent paper by Battaglini (2005) who considers a model with one agent and two types and derives an optimal selling mechanism for a monopolist facing a consumer whose type follows a Markov process. Our results for a model with continuous types indicate that many of his predictions seem specific to the special setting with only two types. We discuss in more detail the differences between the results in the two papers in subsection 4.6.³

Gershkov and Moldovanu (2007) consider both efficient and profit maximizing mechanisms to allocate a fixed set of objects to buyers that arrive randomly over time. While the model has multiple agents, they assume that each agent lives only instantaneously. Hence the problem that each agent faces is actually static. The paper derives a payoff-equivalence result which is essentially a static payoff equivalence result applied separately to each short lived agent. In contrast, we allow the agents to be long lived.⁴

Eso and Szentes (2007) consider a two-period model with many agents but with a single decision in the second period. They propose a different approach than that in Baron and Besanko (1984) and Courty and Li (2000) to the characterization of optimal mechanisms. Their approach consists in using the Probability Integral Transform Theorem to represent an agent’s second-period type as a function of his first-period type and a random shock that is independent of the first-period type. In Section 3.3 we show how the Probability Integral Transform Theorem can be used recursively in a setting with more than two periods to describe the entire stochastic process that governs the evolution of the agents’ private information by means of serially independent shocks. We then show how the independent-shock representation can be used to derive our dynamic payoff formula under a somewhat different set of assumptions. Eso and Szentes also derive a profit-maximizing auction and coin the term “handicapped auction” to describe it. However, in their two-period AR(1) setting, it turns out that any incentive-compatible mechanism, not just a profit-maximizing one, can be viewed as a “handicapped auction.” What we find more surprising is that under the special assumptions of an AR(k) type process and affine payoffs, then even with many periods the optimal mechanism remains an “handicapped mechanism.” The distinguishing feature of such

³See also our companion paper “On the Dynamics of Distortions in Long-term Contracting,” for a further discussion.

⁴Other recent papers that study dynamic profit-maximizing mechanisms include Bognar, Börgers, and Meyerter-Vehn, 2008, and Zhang, 2008. The key difference between these papers and ours is that these papers look at particular issues that can emerge in dynamic environments, such as costly participation, while our abstracts from some of these issues but instead provides a more general characterization of incentive-compatibility.

mechanisms is that the allocation in a given period depends only on that period's reports and the reports in the first period; it is thus independent of the reports in all intermediate periods.⁵

The paper is also related to a more “macroish” literature on optimal dynamic taxation. While the early literature typically assumes i.i.d. shocks (e.g. Green 1987, Thomas and Worrall, 1990, Atkeson and Lucas, 1992), the more recent literature considers the case of persistent private information (e.g. Fernandes and Phelan, 2000, Golosov, Kocherlakota, and Tsyvinski, 2003, Kocherlakota, 2005, Golosov and Tsyvinski, 2006, Kapicka, 2008, Tchistyi, 2006, Biais, Mariotti, Plantin, and Rochet, 2007, Zhang, 2007, Williams, 2008). While our work shares several modelling assumptions with some of the papers in this literature, its key distinctive aspect is the general characterization of incentive compatibility as opposed to the features of the optimal mechanism in the context of specific applications.⁶

Dynamic mechanism design is also inherently related to the literature on multidimensional screening, as noted, e.g., in Rochet and Stole (2003). Indeed, it is the multidimensional nature of the problem that prevents a complete characterization of all implementable allocation rules. Nevertheless, there is a sense in which incentive compatibility is much easier to ensure in a dynamic mechanism than in a static multidimensional mechanism. This is because in a dynamic environment an agent is asked to report each dimension of his private information before learning the subsequent dimensions. By implication there are fewer deviations than in the corresponding static environment in which the agents observe all the dimensions at once. Because of this, the set of allocation rules that are implementable in a dynamic environment proves to be significantly larger than the set of allocation rules that are implementable in the corresponding static multidimensional environment. For example, the profit-maximizing dynamic allocation rules we characterize are typically not implementable if the agents were to observe all of their private information at the outset of the mechanism.

We also touch here upon the issue of transparency in mechanisms. Calzolari and Pavan (2006b) and Calzolari and Pavan (2006a) study its role in environments in which downstream actions (e.g. resale offers in secondary markets, or more generally contract offers in sequential common agency) are not contractible upstream. Pans (2007) also studies the role of transparency in environments where agents take nonenforceable actions such as investment or information acquisition.

⁵Another key difference between the two papers is that, while Eso and Szentes use their model to study primarily the effects of the seller's information disclosures on surplus extraction, here we focus on the characterization of incentive compatibility in general dynamic mechanisms. For this purpose, it is essential to allow for non-Markov processes and non-time-separable preferences, and to permit decisions to affect the type distributions.

⁶Some of the works in this literature limit the analysis to the characterization of local first-order conditions (e.g. the inverse Euler equation) and either leave the dynamics of the optimal mechanism unspecified or they solve it numerically.

3 Single-agent case

3.1 General setup

3.1.1 The Environment

We consider an environment with one agent and finitely many periods, indexed by $t = 1, 2, \dots, T$. In each period t there is a contractible *decision* $y_t \in Y_t$, whose outcome is observed by the agent. (In the next section we apply the model to a more general setup where y_t is the part of the decision taken in period t that is observed by the agent.) Each Y_t is assumed to be a measurable space with the sigma-algebra left implicit. The set of all period- t decision histories is denoted $Y^t \equiv \prod_{\tau=1}^t Y_\tau$.⁷ For the full histories we drop the superscripts so that y is an element of $Y \equiv Y^T$.

Before the period t decision is taken, the agent privately observes his *current type* $\theta_t \in \Theta_t \equiv (\underline{\theta}_t, \bar{\theta}_t) \subset \mathbb{R}$ where $-\infty \leq \underline{\theta}_t \leq \bar{\theta}_t \leq +\infty$. We implicitly endow the set Θ_t with the Borel sigma-algebra. The set of possible type histories at period t is denoted by $\Theta^t \equiv \prod_{\tau=1}^t \Theta_\tau$. An element θ of $\Theta \equiv \Theta^T$ is referred to as the agent's *type*.

The distribution of the current type θ_t may depend on the entire history of types and decisions $(\theta^{t-1}, y^{t-1}) \in \Theta^{t-1} \times Y^{t-1}$. In particular, we assume that the distribution of θ_t is governed by a history-dependent probability measure (“kernel”) $F_t(\cdot | \theta^{t-1}, y^{t-1})$ on Θ_t such that $F_t(A | \cdot) : \Theta^{t-1} \times Y^{t-1} \rightarrow \mathbb{R}$ is measurable for all measurable $A \subset \Theta_t$.⁸ Note that the distribution of θ_t depends only on variables observed by the agent. We denote the collection of kernels by

$$F \equiv \langle F_t : \Theta^{t-1} \times Y^{t-1} \rightarrow \Delta(\Theta_t) \rangle_{t=1}^T,$$

where for any measurable set A , $\Delta(A)$ denotes the set of probability measures on A . We abuse notation by using $F_t(\cdot | \theta^{t-1}, y^{t-1})$ to denote also the cumulative distribution function (cdf) corresponding to the measure $F_t(\cdot | \theta^{t-1}, y^{t-1})$.

The agent is a von Neumann-Morgenstern decision maker whose preferences over lotteries over $\Theta \times Y$ are represented by the expectation of a (measurable) Bernoulli utility function

$$U : \Theta \times Y \rightarrow \mathbb{R}.$$

(Although some form of time separability of U is typically assumed in applications, it is not needed for the general results.)

An often encountered special case in applications is one where private information evolves in a

⁷By convention, all products of measurable spaces encountered in the text are endowed with the product sigma-algebra.

⁸Throughout, we adopt the convention that for any set A , $A^0 \equiv \{\emptyset\}$.

Markovian fashion, and where the agent's payoff is Markovian in the following sense

Definition 1 *The environment is Markov if*

1. for all t , and all $(\theta^{t-1}, y^{t-1}) \in \Theta^{t-1} \times Y^{t-1}$, $F_t(\cdot | \theta^{t-1}, y^{t-1})$ does not depend on θ^{t-2} , and
2. there exists functions $\langle A_t : \Theta^t \times Y^t \rightarrow \mathbb{R}_{++} \rangle_{t=1}^{T-1}$ and $\langle B_t : \Theta_t \times Y^t \rightarrow \mathbb{R} \rangle_{t=1}^T$ such that for all $(\theta, y) \in \Theta \times Y$,

$$U(\theta, y) = \sum_{t=1}^T \left(\prod_{\tau=1}^{t-1} A_\tau(\theta_\tau, y^\tau) \right) B_t(\theta_t, y^t). \quad (1)$$

Condition (1) ensures that in any given period t after observing history (θ^t, y^t) , the agent's von Neumann-Morgenstern preferences over future lotteries depend on his type history θ^t only through the current type θ_t . In particular, it encompasses the case of *additive separable* preferences ($A_t(\theta_t, y^t) = 1$ for all t) as well as the case of *multiplicative separable* preferences ($B_t(\theta_t, y^t) = 0$ for all $t < T$).

3.1.2 Mechanisms

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 and the mechanism responds with a (possibly randomized) decision that may depend on the entire history of messages sent up to period t , and on past decisions. By the Revelation Principle (adapted from Myerson, 1986), for any standard solution concept, any distribution on $\Theta \times Y$ that can be induced as an equilibrium outcome in any mechanism can be induced as an equilibrium outcome of a "direct mechanism" in which the agent is asked to report the current type in each period, and in equilibrium he reports truthfully.

Let $m_t \in \Theta_t$ denote the agent's period- t message, and let $m^t \equiv (m_1, \dots, m_t)$.

Definition 2 *A direct mechanism is a collection*

$$\Omega \equiv \langle \Omega_t : \Theta^t \times Y^{t-1} \rightarrow \Delta(Y_t) \rangle_{t=1}^T$$

such that for all t , and all measurable $A \subset Y_t$, $\Omega_t(A | \cdot) : \Theta^t \times Y^{t-1} \rightarrow [0, 1]$ is measurable.

(The notation $\Omega_t(A | m^t, y^{t-1})$ stands for the probability of the mechanism generating $y_t \in A \subset Y_t$ given history $(m^t, y^{t-1}) \in \Theta^t \times Y^{t-1}$.)

Given a direct mechanism Ω , and a history $(\theta^{t-1}, m^{t-1}, y^{t-1}) \in \Theta^{t-1} \times \Theta^{t-1} \times Y^{t-1}$, the following sequence of events takes place in each period t :

1. The agent privately observes his current type $\theta_t \in \Theta_t$ drawn according to $F_t(\cdot | \theta^{t-1}, y^{t-1})$.

2. The agent sends a message $m_t \in \Theta_t$.
3. The mechanism selects a decision $y_t \in Y_t$ according to $\Omega_t(\cdot|m^t, y^{t-1})$.

A (*pure*) *strategy* for the agent in a direct mechanism is thus a collection of measurable functions

$$\sigma \equiv \langle \sigma_t : \Theta^t \times \Theta^{t-1} \times Y^{t-1} \rightarrow \Theta_t \rangle_{t=1}^T.$$

Definition 3 A strategy σ is truthful if for all t and all $((\theta^{t-1}, \theta_t), m^{t-1}, y^{t-1}) \in \Theta^t \times \Theta^{t-1} \times Y^{t-1}$,

$$\sigma_t((\theta^{t-1}, \theta_t), m^{t-1}, y^{t-1}) = \theta_t.$$

This definition defines a unique strategy that requires the agent to report his current type truthfully following all histories, including non-truthful ones.

In order to describe expected payoffs, it is convenient to develop some more notation. First we define histories. For all $t = 0, 1, \dots, T$, let

$$H_t \equiv (\Theta^t \times \Theta^{t-1} \times Y^{t-1}) \cup (\Theta^t \times \Theta^t \times Y^{t-1}) \cup (\Theta^t \times \Theta^t \times Y^t),$$

where by convention $H_0 = \{\emptyset\}$, and $H_1 = \Theta_1 \cup (\Theta_1 \times \Theta_1) \cup (\Theta_1 \times \Theta_1 \times Y_1)$. Then H_t is the set of all histories terminating within period t , and, accordingly, any $h \in H_t$ is referred to as a *period- t history*. We let

$$H \equiv \bigcup_{t=0}^T H_t$$

denote the set of all histories. A history $(\theta^s, m^t, y^u) \in H$ is a *successor* to history $(\hat{\theta}^j, \hat{m}^k, \hat{y}^l) \in H$ if (1) $(s, t, u) \geq (j, k, l)$, and (2) $(\theta^j, m^k, y^l) = (\hat{\theta}^j, \hat{m}^k, \hat{y}^l)$. A history $h = (\theta^s, m^t, y^u) \in H$ is a *truthful history* if $\theta^t = m^t$.

Fix a direct mechanism Ω , a strategy σ , and a history $h \in H$. Let $\mu[\Omega, \sigma|h]$ denote the (unique) probability measure on $\Theta \times \Theta \times Y$ —the product space of types, messages, and decisions—induced by assuming that following history h in mechanism Ω , the agent follows strategy σ in the future. More precisely, let $h = (\theta^s, m^t, y^u)$. Then $\mu[\Omega, \sigma|h]$ assigns probability one to $(\tilde{\theta}, \tilde{m}, \tilde{y})$ such that $(\tilde{\theta}^s, \tilde{m}^t, \tilde{y}^u) = (\theta^s, m^t, y^u)$. Its behavior on $\Theta \times \Theta \times Y$ is otherwise induced by the stochastic process that starts in period s with history h , and whose transitions are determined by the strategy σ , mechanism Ω , and kernels F . If h is the null history we then simply write $\mu[\Omega, \sigma]$. We also adopt the convention of omitting σ from the arguments of μ when σ is the truthful strategy. Thus $\mu[\Omega]$ is the ex-ante measure induced by truthtelling while $\mu[\Omega|h]$ is the measure induced by the truthful strategy following history h .

Given this notation, we write the agent’s expected payoff when following history h he plays according to strategy σ in the future as $\mathbb{E}^{\mu[\Omega, \sigma]|h}[U(\tilde{\theta}, \tilde{y})]$.⁹

For most of the results we use ex-ante rationality as our solution concept. That is, we require the agent’s strategy to be optimal when evaluated at date zero, before learning θ_1 . In a direct mechanism this corresponds to ex-ante incentive compatibility defined as follows.

Definition 4 *A direct mechanism Ω is ex-ante incentive compatible (ex-ante IC) if for all strategies σ ,*¹⁰

$$\mathbb{E}^{\mu[\Omega]}[U(\tilde{\theta}, \tilde{y})] \geq \mathbb{E}^{\mu[\Omega, \sigma]}[U(\tilde{\theta}, \tilde{y})].$$

This notion of IC is arguably the weakest for a dynamic environment. Thus deriving necessary conditions for this notion gives the strongest results. However, for certain results it is convenient to define IC at a given history.

Definition 5 *Given a direct mechanism Ω , the agent’s value function is a mapping $V^\Omega : H \rightarrow \mathbb{R}$ such that for all $h \in H$,*

$$V^\Omega(h) = \sup_{\sigma} \mathbb{E}^{\mu[\Omega, \sigma]|h}[U(\tilde{\theta}, \tilde{y})].$$

Definition 6 *Let $h \in H$. A direct mechanism Ω is incentive compatible at history h (IC at h) if*

$$\mathbb{E}^{\mu[\Omega]|h}[U(\tilde{\theta}, \tilde{y})] = V^\Omega(h).$$

In words, Ω is IC at h if truthful reporting in the future maximizes the agent’s expected continuation payoff following history h . This definition is flexible in that it allows us to generate different notions of IC as special cases by requiring IC at all histories in a particular subset. For example, ex-ante IC is equivalent to requiring IC only at the null history. Or in a static model (i.e., if $T = 1$), the standard definition of interim incentive compatibility obtains by requiring Ω to be IC at all histories where the agent knows only his type. In a dynamic model a natural alternative is to require that if the agent has been truthful in the past, he finds it optimal to continue to report truthfully. This is obtained by requiring Ω to be IC at all truthful histories.

The Principle of Optimality implies the following lemma.

Lemma 1 *If Ω is IC at h , then for $\mu[\Omega]|h$ -almost all successors h' to h , Ω is IC at h' .*

⁹Throughout we use “tildes” to denote random variables with the same symbol without the tilde corresponding to a particular realization.

¹⁰Restricting attention to pure strategies is without loss: By the Revelation Principle the agent can be assumed to follow the truthful pure strategy in equilibrium. As for deviations, a mixed strategy (or a collection of behavioral strategies) induces a lottery over payoffs from pure strategies. Thus, if there is a profitable deviation to a mixed strategy, then there is also a profitable deviation to a pure strategy in the support of the mixed strategy.

In particular, if Ω is ex-ante IC, then truthtelling is also sequentially optimal at truthful future histories h with probability one, and the agent's equilibrium payoff at such histories is given by $V^\Omega(h)$ with probability one. We will sometimes find it convenient to focus on such histories, and they are the only ones that matter for ex-ante expectations.

3.2 Necessary Conditions for IC: Recursive Approach

3.2.1 Backward-Induction Formula

We now set out to derive a recursive formula for (the derivative of) the agent's expected payoff in an incentive compatible mechanism. This formula extends to dynamic models the standard use of the envelope theorem in static models to pin down the dependence of the agent's equilibrium utility on his true type (see, e.g., Milgrom and Segal, 2002). We begin with a heuristic derivation of the result. First recall the standard approach with $T = 1$, which expresses the derivative of the agent's equilibrium payoff in an IC mechanism with respect to his type as the partial derivative of his utility function with respect to the true type holding the truthful equilibrium message fixed:

$$\frac{dV^\Omega(\theta_1)}{d\theta_1} = \int_{Y_1} \frac{\partial U(\theta_1, y_1)}{\partial \theta_1} d\Omega_1(y_1|\theta_1) = \mathbb{E}^{\mu^{[\Omega]}|\theta_1} \left[\frac{\partial U(\tilde{\theta}_1, \tilde{y}_1)}{\partial \theta_1} \right].$$

(For the moment we ignore the precise conditions for the argument to be valid.)

With $T > 1$, we may be interested in evaluating the equilibrium payoff starting from any period t . In general, the agent's continuation utility from truthtelling following a truthful history $h = (\theta^t, \theta^{t-1}, y^{t-1})$ is

$$\mathbb{E}^{\mu^{[\Omega]}|h} \left[U(\tilde{\theta}, \tilde{y}) \right] = \int U(\theta, y) dF_{T+1}(\theta_{T+1}|\theta^T, y^T) d\Omega_T(y_T|m^T, y^{T-1}) \cdots dF_{t+1}(\theta_{t+1}|\theta^t, y^t) d\Omega_t(y_t|m^t, y^{t-1}) \Big|_{m=\theta},$$

where $dF_{T+1}(\theta_{T+1}|\theta^T, y^T) \equiv 1$. Assume for the moment that this expression is sufficiently well-behaved so that the derivatives encountered below exist. Suppose one now replicates the argument from the static case. That is, consider the agent's problem of choosing a continuation strategy given truthful history $(\theta^t, \theta^{t-1}, y^{t-1})$. Assuming that an envelope argument applies, we differentiate with respect to the agent's current type θ_t holding the agent's truthful future messages fixed. The current type directly enters the payoff in two ways. First, it enters the agent's utility function U . This gives the term $\mathbb{E}^{\mu^{[\Omega]}|h} [\partial U(\tilde{\theta}, \tilde{y}) / \partial \theta_t]$. Second, it enters the kernels F . This gives (after integrating

by parts and differentiating within the integral) for each $\tau > t$ the term

$$-\mathbb{E}^{\mu[\Omega]|h} \left[\int \frac{\partial F_\tau(\theta_\tau|\tilde{\theta}^{\tau-1}, \tilde{y}^{\tau-1})}{\partial \theta_t} \frac{\partial V^\Omega((\tilde{\theta}^{\tau-1}, \theta_\tau), \tilde{\theta}^{\tau-1}, \tilde{y}^{\tau-1})}{\partial \theta_\tau} d\theta_\tau \right].$$

This suggests that a marginal change in the current type effects the equilibrium payoff through two different channels. First, it changes the agent's payoff from any allocation. Second, it changes the distribution of future types in all periods $\tau > t$, and hence leads to a change in the period- τ continuation utility captured by the derivative of the value function V^Ω evaluated at the appropriate history.

While the above heuristic derivation isolates the effects of the current type on the agent's equilibrium payoff, it does not address the technical conditions for the derivation to be valid. In fact, in general the derivatives of the future value function can not be assumed to exist so that the actual formal argument is more involved. In particular, we do not want to impose any restriction on the mechanism Ω to guarantee differentiability of the value function. This would clearly be restrictive, for example, for the purposes of deriving implications for optimal mechanisms. Instead, we seek to identify *properties of the environment* that guarantee that the value function is sufficiently well behaved.

Our derivation makes use of the following key assumptions.

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

Assumption 2 For all t , and all $(\theta^{t-1}, y^{t-1}) \in \Theta^{t-1} \times Y^{t-1}$, $\int |\theta_t| dF_t(\theta_t|\theta^{t-1}, y^{t-1}) < +\infty$.

Assumption 3 For all t , and all $(\theta^{t-1}, y^{t-1}) \in \Theta^{t-1} \times Y^{t-1}$, the cdf $F_t(\cdot|\theta^{t-1}, y^{t-1})$ is strictly increasing on Θ_t .

Assumption 4 For all t , and all $(\theta, y) \in \Theta \times Y$, $\partial U(\theta, y)/\partial \theta_t$ exists and is bounded uniformly in (θ, y) .

Assumption 5 For all t , all $\tau < t$, and all $(\theta^t, y^{t-1}) \in \Theta^t \times Y^{t-1}$, $\partial F_t(\theta_t|\theta^{t-1}, y^{t-1})/\partial \theta_\tau$ exists. Furthermore, for all t , there exists integrable $B_t : \Theta_t \rightarrow \mathbb{R} \cup \{+\infty\}$ such that for all $\tau < t$, and all $(\theta^t, y^{t-1}) \in \Theta^t \times Y^{t-1}$,

$$|\partial F_t(\theta_t|\theta^{t-1}, y^{t-1})/\partial \theta_\tau| \leq B_t(\theta_t).$$

Assumption 6 For all t , and all $y^{t-1} \in Y^{t-1}$, the probability measure $F_t(\cdot|\theta^{t-1}, y^{t-1})$ is continuous in θ^{t-1} in the total variation metric.¹¹

¹¹See, e.g., Stokey and Lucas (1989) for the definition of the total variation metric.

Assumptions 1 and 4 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 2 and 3 are also typically made in static models. Assumption 2 about the existence of the expectation is trivially satisfied if Θ_t is bounded. Assumption 3 is a full support assumption, which is related to Assumption 1. While Assumption 1 requires that the set Θ_t of all feasible types be connected, Assumption 3 implies that the set of relevant types is a connected set.¹²

Assumption 5 requires that the distribution of the current type depend sufficiently smoothly on past types. The motivation for it is essentially the same as for requiring that, even in static settings, utility depends smoothly on types (i.e., Assumption 4). In a dynamic model the agent’s expected payoff depends on his true type both through the utility function U and the kernels F . For the expected payoff to depend smoothly on types, both U and F need to have this property.¹³ Since this assumption does not have an immediate counterpart in the static model, it is instructive to consider what restrictions it imposes on the stochastic process for θ_t . In particular, it implies that the partial derivative of the expected current type with respect to any past type θ_τ , $\frac{\partial}{\partial \theta_\tau} \mathbb{E}[\theta_t | \theta^{t-1}, y^{t-1}]$, exists and is bounded uniformly in (θ^{t-1}, y^{t-1}) —see Lemma A1 in the Appendix.

It turns out that for non-Markov models Assumption 5 by itself does not impose enough regularity on the dependence of the kernels on past types, and hence we impose also Assumption 6.

We are now ready to state our first main result.

Proposition 1 *Suppose Assumptions 1-6 hold. (In the Markov case, Assumption 6 can be dispensed with.) If Ω is IC at the truthful history $h^{t-1} \equiv (\theta^{t-1}, \theta^{t-1}, y^{t-1})$, then*

$$\begin{aligned}
 & V^\Omega(\theta_t, h^{t-1}) \text{ is Lipschitz continuous in } \theta_t, \text{ and for a.e. } \theta_t, \\
 & \frac{\partial V^\Omega(\theta_t, h^{t-1})}{\partial \theta_t} = \mathbb{E}^{\mu[\Omega] | (\theta_t, h^{t-1})} \left[\frac{\partial U(\tilde{\theta}, \tilde{y})}{\partial \theta_t} - \sum_{\tau=t+1}^T \int \frac{\partial F_\tau(\theta_\tau | \tilde{\theta}^{\tau-1}, \tilde{y}^{\tau-1})}{\partial \theta_t} \frac{\partial V^\Omega((\tilde{\theta}^{\tau-1}, \theta_\tau), \tilde{\theta}^{\tau-1}, \tilde{y}^{\tau-1})}{\partial \theta_\tau} d\theta_\tau \right]. \tag{IC-FOC}
 \end{aligned}$$

The recursive formula (IC-FOC) pins down how the agent’s equilibrium utility varies as a function of the current type θ_t . It is a dynamic generalization of the static envelope theorem formula sometimes referred to as the “Mirrlees’s trick” (Mirrlees, 1971). (Of course, the static result obtains as a special case when $T = t = 1$.) As suggested in the heuristic derivation preceding

¹²Depending on the notion of IC used, full support may not be needed as long as IC is imposed for all types in Θ_t . However, without it, the interpretation becomes an issue. For example, consider a static model where $\Theta_1 = [0, 1]$ but where F assigns probability one to the set $\{0, 1\}$. Is this a model with a continuous type space in which IC is imposed for all $\theta_1 \in [0, 1]$, or a model with two types with IC imposed only on $\theta_1 \in \{0, 1\}$?

¹³This presumes the assumptions have to be stated separately for the primitives U and F . A weaker joint (or “reduced form”) assumption imposing restrictions directly on the expected payoff would suffice.

the result, an infinitesimal change in the current type has two kinds of effects in a dynamic model. First, there is a direct effect on the final utility from decisions, which is captured by the partial derivative of U with respect to θ_t . This is the only effect present in static models. With more than one period, there is a second, indirect, effect through the impact of the current type on the distribution of future types. This is captured by the sum within the expectation. The effect of the current type θ_t on the distribution of period τ type is captured by the partial derivative of F_τ with respect to θ_t . The induced change in utility is evaluated by considering the partial derivative of the period τ value function V_τ with respect to θ_τ .

Remark 1 *We have assumed that the information the agent receives in each period (his current type) is one-dimensional. If in a given period the agent's current type were multidimensional, we could still derive the same necessary condition (IC-FOC) for incentive compatibility by restricting the agent to observing each dimension of his current type at a time and reporting each dimension before observing the subsequent ones. (This restriction only reduces the set of possible deviations and therefore preserves incentive compatibility.) However, incentive compatibility is harder to ensure when the agent observes several dimensions at once (see Remark 2 for more detail).*

3.2.2 Role of the assumptions

To better appreciate the role of the assumptions in Proposition 1, it is useful to consider a few counterexamples. The first one illustrates the role of Assumptions 1 and 3. The other two illustrate the role of assumption 5.

Example 1 (Lack of full support) *Consider the following simple quasi-linear environment where $T = 2$, $\Theta_1 = (0, 1)$, $\Theta_2 = (0, 3)$, $Y_1 = \emptyset$, $y_2 = (x, p) \in Y_2 = \{0, 1\} \times \mathbb{R}$, and*

$$F_2(\theta_2|\theta_1) = \begin{cases} 0 & \text{if } \theta_2 < 0 \\ (1 - \theta_1)\theta_2 & \text{if } \theta_2 \in [0, 1) \\ 1 - \theta_1 & \text{if } \theta_2 \in [1, 2) \\ 1 - \theta_1 + \theta_1(\theta_2 - 2) & \text{if } \theta_2 \in [2, 3) \\ 1 & \text{if } \theta_2 \geq 3 \end{cases}$$

The agent's payoff is $U(\theta, y) = \theta_2 x - p$. This environment corresponds, for example, to a setting where the agent is a buyer whose period-1 type represents the probability he assigns to his period-2 valuation for an indivisible object (denoted by θ_2) being higher than 2. Now consider the following deterministic mechanism

$$\Omega(\theta_1, \theta_2) = \begin{cases} (1, p) & \text{if } \theta_2 \in [p, 3) \\ (0, 0) & \text{otherwise} \end{cases}$$

with $p \in [2, 3]$.¹⁴ That is, there is a posted price p in period 2. It is easy to see that this mechanism is IC at any history. The value function, evaluated at period-one history, is thus $V^\Omega(\theta_1) = \mathbb{E}[\theta_2 | \theta_2 \in [p, 3]] \Pr(\theta_2 \geq p | \theta_1) = \frac{p+3}{2} \theta_1 (3-p)$. The derivative of this function with respect to θ_1 depends on p , which is in contrast with what is predicted by (IC-FOC). The example also illustrates the failure of the revenue equivalence result for quasi-linear settings documented in the static literature; we will come back to the relation between this result and Proposition 1 in Section 4.

Example 2 (Discontinuous transitions) Next, consider the same example discussed above but now assume that $\Theta_1 = \Theta_2 = (0, 1)$ and that

$$F_2(\theta_2 | \theta_1) = \begin{cases} \theta_2 & \text{if } \theta_1 < 1/2 \\ \theta_2^2 & \text{if } \theta_1 \geq 1/2 \end{cases}$$

Now consider the following deterministic mechanism:

$$\Omega(\theta_1, \theta_2) = \begin{cases} (1, p) & \text{if } \theta_1 \in [1/2, 1) \\ (0, 0) & \text{otherwise} \end{cases}$$

with $p \in (1/2, 2/3)$. That is, there is now a forward contract offered in period 1 at price p for delivery at period 2. This mechanism is clearly IC at any history. The corresponding value function is

$$V^\Omega(\theta_1) = \begin{cases} 0 & \text{if } \theta_1 < \frac{1}{2} \\ \frac{2}{3} - p & \text{if } \theta_1 \geq \frac{1}{2} \end{cases}$$

The value function is thus not Lipschitz continuous in this example and, once again, revenue equivalence fails to obtain.

Example 3 (Lack of equi-Lipschitz continuity) As another example of the role that assumption 5 plays for the result in Proposition 1, consider an environment in which $Y_1 = (0, +\infty)$, $Y_2 = \emptyset$, $\Theta_1 = \Theta_2 = (0, 1)$ and where, for any y_1 , $F_2(\theta_2 | \theta_1, y_1)$ is continuously differentiable in both θ_1 and θ_2 but is not equi-Lipschitz continuous in θ_1 . The agent's payoff is $U(\theta, y) = \theta_2$. Then consider the following mechanism

$$\Omega(\theta_1) = \arg \max_{y_1 \in Y_1} \int \theta_2 dF_2(\theta_2 | \theta_1, y_1)$$

By construction, the mechanism is IC at any history. Furthermore, by assumption, for any y_1 , the function $g(\theta_1, y_1) \equiv \int \theta_2 dF_2(\theta_2 | \theta_1, y_1)$ is continuously differentiable in θ_1 . Following Example 1 in

¹⁴In this example, we are abusing notation by letting $\Omega(x, p)$ denote the distribution that assigns measure one to (x, p) .

Milgrom and Segal (2002), one can then find transitions F_2 such that the derivative of $g(\theta_1, y_1)$ with respect to θ_1 is not bounded by any integral function which make the value function discontinuous in θ_1 .

3.2.3 Closed-form expression for expected payoff derivative

The recursive formula for the partial derivative of V^Ω with respect to current type θ_t in Proposition 1 can be iterated backwards to get a closed form formula. Although this can in principle be done under the assumptions of the proposition, a more compact expression obtains if we make the following additional assumption.

Assumption 7 For all t and all $(\theta^{t-1}, y^{t-1}) \in \Theta^{t-1} \times Y^{t-1}$, the function $F_t(\cdot | \theta^{t-1}, y^{t-1})$ is absolutely continuous and its density satisfies $f_t(\theta_t | \theta^{t-1}, y^{t-1}) > 0$ for a.e. $\theta_t \in \Theta_t$.

The existence of a strictly positive density allows us to write the formula in terms of expectation operators rather than integrals. Using iterated expectations then yields the following result.

Proposition 2 Suppose Assumptions 1-7 hold. (In the Markov case, Assumption 6 can be dispensed with.) If Ω is IC at the truthful history $h^{t-1} \equiv (\theta^{t-1}, \theta^{t-1}, y^{t-1})$, then

$$V^\Omega(\theta_t, h^{t-1}) \text{ is Lipschitz continuous in } \theta_t, \text{ and for a.e. } \theta_t, \quad (2)$$

$$\frac{\partial V^\Omega(\theta_t, h^{t-1})}{\partial \theta_t} = \mathbb{E}^{\mu[\Omega] | (\theta_t, h^{t-1})} \left[\sum_{\tau=t}^T J_\tau^\tau(\tilde{\theta}^\tau, \tilde{y}^{\tau-1}) \frac{\partial U(\tilde{\theta}, \tilde{y})}{\partial \theta_\tau} \right],$$

where $J_t^t(\tilde{\theta}^t, \tilde{y}^{t-1}) \equiv 1$ and

$$J_t^\tau(\theta^\tau, y^{\tau-1}) \equiv \sum_{\substack{K \in \mathbb{N}, l \in \mathbb{N}^K: \\ t=l_0 < \dots < l_K = \tau}} \prod_{k=1}^K I_{l_{k-1}}^{l_k}(\theta^{l_k}, y^{l_k-1}) \text{ for } \tau > t,$$

with

$$I_l^m(\theta^m, y^{m-1}) \equiv - \frac{\partial F_m(\theta_m | \theta^{m-1}, y^{m-1}) / \partial \theta_l}{f_m(\theta_m | \theta^{m-1}, y^{m-1})}.$$

Proof. We proceed by backward induction. For $t = T$ the claim follows immediately from Proposition 1. Suppose now that it holds for all $\tau > t$ for some $t \in \{1, \dots, T-1\}$. We will show that it

holds for t . Using iterated expectations and the induction hypothesis, (IC-FOC) can be written as

$$\begin{aligned}
\frac{\partial V^\Omega(\theta_t, h^{t-1})}{\partial \theta_t} &= \mathbb{E}^{\mu[\Omega]|\theta_t, h^{t-1}} \left[\frac{\partial U(\tilde{\theta}, \tilde{y})}{\partial \theta_t} + \sum_{\tau=t+1}^T I_t^\tau(\tilde{\theta}^\tau, \tilde{y}^{\tau-1}) \frac{\partial V^\Omega(\tilde{\theta}^\tau, \tilde{\theta}^{\tau-1}, \tilde{y}^{\tau-1})}{\partial \theta_\tau} \right] \\
&= \mathbb{E}^{\mu[\Omega]|\theta_t, h^{t-1}} \left[\frac{\partial U(\tilde{\theta}, \tilde{y})}{\partial \theta_t} + \sum_{\tau=t+1}^T I_t^\tau(\tilde{\theta}^\tau, \tilde{y}^{\tau-1}) \sum_{s=\tau}^T J_\tau^s(\tilde{\theta}^s, \tilde{y}^{s-1}) \frac{\partial U(\tilde{\theta}, \tilde{y})}{\partial \theta_s} \right] \\
&= \mathbb{E}^{\mu[\Omega]|\theta_t, h^{t-1}} \left[\sum_{\tau=t}^T J_t^\tau(\tilde{\theta}^\tau, \tilde{y}^{\tau-1}) \frac{\partial U(\tilde{\theta}, \tilde{y})}{\partial \theta_\tau} \right],
\end{aligned}$$

where the last equality follows by the definition of $J_t^\tau(\tilde{\theta}^\tau, \tilde{y}^{\tau-1})$. ■

Intuition for (2) is as follows. I_l^m can be interpreted as the “direct informational index” of signal θ_l about signal θ_m . J_t^τ can be interpreted as “total informational index” of θ_τ about θ_t . It incorporates all the ways in which θ_t can affect θ_τ , both directly and through the intermediate signals observed by the agent. Note that in calculating J_t^τ each possible chain of effect must be counted exactly once. For example, in the Markov case, $I_l^m = 0$ for $l < m - 1$, and hence $J_t^\tau(\theta^\tau, y^{\tau-1}) = \prod_{k=t+1}^{\tau} I_{k-1}^k(\tilde{\theta}^k, \tilde{y}^{k-1})$. More generally, the following example suggests that the total informational indices could be viewed as “impulse responses” of the stochastic process for θ to an infinitesimal change in θ_t .

Example 4 Let θ_t evolve according to an AR(k) process:

$$\theta_t = \sum_{j=1}^k \phi_j \theta_{t-j} + \varepsilon_t,$$

where $\theta_t = 0$ for any $t \leq 0$, $\phi_j \in \mathbb{R}$ for any $j = 1, \dots, k$, and ε_t is the realization of the random variable $\tilde{\varepsilon}_t$ distributed according to some cdf G_t with strictly positive density over \mathbb{R} , independent from all $\tilde{\varepsilon}_s$, $s \neq t$. For convenience, hereafter we let $\phi_j \equiv 0$ for all $j > k$. Then

$$F_\tau(\theta_\tau | \theta^{\tau-1}, y^{\tau-1}) = G_\tau \left(\theta_\tau - \sum_{j=1}^k \phi_j \theta_{\tau-j} \right),$$

so that for any $\tau > t$,

$$I_t^\tau(\theta^\tau, y^{\tau-1}) \equiv - \frac{\partial F_\tau(\theta_\tau | \theta^{\tau-1}, y^{\tau-1}) / \partial \theta_t}{f_\tau(\theta_\tau | \theta^{\tau-1}, y^{\tau-1})} = \phi_{\tau-t},$$

and

$$J_t^\tau(\theta^\tau, y^{\tau-1}) = \sum_{M \in \mathbb{N}, l \in \mathbb{N}^M: t=l_0 < \dots < l_M=\tau} \prod_{m=1}^M \phi_{l_m - l_{m-1}}.$$

Thus in this case the total informational index $J_t^\tau(\theta^\tau, y^{\tau-1})$ is simply the “impulse response function” for the AR(k) process. Note also that here the total informational index is only a function of t and τ but not of (θ, y) . In the special case of an AR(1) process we have

$$I_t^\tau(\theta^\tau, y^{\tau-1}) = \begin{cases} \phi_1 & \text{if } \tau = t + 1 \\ 0 & \text{otherwise,} \end{cases}$$

which implies that $J_t^\tau(\theta^\tau, y^{\tau-1}) = (\phi_1)^{\tau-t}$.

3.3 Necessary Conditions for IC: Independent Shocks

In this section, we illustrate an alternative approach to the characterization of the agent’s payoff in a dynamic mechanism. This approach is based on the idea that any stochastic process admits an equivalent representation in which the information the agent receives over time can be described as a function of “shocks” that are serially independent (see also Eso and Szentes, 2007, for a similar approach in a two-period-one-decision model). This approach complements the one illustrated in the previous section in that it leads to a different set of assumptions on the primitive environment that guarantee that the agent’s payoff in any incentive-compatible mechanism can be pinned down by an envelope condition.

We start by defining what we mean when we say that a process admits an independent-shock representation. Next, we define in what sense this representation is “equivalent” to the original one and hence can be used as an alternative approach to the characterization of incentive-compatible mechanisms. We then proceed by showing how the formula for the (derivative of the) agent’s payoff identified in the previous section simplifies when the agent is asked to report the shocks instead of his types. Finally, we conclude by showing that *any* stochastic process admits a particular independent-shock representation. We then use this canonical representation to identify conditions for the primitive environment that guarantee that in the corresponding independent-shock representation the agent’s reduced-form payoff satisfies the analog of the envelope formula derived in the previous section. While these conditions differ from the ones identified above they lead to the same payoff formula when the latter is expressed in terms of the primitive representation.

Definition 7 Let $\tilde{\varepsilon} \equiv (\tilde{\varepsilon}_t)_{t=1}^T$ denote a random vector with support $\mathcal{E} \equiv \times_{t=1}^T \mathcal{E}_t \subset \mathbb{R}^T$ and distribution G , and let $z \equiv \langle z_t : \mathcal{E}^t \times Y^{t-1} \rightarrow \Theta_t \rangle_{t=1}^T$. We say that (G, z) is an independent-shock representation for $F = \langle F_t : \Theta^{t-1} \times Y^{t-1} \rightarrow \Delta(\Theta_t) \rangle_{t=1}^T$ if

(i) for each t , there exists a probability measure G_t on \mathcal{E}_t such that, for any $\varepsilon \in \mathcal{E}$, $G(\varepsilon) = \prod_{t=1}^T G_t(\varepsilon_t)$; and

(ii) for any t and any $\varepsilon^{t-1} \in \mathcal{E}^{t-1}$, the distribution of $z_t(\tilde{\varepsilon}^t; y^{t-1})$ given $\tilde{\varepsilon}^{t-1} = \varepsilon^{t-1}$ is the same as the distribution of θ_t given y^{t-1} and $\theta^{t-1} = z^{t-1}(\varepsilon^{t-1}; y^{t-2}) \equiv (z_1(\varepsilon_1), z_2(\varepsilon^2; y_1), \dots, z_{t-1}(\varepsilon^{t-1}; y^{t-2}))$.

Together, conditions (i) and (ii) simply say that, for any y^{t-1} , one can think of the agent's primitive information θ_t as generated by independent "shocks" $\tilde{\varepsilon}^t$.¹⁵

Example 5 Consider the AR(k) process described in 4. In this example, the functions z_t do not depend on y . They are given by

$$\begin{aligned} z_1(\varepsilon_1) &= \varepsilon_1 \\ z_2(\varepsilon^2) &= \phi_1 \varepsilon_1 + \varepsilon_2 \\ z_3(\varepsilon^3) &= \phi_1(\phi_1 \varepsilon_1 + \varepsilon_2) + \phi_2 \varepsilon_1 + \varepsilon_3 = (\phi_1^2 + \phi_2) \varepsilon_1 + \phi_1 \varepsilon_2 + \varepsilon_3 \\ &\dots \\ z_t(\varepsilon^t) &= \sum_{j=1}^t \left[\sum_{M \in \mathbb{N}, l \in \mathbb{N}^M : j=l_0 < \dots < l_M=t} \prod_{m=1}^M \phi_{l_m - l_{m-1}} \right] \varepsilon_j. \end{aligned}$$

Suppose now that θ is generated by independent shocks ε . Assume further that the agent observes not only θ , but also the shocks ε . Let his payoff (defined over ε and y) be described by the function

$$\begin{aligned} \hat{U}(\varepsilon, y) &\equiv U(z(\varepsilon; y), y) \\ &= U(z_1(\varepsilon_1), z_2(\varepsilon^2; y_1), \dots, z_t(\varepsilon^t; y^{t-1}), \dots, z_T(\varepsilon^T; y^{T-1}), y^T). \end{aligned} \tag{3}$$

Next, consider a (randomized direct) mechanism

$$\hat{\Omega} \equiv \left\langle \hat{\Omega}_t : \mathcal{E}^t \times Y^{t-1} \rightarrow \Delta(Y_t) \right\rangle_{t=1}^T,$$

in which the agent reports the shocks ε instead of his primitive payoff-relevant information θ . The primitive representation (U, F) is equivalent to the representation (\hat{U}, G, Z) in the following sense.

For any $y^{t-1} \in Y^{t-1}$, let $\hat{G}^t(\cdot | z^t(\tilde{\varepsilon}^t; y^{t-1}))$ denote the regular conditional probability distribution of the vector $\tilde{\varepsilon}^t$ given the sigma-algebra $\Sigma(z^t(\tilde{\varepsilon}^t; y^{t-1}))$ generated by the random variable $z^t(\tilde{\varepsilon}^t; y^{t-1})$.¹⁶

¹⁵ A more general definition of an independent-shock representation allows the shocks $\tilde{\varepsilon}$ to depend on the decisions y . However, as will become clear below, for any such representation there is an equivalent one (in the sense of Lemma 2) where the shocks do not depend on y .

¹⁶ Such a regular conditional probability distribution here exists since $\varepsilon^t \in \mathbb{R}^t$. See, e.g., Dudley (2002).

Lemma 2 (a) *Given any ex-ante IC mechanism Ω for the primitive representation (U, F) , there exists an ex-ante IC mechanism $\hat{\Omega}$ for the corresponding independent-shock representation (\hat{U}, G, z) such that, for any t , any measurable $A \subset Y_t$, and any (θ^t, y^{t-1}) ,*

$$\int \hat{\Omega}(A|\varepsilon^t, y^{t-1}) d\hat{G}^t(\varepsilon^t|z^t(\varepsilon^t; y^{t-1}) = \theta^t) = \Omega_t(A|\theta^t, y^{t-1}). \quad (4)$$

(b) *Given any ex-ante IC mechanism $\hat{\Omega}$ for the independent-shock representation (\hat{U}, G, z) , there exists an ex-ante IC mechanism Ω for the primitive representation (U, F) such that, for any t , any measurable $A \subset Y_t$, and any (θ^t, y^{t-1}) , (4) holds.*

Hence any outcome (i.e., any joint distribution over $\Theta \times Y$) that can be sustained by having the agent report θ can also be sustained by having him report the shocks ε , and vice versa. Note that Part (a) follows directly from the fact that if the mechanism Ω is ex-ante IC, then the mechanism $\hat{\Omega}$ defined by

$$\hat{\Omega}(\cdot|\varepsilon^t, y^{t-1}) = \Omega_t(\cdot|z^t(\varepsilon^t; y^{t-1}), y^{t-1}) \quad \forall (\varepsilon^t, y^{t-1}) \quad (5)$$

is also ex-ante IC. This mechanism de facto uses the same information as Ω , in the sense that it depends on ε only through $z(\varepsilon; y)$. Part (b) is also trivially satisfied. It suffices to construct Ω from $\hat{\Omega}$ using the transformation defined in (4). To see that if $\hat{\Omega}$ is ex-ante IC, so is Ω , it suffices to note that (i) payoffs depend on the shocks ε only through $z(\varepsilon; y)$, (ii) Ω induces the same distribution over $\Theta \times Y$ as $\hat{\Omega}$, and (iii) any distribution over $\Theta \times Y$ that the agent can induce given Ω could also have been induced given $\hat{\Omega}$.

Suppose now that an independent-shock representation exists. (We will show below that this is always the case.) One can then use this representation as an alternative route to the characterization of the properties of dynamic incentive-compatible mechanisms. In particular, one can treat the shocks as the agent's private information and then use the formula in Proposition 1 to pin down the (derivative of the) value function. To this aim, let

$$\hat{H} \equiv \{(\varepsilon^s, m^t, y^u) \in \mathcal{E}^s \times \mathcal{E}^t \times Y^u \quad \text{with } s \geq t \geq u \geq s-1\}$$

denote the set of all possible histories in the extensive form corresponding to $\hat{\Omega}$. For any $\hat{h} \in \hat{H}$, let $\hat{\mu}[\hat{\Omega}|\hat{h}]$ denote the (unique) probability measure over $\mathcal{E} \times \mathcal{E} \times Y$ induced by assuming that following history \hat{h} in the mechanism $\hat{\Omega}$, the agent reports truthfully at any subsequent information set. Finally, let $\hat{V}^{\hat{\Omega}}(\hat{h})$ denote the agent's value function in $\hat{\Omega}$ evaluated at history \hat{h} . Now assume each \mathcal{E}_t is an interval, with $\int |\varepsilon_t| dG_t < +\infty$, and that the function \hat{U} is equi-Lipschitz continuous and differentiable in each ε_t . We then have that, if $\hat{\Omega}$ is IC at the truthful history $\hat{h}^{t-1} = (\varepsilon^{t-1}, \varepsilon^{t-1}, y^{t-1})$,

then

$$\hat{V}^{\hat{\Omega}}(\varepsilon_t, \hat{h}^{t-1}) \text{ is Lipschitz continuous in } \varepsilon_t, \text{ and for a.e. } \varepsilon_t, \quad (6)$$

$$\frac{\partial \hat{V}^{\hat{\Omega}}(\varepsilon_t, \hat{h}^{t-1})}{\partial \varepsilon_t} = \mathbb{E}^{\hat{\mu}[\hat{\Omega}]_{\varepsilon_t, \hat{h}^{t-1}}} \left[\frac{\partial \hat{U}(\tilde{\varepsilon}, \tilde{y})}{\partial \varepsilon_t} \right].$$

While this formula can be read as a special case of (IC-FOC), the proof for this result is significantly simpler and follows essentially from the same arguments as in a static setting (see, e.g., Milgrom and Segal, 2002).

Condition (6) thus provides an alternative representation of how the agent's payoff must vary with the agent's private information in a dynamic IC mechanism. In certain applications, working with such a representation may facilitate the characterization of the properties of optimal mechanisms.

At this point one may wonder which F admit an independent-shock representation and which environments (U, F) admit an independent-shock representation for which (6) holds (i.e., for which the corresponding reduced-form payoff $\hat{U}(\varepsilon, y)$ is equi-Lipschitz continuous and differentiable in each ε_t). We address each of these questions in turn.

First, we show that any F admits a particular independent-shock representation, which henceforth we refer to as the *canonical representation*. This representation is derived from F as follows. Let $\tilde{\varepsilon}$ denote a vector of independent random variables, each uniformly distributed over $(0, 1)$. Next, for any t and any $\varepsilon \in (0, 1)$, let

$$F_t^{-1}(\varepsilon | \theta^{t-1}, y^{t-1}) \equiv \inf \{ \theta_t : F_t(\theta_t | \theta^{t-1}, y^{t-1}) \geq \varepsilon \}$$

denote the generalized inverse of the kernel F_t . Then let

$$z_t(\varepsilon^t; y^{t-1}) \equiv F_t^{-1}(\varepsilon_t | F_1^{-1}(\varepsilon_1), F_2^{-1}(\varepsilon_2 | F_1^{-1}(\varepsilon_1), y_1), \dots, y^{t-1}). \quad (7)$$

Applying the *quantile function theorem* recursively (see, e.g., Angus, 1994), one can then show that, given any y^{t-1} and any $\varepsilon^{t-1} \in \mathcal{E}^{t-1} \equiv (0, 1)^{t-1}$, the distribution of $z_t(\tilde{\varepsilon}^t; y^{t-1})$ given $\tilde{\varepsilon}^{t-1} = \varepsilon^{t-1}$ is the same as the distribution of θ_t given y^{t-1} and $\theta^{t-1} = (F_1^{-1}(\varepsilon_1), F_2^{-1}(\varepsilon_2 | F_1^{-1}(\varepsilon_1), y_1), \dots, y^{t-1})$. Hence, any process admits an independent-shock representation in which, for any t , G_t is simply the uniform distribution over $(0, 1)$ and where the functions $z_t : \mathcal{E}^t \times Y^{t-1} \rightarrow \Theta_t$ are as in (7).

Using the canonical representation, we can identify conditions on the primitive environment (U, F) that guarantee that the corresponding payoff \hat{U} in the canonical representation is equi-Lipschitz continuous and differentiable in z_t , for any t .

Assumption 8 *For all $y \in Y$, $U(\cdot, y)$ is equi-Lipschitz and continuously differentiable.*

Assumption 9 For all t , all $\varepsilon \in (0, 1)$, and all $y^{t-1} \in Y^{t-1}$, $F_t^{-1}(\varepsilon|\cdot, y^{t-1})$ is equi-Lipschitz and continuously differentiable.

Assumption 10 For all $(\theta^{t-1}, y^{t-1}) \in \Theta^{t-1} \times Y^{t-1}$, $F_t^{-1}(\cdot|\theta^{t-1}, y^{t-1})$ is equi-Lipschitz and continuously differentiable.

We then have the following result.

Proposition 3 Suppose the primitive description of the environment (U, F) satisfies assumptions 1, 2, 8, 9 and 10. Then, in the corresponding canonical representation, the function \hat{U} obtained from (U, F) using the transformation in (3) is equi-Lipschitz continuous and differentiable in ε . It follows that a mechanism $\hat{\Omega}$ is IC at the truthful history $\hat{h}^{t-1} = (\varepsilon^{t-1}, \varepsilon^{t-1}, y^{t-1})$ only if the value function $\hat{V}^{\hat{\Omega}}(\varepsilon_t, \hat{h}^{t-1})$ satisfies (6).

Proof. That, under the Assumptions 8, 9 and 10, $\partial\hat{U}(\varepsilon, y)/\partial\varepsilon_t$ exists and is continuous in ε_t follows from standard results in calculus (see, e.g., Rudin, 1976, Theorem 9.18.). That $\partial\hat{U}(\varepsilon, y)/\partial\varepsilon_t$ is bounded uniformly over (ε, y) follows from the same conditions along with Assumption 2. It follows that the function \hat{U} is equi-Lipschitz continuous and differentiable in ε . The result then follows directly from this property together with the fact that $\mathcal{E}_t = (0, 1)$ is both connected and bounded. ■

Proposition 3 thus identifies a new set of conditions for the primitive environment (U, F) that guarantee that in any IC mechanism, the agent's expected payoff, when expressed using the canonical representation, satisfies the envelope formula of (6). Comparing the conditions in this proposition with those in Proposition 1, one can see that while the assumptions in Proposition 1 rule out, for example, an atom at $\theta_t = \theta_t^\#$ that “shifts” with the past θ^{t-1} (e.g., fully persistent types), such a possibility is accommodated by the assumptions in Proposition 3. On the other hand, the assumptions in Proposition 3 rule out an atom at $\theta_t = \theta_t^\#$ whose measure grows with θ^{t-1} while such a possibility is allowed by the assumptions in Proposition 1. The assumptions in the two propositions are thus not nested and possibly capture different environments.

To see how the formula in (6) compares to the closed-form one in (2), we proceed as follows. Take any mechanism Ω for the primitive representation (U, F) and let $\hat{\Omega}$ be the mechanism in the corresponding independent-shock representation that is obtained from Ω using (5). Because, for any y , the agent's payoff in $\hat{\Omega}$ depends on ε only through $\theta = z(\varepsilon; y)$, we have that, for any y^{t-1} and any ε^t the following identity holds:

$$\hat{V}^{\hat{\Omega}}(\varepsilon^t, \varepsilon^{t-1}, y^{t-1}) = V^\Omega(z^t(\varepsilon^t; y^{t-1}), z^{t-1}(\varepsilon^{t-1}; y^{t-2}), y^{t-1}). \quad (8)$$

Therefore, at any point of differentiability of $\hat{V}_t^{\hat{\Omega}}$ in ε_t ,

$$\frac{\partial \hat{V}_t^{\hat{\Omega}}(\varepsilon^t, \varepsilon^{t-1}, y^{t-1})}{\partial \varepsilon_t} = \left(\frac{\partial V^{\Omega}(z^t(\varepsilon^t; y^{t-1}), z^{t-1}(\varepsilon^{t-1}; y^{t-2}), y^{t-1})}{\partial \theta_t} \right) \left(\frac{\partial z_t(\varepsilon^t; y^{t-1})}{\partial \varepsilon_t} \right). \quad (9)$$

Using (9) and (3), the formula in (6) can then be rewritten as

$$\begin{aligned} & \left(\frac{\partial V^{\Omega}(z^t(\varepsilon^t; y^{t-1}), z^{t-1}(\varepsilon^{t-1}; y^{t-2}), y^{t-1})}{\partial \theta_t} \right) \left(\frac{\partial z_t(\varepsilon^t; y^{t-1})}{\partial \varepsilon_t} \right) \\ &= \mathbb{E}^{\hat{\mu}[\hat{\Omega}]|\varepsilon_t, \hat{h}^{t-1}} \left[\sum_{s=t}^T \frac{\partial U(z^T(\tilde{\varepsilon}^T; \tilde{y}^{T-1}), \tilde{y}^T)}{\partial \theta_s} \frac{\partial z_s(\tilde{\varepsilon}^s; \tilde{y}^{s-1})}{\partial \varepsilon_t} \right]. \end{aligned} \quad (10)$$

Note that this formula applies to all independent-shock representations of F , not only to the canonical one. In the special case in which (G, z) is the canonical representation we have

$$\frac{\partial z_t(\varepsilon^t; y^{t-1})}{\partial \varepsilon_t} = \frac{\partial F_t^{-1}(\varepsilon_t | F_1^{-1}(\varepsilon_1), F_2^{-1}(\varepsilon_2 | F_1^{-1}(\varepsilon_1), y_1), \dots, y^{t-1})}{\partial \varepsilon_t},$$

whereas for any $s > t$,

$$\begin{aligned} \frac{\partial z_s(\varepsilon^s; y^{s-1})}{\partial \varepsilon_t} &= \sum_{j=t}^{s-1} A_j^s(\varepsilon^s, y^{s-1}) \frac{\partial z_j(\varepsilon^j; y^{j-1})}{\partial \varepsilon_t} \\ &= \frac{\partial z_t(\varepsilon^t; y^{t-1})}{\partial \varepsilon_t} [A_t^s(\varepsilon^s, y^{s-1}) + A_{t+1}^s(\varepsilon^s, y^{s-1}) A_t^{t+1}(\varepsilon^{t+1}, y^t) \\ &\quad + A_{t+2}^s(\varepsilon^s, y^{s-1}) (A_t^{t+2}(\varepsilon^{t+2}, y^{t+1}) + A_t^{t+1}(\varepsilon^{t+1}, y^t) A_{t+1}^{t+2}(\varepsilon^{t+2}, y^{t+1})) + \dots], \end{aligned} \quad (11)$$

with

$$A_j^s(\varepsilon^s, y^{s-1}) \equiv \frac{\partial F_s^{-1}(\varepsilon_s | F_1^{-1}(\varepsilon_1), F_2^{-1}(\varepsilon_2 | F_1^{-1}(\varepsilon_1), y_1), \dots, y^{s-1})}{\partial \theta_j}.$$

One can then show that if $F_t(\theta_t | \theta^{t-1}, y^{t-1})$ is strictly increasing and absolutely continuous in θ_t and differentiable in each θ_s , $s \leq t$, then

$$A_j^s(\varepsilon^s, y^{s-1}) = I_j^s(\theta_s | \theta^{s-1}, y^{s-1}) \Big|_{\theta^s = z^s(\varepsilon^s, y^{s-1})},$$

and

$$\frac{\frac{\partial z_s(\varepsilon^s; y^{s-1})}{\partial \varepsilon_t}}{\frac{\partial z_t(\varepsilon^t; y^{t-1})}{\partial \varepsilon_t}} = J_t^s(z^s(\varepsilon^s; y^{s-1}), y^{s-1}),$$

so that (6) coincides with (2)—the details are in the Appendix. We then have the following result.

Proposition 4 *Suppose the primitive environment (U, F) satisfies assumptions 1, 2, 3, 7, 8, 9,*

and 10. Then the conclusions of Proposition 2 hold.

Note that assumptions 1, 2, 3 and 7 are also present in Proposition 2. Assumption 8 is stronger than assumption 7 (it implies the latter). On the other hand, assumptions 5 and 6 are not implied by assumptions 9 and 10. The two propositions thus identify different sets of necessary conditions for the validity of the closed-form formula given in (2).

3.4 Sufficient conditions for IC

While formula (2) summarizes local (first-order) incentive constraints, it does not imply the satisfaction of all (global) incentive constraints. In this section we formulate some sufficient conditions for incentive compatibility. These conditions generalize the well-known monotonicity condition, which together with the first-order condition characterizes incentive-compatible mechanisms in the static model with a one-dimensional type space. The static characterization cannot be extended to the dynamic model, which could be viewed as an instance of multidimensional mechanism design problem, for which the characterization of IC mechanisms is more difficult (see, e.g., Rochet and Stole, 2003). More precisely, there are two sources of difficulty in ensuring incentive compatibility of a dynamic mechanism: (a) in general one needs to consider multiperiod deviations, since once the agent lies in one period, his optimal continuation strategy may require lying in subsequent periods as well;¹⁷ and (b) even if one focuses on single-period deviations, in which the agent misrepresents his current one-dimensional type, the decisions assigned by the mechanism from that period onward form a multidimensional decision space.

While these problems make it hard to have a general *characterization* of incentive compatibility, we can still formulate sufficient conditions for IC that prove useful in a number of applications. Problem (a) is sidestepped by focusing on environments in which we can assure that truth-telling is an optimal continuation strategy even following deviations, and so incentive compatibility can be assured by checking one-period deviations. (While this focus is quite restrictive, it includes all Markov environments, as well as some other interesting cases—see for example the application to sequential auctions with AR(k) values considered in subsection 5.2. Problem (b) is sidestepped by formulating a monotonicity condition that, while not necessary for IC, is sufficient and is easy to check in applications.

¹⁷It is possible to ensure that truth-telling is optimal even after deviations by allowing the agent to re-report his complete history θ^t in each period t , possibly contradicting his earlier reports. This is the version of the revelation principle proposed by Doepke and Townsend (2006). While this approach would allow us to restrict attention to one-stage deviations from truth-telling, the deviations in each period would now be multidimensional, and contingent on possibly inconsistent reporting histories, so it is not clear that this approach would simplify formulation of sufficient conditions.

Proposition 5 *Suppose the environment satisfies either the assumptions of Proposition 2 or those of Proposition 4. Fix any period t and for any period- t history h , let*

$$D^\Omega(h) \equiv \mathbb{E}^{\mu[\Omega]|h} \left[\sum_{\tau=t}^T J_t^\tau(\tilde{\theta}^\tau, \tilde{y}^{\tau-1}) \frac{\partial U(\tilde{\theta}, \tilde{y})}{\partial \theta_\tau} \right].$$

Suppose that for any truthful history $(\theta^{t-1}, \theta^{t-1}, y^{t-1})$,

(i) $\mathbb{E}^{\mu[\Omega]|((\theta^{t-1}, \theta_t), \theta^{t-1}, y^{t-1})} [U(\tilde{\theta}, \tilde{y})]$ is Lipschitz continuous in θ_t , and for a.e. θ_t ,

$$\frac{d}{d\theta_t} \mathbb{E}^{\mu[\Omega]|((\theta^{t-1}, \theta_t), \theta^{t-1}, y^{t-1})} [U(\tilde{\theta}, \tilde{y})] = D^\Omega((\theta^{t-1}, \theta_t), (\theta^{t-1}, \theta_t), y^{t-1}).$$

(ii) For any m_t , for a.e. θ_t ,

$$[D^\Omega((\theta^{t-1}, \theta_t), (\theta^{t-1}, \theta_t), y^{t-1}) - D^\Omega((\theta^{t-1}, \theta_t), (\theta^{t-1}, m_t), y^{t-1})] \cdot (\theta_t - m_t) \geq 0,$$

(iii) Ω is IC at any (possibly non-truthful) period $t+1$ history.

Then Ω is IC at any truthful period- t history.

Propositions 2 and 4 imply that condition (i) in Proposition 5 is a necessary condition for the mechanism to be IC at *all* truthful period- t histories (Recall that this means that the agent's value function at these histories coincides with the expected equilibrium payoff). The addition of conditions (ii) and (iii) is then sufficient (but in general not necessary) for IC at all truthful period- t histories—The proof is based on a lemma in the appendix that extends to a dynamic setting a result by Garcia, 2005 for static mechanism design with one-dimensional type and multidimensional decisions.

The assumption that the mechanism is IC at all period $t+1$ histories, including non-truthful ones, is rather strong, but it can be satisfied in some applications. As one prominent example, in a Markov setting, the history θ^t of the agent's true types does not affect his incentives in period $t+1$ after θ_{t+1} is observed. Thus, any mechanism that is IC at all truthful period $t+1$ histories must also be IC at *all* period $t+1$ histories. In this case, the Proposition can be iterated starting from period $T+1$ moving backward to establish IC in all periods and at all histories.

4 Multi-agent quasilinear case

We now introduce multiple agents. The multi-agent model we consider features three important assumptions: (1) the environment is quasilinear (i.e., the decision taken in each period can be decomposed into an allocation and a vector of monetary payments and the agents' preferences are quasilinear in the payments), (2) the type distributions are independent of past monetary payments

(but they may still depend on past allocations), and (3) types are independent across agents. After setting up the model we show how from the perspective of an individual agent the model reduces to the single-agent case studied in the previous section.

4.1 Quasilinear environment

There are N agents indexed by $i = 1, \dots, N$. In each period $t = 1, \dots, T$, each agent i is shown a nonmonetary “allocation” decision $x_{it} \in X_{it}$ (where X_{it} is a measurable space), and given a payment $p_{it} \in \mathbb{R}$. The set of feasible joint allocation decisions in period t is $X_t \subset \prod_{i=1}^N X_{it}$.^{18,19}

Each agent i observes his own allocations x_{it} but not the others’ allocations $x_{-i,t}$. The observability of x_{it} should be thought of as a technological restriction: A mechanism can reveal more information to agent i in period t than x_{it} , but it cannot conceal x_{it} . As for the payments, because none of the results hinges on the specific information the agents receive about p , we leave the description of the information the agents receive about p unspecified.

As in the single-agent case, histories are denoted using the superscript notation. For example, (x^t, p^t) is an element of $X^t \times \mathbb{R}^{Nt}$, where $X^t \subset \prod_{\tau=1}^t X_\tau$ and $X \subset \prod_{\tau=1}^T X_\tau$.

In each period t , each agent i privately observes his current type $\theta_{it} \in \Theta_{it} \subset \mathbb{R}$. The current type profile is then denoted by $\theta_t \equiv (\theta_{1t}, \dots, \theta_{Nt}) \in \Theta_t \equiv \prod_i \Theta_{it}$. The distribution of the type profile $\theta \in \Theta \equiv \prod_{t=1}^T \Theta_t$ is described in the following definition.

We omit superscripts for full histories, with the exception of $x_i^T \equiv (x_{i1}, \dots, x_{iT})$, $p_i^T \equiv (p_{i1}, \dots, p_{iT})$, and $\theta_i^T \equiv (\theta_{i1}, \dots, \theta_{iT})$ (and the sets they are elements of). This is to avoid confusion between, e.g., $x_t \equiv (x_{1t}, \dots, x_{Nt})$ and $x_i \equiv (x_{i1}, \dots, x_{iT})$.

Agent i ’s payoff function is denoted $U_i : \Theta \times X \times \mathbb{R}^T \rightarrow \mathbb{R}$.

We then define a quasi-linear environment as follows.

Definition 8 *The environment is quasilinear if the following hold:*

1. *There is a sequence of decisions $(x, p) \in X \times \mathbb{R}^{NT}$, where $x = (x_1^T, \dots, x_N^T)$ is an allocation, p is a vector of payments, and for all i , x_i^T is the minimal information about x received by agent i .*
2. *The distribution of the type profile θ is governed by kernels $\langle F_t : \Theta^{t-1} \times X^{t-1} \rightarrow \Delta(\Theta_t) \rangle_{t=1}^T$.*

¹⁸For example, we can have $X_t = \left\{ x_t \in \mathbb{R}_+^N : \sum_i x_{it} \leq \bar{x}_t \right\}$ when the decision is the allocation of a private good among agents, or $X_t = \{ x_t \in \mathbb{R}_+^N : x_{1t} = x_{2t} = \dots = x_{Nt} \}$ when the decision is the provision of a public good.

¹⁹This formulation does not explicitly allow for decisions that are not observed by any agent at the time they are made, but such decisions could still be modeled if need be by creating a fictitious agent observing them.

3. For all i , the payoff function of each agent i , $U_i : \Theta \times X \times \mathbb{R}^T \rightarrow \mathbb{R}$, takes the quasilinear form

$$U_i(\theta, x, p_i^T) = u_i(\theta, x) - \sum_{t=1}^T p_{it}$$

for some measurable $u_i : \Theta \times X \rightarrow \mathbb{R}$.

Note that part 2 restricts the distribution of θ to be independent of the payments. As for part 3, note that for the sake of generality we allow agent i 's utility to depend on things he does not observe, namely x_{-i}^T and θ_{-i}^T .²⁰

Definition 9 Types are Independent if for all t , and all $(\theta^{t-1}, x^{t-1}) \in \Theta^{t-1} \times X^{t-1}$,

$$F_t(\cdot | \theta^{t-1}, x^{t-1}) = \prod_{i=1}^N F_{it}(\cdot | \theta_i^{t-1}, x_i^{t-1}),$$

where for all i , $F_{it}(\cdot | \theta_i^{t-1}, x_i^{t-1})$ is a probability measure on Θ_{it} .

This definition is the proper extension of the Independent-Type assumptions to dynamic settings, allowing us to extend such static results as revenue-equivalence and virtual surplus representation of expected profits. Note that the definition can be broken up into three parts: (i) Conditional on any history (θ^{t-1}, x^{t-1}) , period- t types are independent across agents. (ii) The distribution of agent i 's period- t type, $F_{it}(\cdot | \theta_i^{t-1}, x_i^{t-1})$, does not depend on the other agents' past types (except possibly indirectly through the decision history x_i^{t-1} observed by agent i). (iii) The distribution $F_{it}(\cdot | \theta_i^{t-1}, x_i^{t-1})$ also does not depend on the history of decisions x_{-i}^{t-1} that the agent has not observed. It is easy to see that if the assumptions (i) or (ii) are not satisfied, then a mechanism similar to the one proposed by Cremer and McLean (1988) could be used to extract the agents' information rents. It turns out that a similar extraction of rents is possible if assumption (iii) is not satisfied by using a randomized mechanism—see the discussion after Proposition 6 below.

4.2 Multi-agent mechanisms

For most of the analysis we will focus on the Bayesian Nash Equilibria (BNE) of mechanisms designed for the environment described above. As discussed for the single-agent case, this solution concept imposes the weakest form of rationality on the agents' behavior and thus yields the

²⁰Some readers may feel that an agent must always be able to observe his own final payoff (indeed, this was the case in our model in Section 3). This can still be compatible with an interdependent-value model in which agent i observes x_{-i}^T and θ_{-i}^T at the end of period T and is unable to report them to the mechanism. If we instead allowed the agent to report his observed final payoff in an interdependent-value model to the mechanism, as in Mezzetti (2004), we would effectively convert the model to one with correlated private observations, allowing for full surplus extraction.

strongest necessary conditions for incentive compatibility. The sufficient conditions we offer, will however ensure implementation with a stronger solution concept such as (weak) Perfect Bayesian Equilibrium.

By the revelation principle (adapted from Myerson, 1986), it is without loss of generality to restrict attention to Bayesian incentive compatible “direct mechanisms” (defined more precisely below) where (1) in each period each agent confidentially reports his current type θ_{it} to the mechanism, and (2) the mechanism reports no information back to the agents (i.e., each agent i observes only (θ_i^T, x_i^T) plus whatever is assumed observable about the payments).²¹ The proof for (1) is the familiar one. As for (2), suppose there exists an incentive-compatible direct mechanism where more information is revealed to the agents than what described in (2). Concealing this additional information would then amount to pooling different incentive-compatibility constraints resulting in a new IC mechanism that implements the same outcomes (i.e., the same distribution over $\Theta \times X \times \mathbb{R}^{NT}$).

When exploring the implications of incentive compatibility, it is also convenient to assume that all payments take place at the very end. This is actually without loss of generality. In fact, because postponing payments amounts to hiding information, for any IC mechanism in which some payments are made (and possibly observed) in each period, there exists another IC mechanism in which all payments are postponed to the end which induces the same distribution over $\Theta \times X$ and, for all θ , it induces the same total payments.

For notational simplicity hereafter we restrict attention to deterministic mechanisms. This entails no loss since randomizations could always be generated by introducing a fictitious agent whose type is publicly observed. We will also formulate sufficient conditions under which such randomizations will not be useful.

Definition 10 *A deterministic direct mechanism is a pair $\langle \chi, \psi \rangle$, where $\chi = \langle \chi_t : \Theta^t \rightarrow X_t \rangle_{t=1}^T$ is an allocation rule, and $\psi : \Theta \rightarrow \mathbb{R}^N$ is a (total) payment scheme.*

A deterministic direct mechanism $\langle \chi, \psi \rangle$ defines the following sequence in each period t , following a history θ^{t-1} of type observations and a history $m^{t-1} = (m_1^{t-1}, \dots, m_N^{t-1})$ of type reports by the agents:

1. Each agent i privately observes his current type $\theta_{it} \in \Theta_{it}$ drawn from $F_{it}(\cdot | \theta_i^{t-1}, \chi_i^{t-1}(m^{t-1}))$.
2. Each agent i sends a confidential message $m_{it} \in \Theta_{it}$ to the mechanism.
3. The mechanism implements the decision $\chi_t(m^t)$.
4. Each agent i observes $\chi_{it}(m^t)$.

²¹In our environment there are no actions to be privately chosen by the agents. If the agents were also to choose hidden actions, then a direct mechanism would also send the agents recommendations for the hidden actions.

After period T , the mechanism also implements the payments $\psi(m^T)$.

A mechanism induces an extensive form game between the agents. A (pure) strategy for agent i is a complete contingent plan

$$\sigma_i \equiv \langle \sigma_{it} : \Theta_i^t \times \Theta_i^{t-1} \times X_i^{t-1} \rightarrow \Theta_{it} \rangle_{t=1}^T.$$

Truthful strategies are defined as in the single-agent case.

If all agents play truthful strategies, a deterministic allocation rule χ induces a stochastic process on the agents' types Θ described by the kernels $F_t(\cdot | \theta^{t-1}, \chi^{t-1}(\theta^{t-1}))$. We let $\lambda[\chi]$ denote the resulting probability measure on Θ . Similarly, if all agents but i are playing truthful strategies, while agent i follows a strategy σ_i , this induces a stochastic process on $(\theta, m_i^T) \in \Theta \times \Theta_i^T$, which is described by the kernels F , allocation rule χ , and strategy σ_i . We let $\lambda_i[\chi, \sigma_i]$ denote the resulting probability measure on $\Theta \times \Theta_i^T$. Equipped with this notation, we can define ex-ante incentive compatibility of a mechanism as follows.

Definition 11 *A deterministic direct mechanism $\langle \chi, \psi \rangle$ is ex-ante Bayesian Incentive Compatible (ex-ante BIC) if for all i and all σ_i ,*

$$\mathbb{E}^{\lambda[\chi]}[u_i(\tilde{\theta}, \chi(\tilde{\theta})) - \psi_i(\tilde{\theta})] \geq \mathbb{E}^{\lambda_i[\chi, \sigma_i]}[u_i(\tilde{\theta}, \chi(\tilde{m}_i^T, \tilde{\theta}_{-i}^T)) - \psi_i(\tilde{m}_i^T, \tilde{\theta}_{-i}^T)].$$

That is, a mechanism is ex-ante BIC if the truthful strategies form a Bayesian Nash Equilibrium of the game induced by the mechanism.

4.3 Mapping the multi-agent into the single-agent case

We now show that, from the perspective of each agent, the environment described above can be mapped into the single-agent model considered in Section 3. To see this, fix an arbitrary agent i . Given any deterministic mechanism $\langle \chi, \psi \rangle$, when agents $-i$ are playing truthful strategies, agent i effectively faces a randomized mechanism where the randomizations are due to the uncertainty that agent i faces about the other agents' types. Over the course of the mechanism, agent i only observes $(\theta_i^T, m_i^T, x_i^T)$. However, the mechanism depends on the other agents' types θ_{-i}^T through their equilibrium messages; furthermore, agent i 's utility may depend directly on θ_{-i}^T and x_{-i}^T . Thus evaluating the optimality of i 's strategy requires keeping track of his beliefs about θ_{-i}^T conditional on the observed history.

Formally the problem faced by agent i can be mapped into the single-agent model considered in the previous section as follows. For all $t < T$, let $Y_{it} = X_{it}$, and let $Y_{iT} = X_{iT} \times X_{-i}^T \times \Theta_{-i}^T$. Also, let $Y_{i,T+1} = \mathbb{R}^T$. That is, in periods $t < T$ the decision $y_{it} = x_{it}$ consists of the part of the allocation observed by agent i . In period T , the decision y_{iT} also shows the agent the rest of the

variables affecting his utility (i.e., the part of the allocation x_{-i}^T unobservable to him and the other agents' types θ_{-i}^T). Then in period $T + 1$, which is introduced just as a convenient modelling device, the agent observes his payment p_i^T .

Agent i 's type θ_i^T evolves according to the kernels $F_i = \langle F_{it} : \Theta_i^{t-1} \times X_i^{t-1} \rightarrow \Delta(\Theta_{it}) \rangle_{t=1}^T = \langle F_{it} : \Theta_i^{t-1} \times Y_i^{t-1} \rightarrow \Delta(\Theta_{it}) \rangle_{t=1}^T$, where the equality is by definition of Y_{it} . There is no type in period $T + 1$ (formally, $\Theta_{i,T+1}$ can be taken to be an arbitrary singleton).

In the single-agent setup, agent i 's payoff is defined over $\Theta_i^T \times Y_i^{T+1}$, where $Y_i = \prod_{t=1}^{T+1} Y_{it}$. However, by construction $\Theta_i^T \times Y_i^{T+1}$ is simply a reordering of $\Theta \times X \times \mathbb{R}^T$ —the domain of agent i 's payoff in the multi-agent environment. To highlight this connection, we abuse notation and continue to use U_i with its arguments appropriately reordered.

Agent i faces a randomized mechanism $\Omega_i = \Omega_i[\chi, \psi] \equiv \langle \Omega_{it} : \Theta_i^t \times Y_i^{t-1} \rightarrow \Delta(Y_{it}) \rangle_{t=1}^{T+1}$ constructed as follows. We first construct inductively a consistent family of regular conditional probability distributions (rcpd) that represent the evolution of agent i 's beliefs about θ_{-i}^T conditional on observable allocations and his own messages.²² Fix $t < T$. Suppose that a rcpd $\Gamma_{\tau-1}(\cdot | \chi_i^{\tau-1}(m_i^{\tau-1}, \tilde{\theta}_{-i}^{\tau-1}))$ on $\Theta_{-i}^{\tau-1}$ exists for all $m_i^{\tau-1}$, and all periods $\tau \leq t$. (The conditioning here is on the random variable $\chi_i^{\tau-1}(m_i^{\tau-1}, \tilde{\theta}_{-i}^{\tau-1})$ taking values in $Y_i^{\tau-1}$.) Note that the assumption holds vacuously for $t = 1$. For all m_i^t , the rcpd $\Gamma_{t-1}(\cdot | \chi_i^{t-1}(m_i^{t-1}, \tilde{\theta}_{-i}^{t-1}))$ and the kernels $F_{-i,t}(\cdot | \theta_{-i}^{t-1}, \chi_{-i}^{t-1}(m_i^{t-1}, \theta_{-i}^{t-1}))$ induce a probability measure on Θ_{-i}^t . Since $\Theta_{-i}^t \subset \mathbb{R}^{N-1}$, there exists a rcpd of $\tilde{\theta}_{-i}^t$ given $\Sigma(\chi_i^t(m_i^t, \tilde{\theta}_{-i}^t))$, where $\Sigma(\chi_i^t(m_i^t, \tilde{\theta}_{-i}^t))$ denotes the sigma-algebra generated by the random variable $\chi_i^t(m_i^t, \tilde{\theta}_{-i}^t)$ (see, e.g., Theorem 10.2.2 in Dudley, 2002). We define $\Gamma_t(\cdot | \chi_i^t(m_i^t, \tilde{\theta}_{-i}^t))$ to be this rcpd. Consistency of the family follows by construction. At $t = T$ the decision y_{iT} reveals to the agent θ_{-i}^T , and hence his beliefs are degenerate in periods T and $T + 1$.

Let $t < T$ and fix a history (m_i^t, y_i^{t-1}) . Then for any measurable $A \subset Y_{it}$, the probability that $y_{it} \in A$ is

$$\Omega_{it}(A | m_i^t, y_i^{t-1}) \equiv \int_{\{\theta_{-i}^t \in \Theta_{-i}^t : \chi_{it}(m_i^t, \theta_{-i}^t) \in A\}} dF_{-i,t}(\theta_{-i,t} | \theta_{-i}^{t-1}, \chi_{-i}^{t-1}(m_i^{t-1}, \theta_{-i}^{t-1})) d\Gamma_{t-1}(\theta_{-i}^{t-1} | \chi_i^{t-1}(m_i^{t-1}, \tilde{\theta}_{-i}^{t-1}) = y_i^{t-1}).$$

The measure $\Omega_{it}(\cdot | m_i^T, y_i^{T-1})$ is defined analogously except that the integral is over the set

$$\{\theta_{-i}^T \in \Theta_{-i}^T : (\chi_i(m_i^T, \theta_{-i}^T), \chi_{-i}^T(m_i^T, \theta_{-i}^T), \theta_{-i}^T) \in A\}.$$

Finally, $\Omega_{i,T+1}(\cdot | m_i^T, (x_i^T, x_{-i}^T, \theta_{-i}^T))$ is defined to be a point mass at $\psi(m_i^T, \theta_{-i}^T)$. This defines the randomized direct mechanism $\Omega_i = \Omega_i[\chi, \psi]$.

²²See, e.g., Dudley (2002) for the definition of a regular conditional probability distributions.

Thus, from the perspective of agent i , there is a decision y_{it} in each period t , his type θ_{it} evolves according to kernels F_i , utility is given by U_i , and he is facing a randomized direct mechanism Ω_i . This is the setup considered in the single-agent part. In particular, let $H_i \equiv \{(\theta_i^s, m_i^t, y_i^u) : s \geq t \geq u \geq s - 1\}$ denote the set of the agent's private histories. Then a strategy σ_i and a private history $h_i \in H_i$ induce a probability measure $\mu_i[\Omega_i, \sigma_i] | h_i$ on $\Theta_i^T \times \Theta_i^T \times Y_i$. Since Ω_i is derived from the multi-agent mechanism $\langle \chi, \psi \rangle$, we abuse notation and write $\mu_i[\langle \chi, \psi \rangle, \sigma_i] | h_i$ to emphasize the connection to the original mechanism. For the truthful strategy and the null history the measure is then denoted $\mu_i[\chi, \psi] | h_i$ and $\mu_i[\langle \chi, \psi \rangle, \sigma_i]$, respectively. The agent's payoff from truthtelling following history h_i is thus $\mathbb{E}^{\mu_i[\chi, \psi] | h_i}[U_i(\tilde{\theta}_i, \tilde{y}_i)] = \mathbb{E}^{\mu_i[\chi, \psi] | h_i}[U_i(\tilde{\theta}, \tilde{x}, \tilde{p}_i^T)]$, where the equality is by definition of y_i . We can then define the value function $V_i^{\Omega_i[\chi, \psi]} : H_i \rightarrow \mathbb{R}$ and incentive compatibility at a private history h_i analogously to the single-agent definitions.

It will be convenient to let $\mu_i^T[\chi] | h_i$ denote the marginal of $\mu_i[\chi, \psi] | h_i$ on $\Theta_i^T \times \Theta_i^T \times Y_i^T$ given private history h_i . Thus, $\mu_i^T[\chi] | h_i$ is a process on types, messages, and nonmonetary allocations, but not on the payments (which by our convention are only made in period $T + 1$). The role of this notation is to highlight the fact that the stochastic process over everything but the payments in the quasilinear environment is determined by the allocation rule χ and independently of the payment rule ψ . Since the payment scheme ψ is a deterministic function of the messages (which under $\mu_i^T[\chi] | h_i$ are truthful), we can use $\mu_i^T[\chi] | h_i$ to write agent i 's payoff as $\mathbb{E}^{\mu_i^T[\chi] | h_i}[u_i(\tilde{\theta}, \tilde{x}) + \psi_i(\tilde{\theta})]$.

4.4 Revenue equivalence

Suppose the assumptions in Proposition 1, or alternatively those in Proposition 3 along with assumption 3, hold for any i . We then have that in any mechanism $\langle \chi, \psi \rangle$ that is IC for agent i at a truthful private history $h_i^{t-1} = (\theta_i^{t-1}, \theta_i^{t-1}, x_i^{t-1})$, $V^{\Omega_i[\chi, \psi]}(\theta_{it}, h_i^{t-1})$ is equi-Lipschitz continuous in θ_{it} . Furthermore, under quasi-linearity, the derivative of $V^{\Omega_i[\chi, \psi]}(\theta_{it}, h_i^{t-1})$ with respect to θ_{it} does not depend on the payment scheme. Under the assumptions of Proposition 1, this can be seen by iterating (IC-FOC) backward starting from $t = T$. Under the assumptions of Proposition 3 this can be seen directly from (6).

In a quasi-linear environment, the aforementioned propositions thus imply that in each ex-ante BIC mechanism, the value function of each agent i at any truthful private history $h_i^{t-1} = (\theta_i^{t-1}, \theta_i^{t-1}, x_i^{t-1})$ is pinned down by the allocation rule χ up to a constant $k_i(h_i^{t-1})$ that may depend on h_i^{t-1} , but not on θ_{it} .

Using the law of iterated expectations, one can then also get rid of the dependence of this constant on the history. To see this more clearly, suppose there is a single agent i and assume, for simplicity, that there are only two periods. Now consider any two ex-ante IC deterministic mechanisms $\langle \chi, \psi \rangle$ and $\langle \chi, \hat{\psi} \rangle$ implementing the same allocation rule χ . Then in period two, for

any truthful history $h_{i1} = (\theta_{i1}, \theta_{i1}, \chi(\theta_{i1}))$, there exists a scalar $k_i(h_{i1}) = K_i(\theta_{i1})$ such that, for any θ_{i2} , $V^{\Omega_i[\chi, \psi]}(\theta_{i2}, h_{i1}) - V^{\Omega_i[\chi, \hat{\psi}]}(\theta_{i2}, h_{i1}) = K_i(\theta_{i1})$. A similar result also applies to period one: there exists a scalar K_i such that, for each θ_{i1} , $V^{\Omega_i[\chi, \psi]}(\theta_{i1}) - V^{\Omega_i[\chi, \hat{\psi}]}(\theta_{i1}) = K_i$. Because $V^{\Omega_i[\chi, \psi]}(\theta_{i1}) - V^{\Omega_i[\chi, \hat{\psi}]}(\theta_{i1})$ is simply the expectation of $V^{\Omega_i[\chi, \psi]}(\theta_{i2}, h_{i1}) - V^{\Omega_i[\chi, \hat{\psi}]}(\theta_{i2}, h_{i1})$, we then have that $K_i(\theta_{i1}) = K_i$ for all θ_{i1} .

Clearly, the same result extends to any T . Furthermore, because the value function coincides with the equilibrium payoff with probability one and because the latter is simply the difference between the expectation of $u(\tilde{\theta}_i^T, \chi(\tilde{\theta}_i^T))$ and the expectation of $\psi(\tilde{\theta}_i^T)$, we have that the entire payment scheme ψ is uniquely determined by the allocation rule χ up to a scalar.

This result extends to a setting with multiple agents, provided that types are independent: The total payment that each agent expects to make to the mechanism as a function his period-one type is uniquely determined by the allocation rule χ up to a scalar K_i that does not depend on θ_{i1} . This is the famous "*revenue equivalence*" result extensively documented in static environments. More generally, one can show that the same result extends to any arbitrary period $t \geq 1$ provided that the following condition holds.

Assumption 11 (DNOT) *Decisions do Not Affect Types: For all $i = 1, \dots, N$ and all $t = 2, \dots, T$, $F_{it}(\theta_{it} | \theta_i^{t-1}, x_i^{t-1})$ do not depend on x_i^{t-1} .*

We then have the following result.

Proposition 6 *Suppose types are independent and that the environment satisfies for all i either the assumptions of Proposition 1, or those of Proposition 3 along with assumption (3). Consider any two ex-ante BIC deterministic mechanisms $\langle \chi, \psi \rangle$ and $\langle \chi, \hat{\psi} \rangle$ implementing the same allocation rule χ . Then for all i , there exists a $K_i \in \mathbb{R}$ such that²³*

$$\mathbb{E}^{\lambda[\chi]}[\psi_i(\tilde{\theta}) | \tilde{\theta}_{i1}] - \mathbb{E}^{\lambda[\chi]}[\hat{\psi}_i(\tilde{\theta}) | \tilde{\theta}_{i1}] = K_i. \quad (12)$$

If, in addition, assumption DNOT holds (with $N = 1$, assumption DNOT can be dispensed with), then, for all i and any t, s ,

$$\mathbb{E}^{\lambda}[\psi_i(\tilde{\theta}) | \tilde{\theta}_i^t] - \mathbb{E}^{\lambda}[\hat{\psi}_i(\tilde{\theta}) | \tilde{\theta}_i^t] = \mathbb{E}^{\lambda}[\psi_i(\tilde{\theta}) | \tilde{\theta}_i^s] - \mathbb{E}^{\lambda}[\hat{\psi}_i(\tilde{\theta}) | \tilde{\theta}_i^s]. \quad (13)$$

The value of Proposition 6 is twofold: (a) it sheds light on certain real-world institutions (for example it, can be used to establish revenue-equivalence across different dynamic auctions formats); (b) it facilitates the characterization of profit-maximizing mechanisms by permitting one to express

²³Given a mechanism $\langle \chi, \psi \rangle$, $\mathbb{E}^{\lambda[\chi]}[\psi_i(\tilde{\theta}) | \tilde{\theta}_i^t]$ denotes the expectation of $\psi_i(\tilde{\theta})$ conditional on the random variable $\tilde{\theta}_i^t$, where, as usual, conditional expectations are interpreted as Radon-Nikodym derivatives.

the principal’s expected payoff as expected virtual surplus, as illustrated below. Both (a) and (b) use the result of Proposition 6 only for $t = 1$. However, the property that, when decisions do not affect types, the difference in expected payments remains constant over time in the sense of condition (13) also turns useful in certain applications.

Note also that the result in Proposition 6 can be sharpened by considering a stronger solution concept. Suppose one is interested in mechanisms with the property that each agent finds it optimal to report truthfully even after being shown at the beginning of period one the entire profile of the other agents’ types θ_{-i}^T (which includes current and future types). Then a simple backward induction argument similar to the one used to establish Proposition 6 implies that, for each agent, payments are uniquely determined up to a scalar not only in expectation but for each state $(\theta_i^T, \theta_{-i}^T)$. (We provide sufficient conditions for the resulting mechanism to indeed satisfy this robustness to information leakage in Corollary 1 below.)

Lastly, note that a key assumption in Proposition 6 is that types are *independent*. As mentioned above, this assumption has two parts: First, it requires that, given (θ^{t-1}, x^{t-1}) , current types are independent across agents; Second it requires that the distribution of each agent i ’s current type θ_{it} depends only on objects observable to agent i , that is, on $(\theta_i^{t-1}, x_i^{t-1})$. The importance of the first part for revenue equivalence is well understood. The arguments are the same as in static environments (see, e.g., Cremer and McLean, 1988). The importance of the second part may be less obvious. To see it, suppose for simplicity there are only two periods and assume that the distribution of θ_{i2} depends not only on θ_{i1}, x_{i1} but also on a variable $x_{-i,1}$ that is not directly observed by agent i but which is observed by the principal (or by whoever runs the mechanism). Depending on the application, one may think of $x_{-i,1}$ as the amount of R&D commissioned to a research lab (the principal) by competitive clients (the other agents). Alternatively, one may think of $x_{-i,1}$ as the unobservable quality of a product supplied by the principal to buyer i . If $x_{-i,1}$ is known to the principal but not to agent i and if it is correlated with θ_{i2} , then the principal can extract all the private information that agent i possesses about θ_{i2} for free (the arguments here are once again the same as in the case of correlated types). This clearly precludes revenue equivalence.

4.5 Dynamic virtual surplus and profit-maximizing mechanisms

In a static setting, the envelope formula allows to calculate the agents’ information rents, providing a tool for designing profit-maximizing mechanisms. We show here how this approach extends to a dynamic setting. We start by showing how the dynamic payoff formula derived in Section 3 permits one to compute expected rents and then show how the latter can be used to derive profit-maximizing mechanisms.

Suppose that, in addition to the N agents, there is a “principal” (referred to as “agent 0”) who

designs the mechanism and whose payoff takes the quasilinear form

$$U_0(\theta, x, p) = u_0(x, \theta) + \sum_{i=1}^N p_i$$

for some measurable function $u_0 : \Theta \times X \rightarrow \mathbb{R}$. As standard in the literature, we assume that the principal designs the mechanism and then makes a take-it-or-leave-it offer to the agents in period one after each agent has observed his first-period type.²⁴ We then restrict the principal to offer a mechanism that is accepted in equilibrium by all agents with probability one. Hereafter, we will refer to any such mechanism as an Individually-Rational Bayesian Incentive-Compatible (IR-BIC) mechanism.

The requirement that all agents accept the mechanism gives rise to participation constraints in period 1. In addition, agents might have the ability to quit the mechanism at later stages, which may give rise to participation constraints in subsequent periods. However, the principal can always relax all the participation constraints after the initial acceptance decision by asking each agent to post a bond when accepting the mechanism; this bond is forfeited if the agent quits the mechanism, else is returned to the agent after period T .²⁵ Thus, we can restrict attention to the participation constraints that each agent faces at the moment he is being offered the mechanism. This constraint requires that each agent's expected payoff in the mechanism upon observing his first-period type be at least as high as the payoff the agent obtains by refusing to participate in the mechanism (i.e. his reservation payoff). For simplicity, we assume that reservation payoffs are equal to zero for all agents and all types. The participation constraints can then be written as

$$V^{\Omega_i[\chi, \psi]}(\theta_{i1}) \geq 0 \quad \text{for all } i, \text{ almost all } \theta_{i1} \in \Theta_{i1}. \quad (14)$$

The principal's problem thus consists in choosing an ex-ante BIC mechanism $\langle \chi, \psi \rangle$ that maximizes her expected payoff among those that satisfy the agents' period-1 participation constraints.

While this problem appears quite complicated, it can be simplified by first setting up a "Relaxed Program" that contains only a subset of the constraints, and then providing conditions for a solution to the Relaxed Program to satisfy all of the constraints. In particular, the Relaxed Program replaces all the incentive-compatibility constraints with the local incentive-compatibility conditions embodied in the period-1 dynamic payoff formula derived in Section 3. Specifically, assuming for

²⁴If the principal could approach the agents at the ex-ante stage, before they learn their private information, she could extract all the surplus and hence she would implement an efficient allocation rule.

²⁵The possibility of bonding relies on the following assumptions: (a) unrestricted monetary transfers (in particular, unlimited liability); (b) quasilinear utilities (which rules out any benefit from consumption smoothing); and (c) continuation utilities in the mechanism being bounded from below and continuation utilities from quitting being bounded from above. If these assumptions are not satisfied, one has to consider participation constraints in all periods, which makes the analysis considerably harder.

simplicity that the distributions satisfy Assumption 7, according to Proposition 2, ex-ante IC for agent i implies that

$$V^{\Omega_i[\chi, \psi]}(\theta_{i1}) \text{ is Lipschitz continuous, and for a.e. } \theta_{i1},$$

$$\frac{\partial V^{\Omega_i[\chi, \psi]}(\theta_{i1})}{\partial \theta_{i1}} = \mathbb{E}^{\mu_i^T[\chi]|\theta_{i1}} \left[\sum_{\tau=1}^T J_{i1}^\tau(\tilde{\theta}_i^\tau, \tilde{x}_i^{\tau-1}) \frac{\partial u_i(\tilde{\theta}, \tilde{x})}{\partial \theta_{i1}} \right]. \quad (15)$$

The requirement that $\langle \chi, \psi \rangle$ is ex-ante BIC then implies that, for each $i = 1, \dots, N$, agent i 's ex-ante equilibrium expected payoff must coincide with the expectation of his value function. Condition (15) can then be used to calculate the agents' expected information rents. Using (15) and integrating by parts, then gives

$$\mathbb{E}^{\lambda[\chi]}[U_i(\tilde{\theta}, \chi(\tilde{\theta}), \psi_i(\tilde{\theta}))] =$$

$$\mathbb{E}^{\lambda[\chi]} \left[\frac{1}{\eta_{i1}(\tilde{\theta}_{i1})} \sum_{\tau=1}^T J_1^\tau(\tilde{\theta}_i^\tau, \chi_i^\tau(\tilde{\theta}^\tau)) \frac{\partial u_i(\tilde{\theta}, \chi(\tilde{\theta}))}{\partial \theta_{i\tau}} \right]$$

$$+ V^{\Omega_i[\chi, \psi]}(\underline{\theta}_{i1}). \quad (16)$$

As for the participation constraints, the Relaxed Program considers only those for the lowest types $\underline{\theta}_{i1}$:

$$V^{\Omega_i[\chi, \psi]}(\underline{\theta}_{i1}) \geq 0 \quad (17)$$

Finally, the relaxed program treats the functions $V^{\Omega_i[\chi, \psi]}(\underline{\theta}_{i1})$ in (16) and (17) as control variables that can be chosen independently from (χ, ψ) . Formally, the Relaxed Program can thus be stated as follows.

$$\mathcal{P}^r : \begin{cases} \max_{\chi, \psi, (V^{\Omega_i[\chi, \psi]}(\underline{\theta}_{i1}))_{i=1}^N} \mathbb{E}^{\lambda[\chi]}[U_0(\tilde{\theta}, \chi(\tilde{\theta}), \psi(\tilde{\theta}))] \\ \text{s.t., for all } i = 1, \dots, N, \text{ (16) and (17) hold} \end{cases}$$

Substituting (16) into the principal's payoff then gives the following result.

Lemma 3 *Suppose the environment is quasilinear, that types are independent, and that for each $i = 1, \dots, N$, the assumptions of either Proposition 2 or Proposition 4 hold, and $\underline{\theta}_{i1} > -\infty$. Then*

the principal's expected payoff in any IR-BIC mechanism $\langle \chi, \psi \rangle$ equals

$$\begin{aligned} \mathbb{E}^{\lambda[\chi]}[U_0(\tilde{\theta}, \chi(\tilde{\theta}), \psi(\tilde{\theta}))] = & \\ & \mathbb{E}^{\lambda[\chi]} \left[\sum_{i=0}^N u_i(\tilde{\theta}, \chi(\tilde{\theta})) - \sum_{i=1}^N \frac{1}{\eta_{i1}(\tilde{\theta}_{i1})} \sum_{t=1}^T J_{i1}^t(\tilde{\theta}, \chi(\tilde{\theta})) \frac{\partial u_i(\tilde{\theta}, \chi(\tilde{\theta}))}{\partial \theta_{it}} \right] \\ & - \sum_{i=1}^N V^{\Omega_i[\chi, \psi]}(\underline{\theta}_{i1}), \end{aligned}$$

where $\eta_{i1}(\theta_{i1}) \equiv f_{i1}(\theta_{i1})/(1 - F_{i1}(\theta_{i1}))$ is the hazard rate of the distribution F_{i1} .

In what follows we will refer to the expression

$$\mathbb{E}^{\lambda[\chi]} \left[\sum_{i=0}^N u_i(\tilde{\theta}, \chi(\tilde{\theta})) - \sum_{i=1}^N \frac{1}{\eta_{i1}(\tilde{\theta}_{i1})} \sum_{t=1}^T J_{i1}^t(\tilde{\theta}, \chi(\tilde{\theta})) \frac{\partial u_i(\tilde{\theta}, \chi(\tilde{\theta}))}{\partial \theta_{it}} \right], \quad (18)$$

as the “*expected dynamic virtual surplus*.” It is then immediate that a necessary and a sufficient condition for $(\chi, \psi, (V^{\Omega_i[\chi, \psi]}(\underline{\theta}_{i1}))_{i=1}^N)$ to solve the Relaxed Program is that the allocation rule χ maximizes the expected dynamic virtual surplus, that the participation constraints of the lowest period-1 types bind, i.e.

$$V^{\Omega_i[\chi, \psi]}(\underline{\theta}_{i1}) = 0 \quad \text{for all } i. \quad (19)$$

and that the payments ψ satisfy (16). Clearly, if the solution to the relaxed program satisfies all the incentive and participation constraints, then it also solves the “Full Program” that consists in maximizing the principal's ex-ante expected payoff among all mechanisms that are IR-BIC. We then have the following result.

Proposition 7 *Suppose the environment is quasilinear, that types are independent, and that for each $i = 1, \dots, N$, the assumptions of either Proposition 2 or Proposition 4 hold, and $\underline{\theta}_{i1} > -\infty$. Suppose there exists an IR-BIC mechanism $\langle \chi, \psi \rangle$ such that the allocation rule χ maximizes the “expected dynamic virtual surplus” (18), the lowest types' participation constraints (19) bind, and all the participation constraints (14) are satisfied. Then the following are true:*

- (i) *the mechanism $\langle \chi, \psi \rangle$ solves the Full Program;*
- (ii) *in any mechanism that solves the Full Program, the allocation rule must maximize the expected dynamic virtual surplus (18);*
- (iii) *the principal's expected payoff cannot be increased using randomized mechanisms.*

Proof. Parts (i) and (ii) follow directly from Lemma 3. As for part (iii), note that, from the perspective of each single agent, a randomized mechanism is equivalent to a mechanism that conditions on the types of some fictitious agent $N + 1$. The characterization of the necessary conditions

for incentive compatibility in a stochastic mechanism thus parallels that for deterministic ones. Because the principal's payoff under a stochastic mechanism can always be expressed as a convex combination of her payoffs under different deterministic mechanisms, it is then immediate that stochastic mechanisms cannot raise the principal's expected payoff. (This point was made in static mechanism design by Strausz, 2006). ■

Of course, Proposition 7 is only useful if one can indeed ensure that a solution to the Relaxed Program satisfies all the incentive and participation constraints. We will give some sufficient conditions for this in subsection 4.7. Below we focus on the Relaxed Program and characterize the distortions that a profit-maximizing principal creates relative to the efficient allocation rule.

4.6 Distortions

To begin with, we consider a special class of environments in which the expected virtual surplus (18) can be maximized separately for all times and states without the need to solve a dynamic programming problem. This occurs when, in addition to assumption DNOT, the following property holds.

Assumption 12 (USEP) *Utilities Time-Separable in Decisions:* For all $i = 0, \dots, N$ and all $t = 1, \dots, T$, $u_i(x, \theta) = \sum_{t=1}^T u_{it}(\theta_t, x_t)$.

Recall that, under assumption DNOT, the stochastic process λ over Θ is exogenous and does not depend on the mechanism. If in addition USEP holds, the Relaxed Program is solved by requiring that for all periods t , for λ -almost all θ^t ,

$$\chi_t(\theta^t) \in \arg \max_{x_t \in X_t} \left[\sum_{i=0}^N u_{it}(\theta_t, x_t) - \sum_{i=1}^N \frac{1}{\eta_{i1}(\theta_{i1})} J_{i1}^t(\theta_i^t) \frac{\partial u_{it}(\theta_t, x_t)}{\partial \theta_{it}} \right] \quad (20)$$

It is then easy to compare an allocation rule that satisfies (20) with an efficient allocation rule χ^* , where, by definition, for all periods t and λ -almost all θ^t the latter is such that

$$\chi_t^*(\theta^t) \in \arg \max_{x_t \in X_t} \left[\sum_{i=0}^N u_{it}(\theta_t, x_t) \right]. \quad (21)$$

For simplicity, focus on the case of a single agent: $N = 1$. First, note that when Θ_{1t} is bounded and either $\theta_{1t} = \underline{\theta}_{1t}$ or $\theta_{1t} = \bar{\theta}_{1t}$, then by construction the information index $J_{11}^t(\theta_1^t) = 0$, and so it is optimal to set $\chi_t(\theta_1^t) = \chi_t^*(\theta_1^t)$. Intuitively, when only period-1 participation constraints are relevant, the principal distorts the decisions only to reduce the agent's period-1 information rents. With time-separable utilities, distorting the allocations in period t is then useful only to the extent that the type in period t is informationally linked to the type in period one. When the agent's

type in period t coincides with either the highest or the lowest possible type for that period, the informational link disappears, in which case there is no reason to distort the period- t decision. (In a Markov model, in which $J_{11}^t(\theta_1^t) = \Pi_{\tau=1}^{t-1} J_{1\tau}^{\tau+1}(\theta_1^{\tau+1})$, following $\theta_{1t} = \underline{\theta}_{1t}$ or $\theta_{1t} = \bar{\theta}_{1t}$ distortions then vanish also in all subsequent periods, since the informational link with period 1 is severed.)

It is interesting to contrast this finding with the conclusions of Battaglini (2005), who studies a single-agent model satisfying USEP and DNOT in which the agent's type space in each period has only two elements and where the evolution of the agent's type is governed by a Markov process. In his model, from the moment the agent's type turn out to be high then the optimal mechanism entails no distortions in all subsequent periods (this result is referred to as Generalized No Distortion at the Top, or GNDT). Furthermore, the distortions that the agent experiences when his type remains low are monotonically decreasing in time and vanish in the limit as $T \rightarrow \infty$ (this result is referred to as Vanishing Distortions at the Bottom, or VDB). As the analysis above suggests, while the result of GNDT is quite robust in models satisfying DNOT and USEP, the result of VDB need not be. In particular, distortions need not be monotonic neither in type nor in time and should not be expected to vanish in the long-run.²⁶

On the other hand, for intermediate values of θ_{1t} , distortions are determined by the interaction between the information index, $J_{it}^r(\theta_i^r, x_i^{r-1})$, and the partial derivative of the flow utility $u_{it}(\theta_t, x_t)$ with respect to θ_{it} . For example, suppose that, in addition to the aforementioned assumptions, the following holds.

Assumption 13 (FOSD) *First-Order Stochastic Dominance: For all $i = 1, \dots, N$ and all $t = 2, \dots, T$, $F_{it}(\theta_{it}|\theta_i^{t-1}, x_i^{t-1})$ is nonincreasing in θ_{it}^{t-1} .*

Note that FOSD implies that $J_{it}^r(\theta_i^r, x_i^{r-1}) \geq 0$; comparing the Relaxed Program (20) with the Efficient Program (21), one can then see that in the Relaxed Program the principal distorts x_t to reduce the partial derivative $\partial u_{it}(\theta_t, x_t) / \partial \theta_{it}$. In the standard case in which x_t is one-dimensional and the agent's utility $u_{it}(\theta_t, x_t)$ has the standard single-crossing property, this partial derivative is reduced by reducing x_t . Thus, the solution to the Relaxed Program involves downward distortions in all periods $t > 1$ for intermediate types (and in period $t = 1$ for all but the lowest type). Intuitively, FOSD means that the type in each period $t > 1$ is positively informationally linked to the period-1 type. Then, under the single-crossing property, a downward distortion in period t , by reducing the agent's information rent in period t , then also reduces his information rent in period 1, thus raising the principal's expected payoff.

This result of downward distortions can be extended to settings that do not satisfy assumption USEP and that have many agents, under the following generalization of the single-crossing property.

²⁶We refer the reader to our companion paper, Pavan, Segal and Toikka, 2008, for a further discussion of the dynamics of distortions in profit-maximizing mechanisms.

Assumption 14 (SCP) *Single Crossing Property:* for each t , X_t is a lattice and for each $i = 1, \dots, N$, $u_i(\theta, x)$ has increasing differences in (θ_i, x) .

The assumption that X_t is a lattice is reasonable with one agent. With many agents, it is reasonable, say, if x_t describes the provision of public goods, but it need not hold if x_t is the allocation of a private good (see footnote 18 above for both examples). The lattice structure on each X_t induces a product lattice structure on the set \mathcal{X} of all (measurable) decision rules.

Proposition 8 *Let $\mathcal{X}^0 \subset \mathcal{X}$ denote the set of decision rules solving the Relaxed Program and $\mathcal{X}^* \subset \mathcal{X}$ denote the set of decision rules maximizing expected total surplus. Suppose that, for all $i = 0, \dots, N$, assumptions DNOT, FOSD, and SCP hold, and in addition,*

(i) $u_i(\theta, x)$ is supermodular in x ,

(ii) $\frac{\partial u_i(\theta, x)}{\partial \theta_{it}}$ is submodular in x .

Then $\mathcal{X}^0 \leq \mathcal{X}^*$ in the strong set order.

Proof. Define $g : \mathcal{X} \times \{-1, 0\} \rightarrow \mathbb{R}$ as

$$g(\chi, z) \equiv \mathbb{E}^\lambda \left[\sum_{i=0}^N u_i(\tilde{\theta}, \chi(\tilde{\theta})) + z \sum_{i=1}^N \frac{1}{\eta_{i1}(\tilde{\theta}_{i1})} \sum_{t=1}^T J_{i1}^t(\tilde{\theta}_i) \frac{\partial u_i(\tilde{\theta}, \chi(\tilde{\theta}))}{\partial \theta_{it}} \right].$$

Then $g(\chi, 0)$ is the expected total surplus and $g(\chi, -1)$ is the expected virtual surplus. (Assumption DNOT ensures that the stochastic process $\lambda[\chi]$ doesn't depend on χ and that $J_{i1}^t(\theta_i, x_i)$ does not depend on x_i , which is reflected in the formula.) The assumption of FOSD ensures that $J_{i1}^t(\tilde{\theta}_i) \geq 0$. Together with SCP, this ensures that g has increasing differences in (χ, z) . Together with (i) and (ii), this ensures that g is supermodular in χ . The result then follows from Topkis's Theorem (see, e.g., Topkis, 1998). ■

The result means that if χ^0 solves the relaxed problem and χ^* is efficient, then the decision rule $(\chi^0 \vee \chi^*)_t(\theta) = \chi_t^0(\theta) \vee \chi_t^*(\theta)$ is efficient and the decision rule $(\chi^0 \wedge \chi^*)_t(\theta) = \chi_t^0(\theta) \wedge \chi_t^*(\theta)$ solves the relaxed problem. In particular, if χ^0 and χ^* are defined uniquely with probability one, then $\chi^0(\theta) \leq \chi^*(\theta)$ with probability one.

Note that condition (ii) in Proposition 8 is a 3^{rd} -derivative assumption. Also note that (i) and (ii) hold trivially when each X_t is a chain (e.g., $X_t \subset \mathbb{R}$) and USEP holds.

4.7 Sufficiency and Robustness

We now turn to sufficient conditions for incentive compatibility. As anticipated in the introduction, a complete characterization is evasive because of the multidimensional decision space of the problem. Hereafter, we propose some sufficient conditions for a solution to the Relaxed Program to satisfy all of the incentive and participation constraints that we believe can help in applications.

First we provide sufficient conditions for the participation constraints of all types above the lowest type to be redundant.

Proposition 9 *Suppose for each $i = 1, \dots, N$, $u_i(\theta, x)$ is increasing in each θ_{it} and assumption FOSD holds. Then any mechanism $\langle \chi, \psi \rangle$ satisfying the lowest types' participation constraints (19) and the dynamic payoff formula (15) for period one, satisfies all the participation constraints (14).*

Proof. Under the assumptions in the proposition, $J_{i1}^t(\theta, \chi(\theta)) \geq 0$ and $\partial u_i(\theta, x) / \partial \theta_{it} \geq 0$, hence, by (15), $V^{\Omega_i[\chi, \psi]}(\theta_{i1})$ is nondecreasing in θ_{i1} . ■

Next, consider incentive constraints. In what follows we provide conditions ensuring not only that a mechanism is ex-ante incentive compatible, but that it is also incentive compatible on the equilibrium path. That is, the value function of each agent i at any of his *truthful* private history h_i coincides with his equilibrium expected payoff:

$$V^{\Omega_i[\chi, \psi]}(h_i) = \mathbb{E}^{\mu_i[\chi, \psi]|h_i}[u_i(\tilde{\theta}, \tilde{x}) - \tilde{p}_i].$$

This stronger version of incentive-compatibility thus guarantees that the allocation rule χ is implementable also under a stronger solution concept such as weak Perfect Bayesian Equilibrium.

First observe that, for any given allocation rule χ , one can construct payment schemes ψ such that the resulting utility that each agent obtains in equilibrium (i.e., under truthtelling by all agents) satisfies all the necessary conditions of (15): i.e., at any truthful history $h_{i,t-1} = (\theta_i^{t-1}, \theta_i^{t-1}, x_i^{t-1})$,

$$\begin{aligned} \Phi_{it}(\theta_{it}, h_{i,t-1}) &\equiv \mathbb{E}^{\mu_i[\chi, \psi]|(\theta_{it}, h_{i,t-1})} [u_i(\tilde{\theta}, \tilde{x}) - \tilde{p}_i] \text{ is Lipschitz continuous in } \theta_{it}, \text{ and for a.e. } \theta_{it}, \\ \frac{\partial \Phi_{it}(\theta_{it}, h_{i,t-1})}{\partial \theta_{it}} &= \mathbb{E}^{\mu_i^T[\chi]|(\theta_{it}, h_{i,t-1})} \left[\sum_{\tau=t}^T J_{i1}^\tau(\tilde{\theta}_i^\tau, \tilde{x}_i^{\tau-1}) \frac{\partial u_i(\tilde{\theta}, \tilde{x})}{\partial \theta_{i\tau}} \right]. \end{aligned} \quad (22)$$

(Recall that $\mu_i^T[\chi]|h_i$ denotes the probability distribution on $\Theta^T \times \Theta_i^T \times X$ induced by the allocation rule χ when all agents other than i play truthful strategies, agent i 's private history is h_i , and agent i reports truthfully in the future.) To construct these payments, for all i , all $(\theta_i^t, x_i^{t-1}) \in \Theta_i^t \times X_i^{t-1}$, and all $m_{it} \in \Theta_{it}$, let

$$D_i^{[\chi]}(\theta_i^t, (\theta_i^{t-1}, m_{it}), x_i^{t-1}) \equiv \mathbb{E}^{\mu_i^T[\chi]|(\theta_i^t, (\theta_i^{t-1}, m_{it}), x_i^{t-1})} \left[\sum_{\tau=t}^T J_{it}^\tau(\tilde{\theta}_i^\tau, \tilde{x}_i^\tau) \frac{\partial u_i(\tilde{\theta}, \tilde{x})}{\partial \theta_{i\tau}} \right]. \quad (23)$$

We then have the following result.

Lemma 4 *Suppose the environment is quasilinear, that types are independent, and that for each $i = 1, \dots, N$, the assumptions of either Proposition 2 or Proposition 4 hold. Let $\langle \chi, \psi \rangle$ be any deterministic direct mechanism. Fix a period t . Consider the payment scheme $\hat{\psi}$ obtained from $\langle \chi, \psi \rangle$ by setting for all i and all $\theta \in \Theta$,*

$$\begin{aligned} \hat{\psi}_i(\theta) &= \psi_i(\theta) + \delta_i(\theta_i^t, \chi_i^{t-1}(\theta^{t-1})), \text{ where} \\ \delta_i(\theta_i^t, \chi_i^{t-1}) &\equiv \mathbb{E}^{\mu_i^T[\chi] | (\theta_i^t, \theta_i^t, \chi_i^{t-1})} \left[u_i(\tilde{\theta}, \tilde{x}) - \psi_i(\tilde{\theta}) \right] - \int_{\underline{\theta}_{it}}^{\theta_{it}} D_i^{[\chi]}((\theta_i^{t-1}, z), (\theta_i^{t-1}, z), \chi_i^{t-1}) dz. \end{aligned}$$

Then for all i , and for all truthful private histories $h_{i,t-1} = (\theta_i^{t-1}, \theta_i^{t-1}, \chi_i^{t-1}) \in H_{i,t-1}$, in period t the mechanism $\langle \chi, \hat{\psi} \rangle$ satisfies condition (22).

Proof. By construction, for all truthful private histories $h_{i,t-1} = (\theta_i^{t-1}, \theta_i^{t-1}, \chi_i^{t-1})$,

$$\begin{aligned} \mathbb{E}^{\mu_i[\chi, \hat{\psi}] | (\theta_{it}, h_{i,t-1})} [u_i(\tilde{\theta}, \tilde{x}) - \tilde{p}_i] &= \mathbb{E}^{\mu_i^T[\chi] | (\theta_i^t, \theta_i^t, \chi_i^{t-1})} [u_i(\tilde{\theta}, \tilde{x}) - \psi_i(\tilde{\theta})] \\ &\quad - \mathbb{E}^{\mu_i^T[\chi] | (\theta_i^t, \theta_i^t, \chi_i^{t-1})} \left[\delta_i(\theta_i^t, \chi_i^{t-1}(\tilde{\theta}^{t-1})) \right] \\ &= \int_{\underline{\theta}_{it}}^{\theta_{it}} D_i^{[\chi]}((\theta_i^{t-1}, z), (\theta_i^{t-1}, z), \chi_i^{t-1}) dz, \end{aligned}$$

The first equality follows from the fact that $h_{i,t-1}$ is truthful and the fact that $\mu_i^T[\chi]$ corresponds to the distribution over $\Theta^T \times \Theta_i^T \times X$ under truth-telling (by all agents). The second equality follows directly from the definition of $\delta_i(\theta_i^t, \chi_i^{t-1})$. Note that the function $D^{[\chi]}((\theta_i^{t-1}, \cdot), (\theta_i^{t-1}, \cdot), \chi_i^{t-1})$ is measurable and bounded and therefore integrable. Thus the mechanism $\langle \chi, \hat{\psi} \rangle$ satisfies (22) in period t . ■

Note that the construction achieves the satisfaction of condition (22) in period t by adding to the original payment scheme $\psi_i(\theta)$ a payment term that depends only on reports up to period t ; by implication, this construction does not affect the agents' incentives in subsequent periods. Thus, for any given allocation rule χ , iterating the construction of the payments backward from period T to period one yields a mechanism that, in any period, after any truthful history $h_{i,t-1}$ satisfies condition (22).

Now, using the payments constructed in Lemma 4, we provide a sufficient condition for the allocation rule χ to be implementable, which is obtained by specializing Proposition 5 to quasilinear environments.

Proposition 10 *Suppose the environment is quasilinear, that types are independent, and that for each $i = 1, \dots, N$, the assumptions of either Proposition 2 or Proposition 4 hold. Suppose the mechanism $\langle \chi, \psi \rangle$ is IC at any (possibly non-truthful) period $t + 1$ private history. If for all i , all*

$(\theta_i^t, x_i^{t-1}),$

$D_i^{[x]}(\theta_i^t, (\theta_i^{t-1}, m_{it}), x_i^{t-1})$ is nondecreasing in m_{it} ,

then there exists payment rule $\hat{\psi}$ such that the mechanism $\langle \chi, \hat{\psi} \rangle$ is IC at (i) any truthful period t private history, and (b) at any (possibly non-truthful) period $t + 1$ private history.

Proof. Let $\hat{\psi}$ be the payment rule that is obtained from $\langle \chi, \psi \rangle$ using the construction indicated in the proof of Lemma 4. By construction, $\hat{\psi}$ preserves the agents' incentives at all period $t+1$ histories. Hence the mechanism $\langle \chi, \hat{\psi} \rangle$ satisfies condition (iii) of Proposition 5. The payment scheme $\hat{\psi}$ also ensures that, after any truthful private history $h_{i,t-1} = (\theta_i^{t-1}, \theta_i^{t-1}, x_i^{t-1})$, the mechanism $\langle \chi, \hat{\psi} \rangle$ satisfies condition (22) in period t . This establishes condition (i) of Proposition 5 for period t . The assumption that $D_i^{[x]}(\theta_i^t, (\theta_i^{t-1}, m_{it}), x_i^{t-1})$ is nondecreasing in m_{it} then implies that condition (ii) of Proposition 5 is also verified. The result then follows from Proposition 5. ■

To understand this result intuitively, fix a truthful history $(\theta_i^{t-1}, \theta_i^{t-1}, x_i^{t-1})$, and let $\Psi_t(\theta_{it}, m_{it})$ denote agent i 's expected utility at this history as a function of his new type θ_{it} and his new report m_{it} . One can think of m_{it} as a one-dimensional ‘‘allocation’’ chosen by agent i in period t . Note that $\partial \Psi_t(\theta_{it}, m_{it}) / \partial \theta_{it} = D_i^{[x]}(\theta_i^t, (\theta_i^{t-1}, m_{it}), x_i^{t-1})$; because the mechanism $\langle \chi, \psi \rangle$ is IC at any (possibly non-truthful) period $t + 1$ history, this follows from the dynamic payoff formula (2) applied to the modified mechanism in which agent i 's report of θ_{it} is ignored and replaced with the message m_{it} . If this expression is nondecreasing in m_{it} , then Ψ_t has the single-crossing property (formally, increasing differences). By standard static one-dimensional screening arguments, the monotonic ‘‘allocation rule’’ $m_{it}(\theta_{it}) = \theta_{it}$ is then implementable (using payments constructed from the dynamic payoff formula using the construction in Lemma 4).

The proposition cannot in general be iterated backward, since it assumes IC at all period $t + 1$ histories but derives IC only at truthful period t histories. This reflects a fundamental problem with ensuring incentives in dynamic mechanisms: once an agent has lied once, he may find it optimal to continue lying, and it is hard to characterize his continuation strategy. However, the proposition can still be applied to some interesting special cases. In particular, in a Markov environment, an agent's true past types are irrelevant for incentives given his current type. This implies that IC at truthful histories implies IC at *all* histories. Then the proposition can be rolled backward to show that the mechanism is IC at all histories. This result also implies that truthful strategies, together with the beliefs over the other agents' types constructed from the mechanism $\langle \chi, \psi \rangle$ as rcpd as indicated in subsection 4.3, form a *weak PBE* of the mechanism.

The result in Proposition 10 may also turn useful in certain non-Markov environments, as illustrated in subsection 5.2 below.

Note that Proposition 10 can also be used to analyze the effects of disclosing to the agents in

the course of the mechanism information in addition to the minimal one, as captured by x_{it} . Such disclosure can be captured formally by introducing a measurable space X_{it}^d of possible disclosures to agent i in period t , and then considering the extended set $\hat{X}_{it} = X_{it} \times X_{it}^d$, so that $\hat{x}_{it} = (x_{it}, x_{it}^d)$. While the payoff and the stochastic process describing the evolution of agent i 's type continues to depend on \hat{x}_{it} only through x_{it} , the role of x_{it}^d is to capture the additional information that the mechanism discloses to agent i about the other agents' reports (and hence about the decisions x_{-it} as well). The result in Proposition 10 can then be extended to this environment by redefining $D_i^{[\chi]}$ so that the expectation in (23) is now made conditional on $\hat{x}_{it} = (x_{it}, x_{it}^d)$ instead of just x_{it} . Clearly, the monotonicity condition in the proposition is harder to satisfy when more information is disclosed, but it may still be possible.

In particular, we can formulate a simple condition on the allocation rule that ensures robustness to an extreme form of disclosure. Namely, suppose that each agent i somehow learns at the beginning of period t (i.e before sending his period- t report) all the other agents' types θ_{-i} (note that this includes past, current and future ones). Formally, this can be captured through a disclosure $x_{it}^d = \theta_{-i}$. We then say that the mechanism is *Other-Ex-Post IC (OEP-IC)* if truthtelling remains an optimal strategy in this mechanism at any history. It turns out that some allocation rules can be implemented in an OEP-IC mechanism, under some additional assumptions.

Assumption 15 (PDPD) *Payoffs Depend on Private Decisions: $u_i(\theta, x)$ depends on x only through x_i .*

Corollary 1 *Suppose the environment is quasilinear, that types are independent, and that for each $i = 1, \dots, N$, the assumption of either Proposition 2 or Proposition 4 hold. Suppose in addition that assumptions DNOT, FOSD, SCP and PDPD hold and that the mechanism $\langle \chi, \psi \rangle$ is OEP-IC at any (possibly non-truthful) period $t + 1$ private history. If for all i and all $\tau \geq t$,*

$$\chi_{i\tau}(\theta^\tau) \text{ is nondecreasing in } (\theta_{it}, \dots, \theta_{i\tau}) \text{ for all } \theta_i^{t-1}, \theta_{-i}^\tau, \quad (24)$$

then there exists a payment rule $\hat{\psi}$ such that the mechanism $\langle \chi, \hat{\psi} \rangle$ is OEP-IC at (i) any truthful period t private history, and (ii) at any (possibly non-truthful) period $t + 1$ private history.

Proof. Under assumption DNOT, the stochastic process $\lambda[\chi]$ over Θ does not depend on the allocation rule χ and hence can be written as λ . Furthermore, because types are independent, then λ is the product of each individual agent i 's stochastic process over Θ_i , which henceforth we denote by λ_i . For any θ_i^t , we then denote by $\lambda_i|\theta_i^t$ the distribution over Θ_i given θ_i^t .

The payment rule $\hat{\psi}$ is obtained by adapting the construction of Lemma 4 to the situation where agent i has observed θ_{-i} and faces a stochastic process λ_i over his own types (which is essentially

a single-agent situation):

$$\begin{aligned}\hat{\psi}_i(\theta) &= \psi_i(\theta) + \delta_i(\theta_i^t, \theta_{-i}), \text{ where} \\ \delta_i(\theta_i^t, \theta_{-i}) &= \mathbb{E}^{\lambda_i|\theta_i^t} \left[u_i(\tilde{\theta}_i, \theta_{-i}, \chi(\tilde{\theta}_i, \theta_{-i})) - \psi_i(\tilde{\theta}_i, \theta_{-i}) \right] \\ &\quad - \int_{\underline{\theta}_{it}}^{\theta_{it}} D_i^{[\chi]}((\theta_i^{t-1}, z), (\theta_i^{t-1}, z), \theta_{-i}) dz,\end{aligned}$$

and where

$$D_i^{[\chi]}(\theta_i^t, (\theta_i^{t-1}, m_{it}), \theta_{-i}) \equiv \mathbb{E}^{\lambda_i|\theta_i^t} \left[\sum_{\tau=t}^T J_{it}^\tau(\tilde{\theta}_i^\tau) \frac{\partial u_i((\tilde{\theta}_i, \theta_{-i}), \chi((m_{it}, \tilde{\theta}_{i,-t}), \theta_{-i}))}{\partial \theta_{i\tau}} \right].$$

Note that, under assumption DNOT, $J_{it}^\tau(\theta, x_i)$ does not depend on x_i . By FOSD, $J_{it}^\tau(\theta) \geq 0$. By SCP, PDPD, and (24), $\partial u_i(\theta, \chi((m_{it}, \theta_{i,-t}), \theta_{-i})) / \partial \theta_{i\tau}$ is nondecreasing in m_{it} for all θ_{-i} . This implies that $D_i^{[\chi]}(\theta_i^t, (\theta_i^{t-1}, m_{it}), \theta_{-i})$ is nondecreasing in m_{it} for all θ_i^t and all θ_{-i} . The result then follows from Proposition 10 applied to this setting. ■

For example, in a Markov environment, backward iteration of the Corollary implies that under its assumptions, any allocation rule that is “strongly monotone” in the sense that each $\chi_{it}(\theta_i^t, \theta_{-i}^t)$ is nondecreasing in θ_i^t for any given θ_{-i}^t (which Matthews and Moore (1987) call “attribute monotonicity”) is implementable in an OEP-IC mechanism, and therefore in an BIC mechanism under any possible information disclosure. While it should be clear from Proposition 10 that strong monotonicity is not necessary for implementability, it is particularly easy to check it in applications and it does ensure nice robustness to any kind of information disclosure in the mechanism. Subsections 5.2 and 5.3 provide examples of applications where the profit-maximizing allocation rule turns out to be strongly monotone.

Remark 2 *At this point, the reader may wonder whether we could also ensure robustness to an agent observing his own future types from the outset. This is not likely. Indeed, if agent i observes all of his types from the outset, his IC would be characterized as in a multidimensional screening problem. It is well known that incentives are harder to ensure in this setting. For example, in the special case with a single agent with linear utility $u(\theta, x) = \sum_{t=1}^T \theta_t x_t$, a necessary condition for implementability of allocation rule χ is the “Law of Supply”*

$$\sum_{t=1}^T (\chi_t(\theta'_i) - \chi_t(\theta_i)) (\theta'_i - \theta_i) \geq 0 \text{ for all } \theta', \theta \in \Theta.$$

Because the profit-maximizing allocation rules derived in applications typically fail to satisfy this condition, one cannot obtain robustness to the agents’ observations of their own future types “for

free.” Thus, while some authors have drawn analogies between dynamic mechanism design and static multidimensional mechanism design problems (see, e.g., Courty and Li, 2000 and Rochet and Stole, 2003), here we highlight an important difference: significantly more allocation rules are implementable in a dynamic setting in which the agents learn (and report) the dimensions of their types sequentially over time than in a static setting in which they observe (and report) all dimensions at once.

Remark 3 The reader may also wonder whether there are simple conditions on the payoffs and the kernels that ensure that the allocation rule solving the Relaxed Program 18 is strongly monotone. Unfortunately, any such conditions would have to be very restrictive. Indeed, recall from Subsection 4.6 that in a separable environment (i.e. under USEP) at any period $t > 1$, the distortion in x_{it} is determined by the information index $J_{i1}^t(\theta_i^t)$ which need not be monotonic in θ_{it} ; in particular, when Θ_{it} is bounded, the distortion is zero at both $\theta_{it} = \underline{\theta}_{it}$ and $\theta_{it} = \bar{\theta}_{it}$ and downward at intermediate θ_{it} . Thus, because of this nonmonotonic downward distortion, we can have $\chi_{it}(\theta_{it}, \theta_i^{t-1}, \theta_{-i}^t) < \chi(\underline{\theta}_{it}, \theta_i^{t-1}, \theta_{-i}^t)$ for some $\theta_{it} > \underline{\theta}_{it}$. Indeed, it is to ensure that the solution to the Relaxed Program is implementable that Eso and Szentes (2007) make their Assumption 1 that amounts to requiring that $J_{i1}^2(\theta_{i1}, \theta_{i2})$ is nondecreasing in θ_{i2} . However, note that with a bounded type space Θ_{i2} , this assumption can be satisfied only when the information index is identically zero so that θ_{i1} and θ_{i2} are independent. In the applications below we will consider AR(k) processes with unbounded type spaces in which case the information indices are constant—this helps ensuring strong monotonicity of the solution to the Relaxed Program.

5 Applications

We now show how the results in the previous sections can be put to work by examining a few applications where the agents’ types evolve according to linear AR(k) processes. First, we consider a class of problems where the profit-maximizing mechanism takes the form of a quasi-efficient, or handicapped, mechanism where distortions depend only on the agents’ first period types. Next, we consider environments where payoffs separate over time as it is often assumed in applications. Lastly, we consider a setting where agents refine their valuations through consumption, as in the case of experience goods.

5.1 Handicapped mechanisms

Consider an environment where in each period the set of feasible allocations is $X_t \subset \mathbb{R}^{N+1}$. The utility to each agent $i = 1, \dots, N$ (gross of payments) is

$$u_i(\theta, x) = \sum_{t=1}^T \theta_{it} x_{it} - c_i(x), \quad (25)$$

where $c_i : \mathbb{R}^{(N+1)T} \rightarrow \mathbb{R}$ is an intertemporal cost function. The principal's (gross) payoff function is $u_0(\theta, x) = v_0(x)$. Note that the cost functions c_i and the principal's payoff v_0 need not be time-separable; this permits us to accommodate dynamic aspects such as intertemporal capacity constraints, habit formation, and learning-by-doing. The private information of each agent $i = 1, \dots, N$ is assumed to evolve according to a linear AR(k) process, as in Example 4. The total information indices $J_{i1}^t(\theta, x)$ are thus independent of (θ, x) and coincide with the "impulse response functions" J_{i1}^t for the AR(k) process. We assume that the support of the first period innovation ε_{i1} (and hence that of θ_{i1}) is bounded from below.

In this environment, the expected dynamic virtual surplus takes the form

$$\mathbb{E}^\lambda \left[v_0(\chi(\tilde{\theta})) + \sum_{i=1}^N \left[\sum_{t=1}^T \left(\tilde{\theta}_{it} \chi_{it}(\tilde{\theta}^t) - J_{i1}^t \eta_{i1}^{-1}(\tilde{\theta}_{i1}) \chi_{it}(\tilde{\theta}^t) \right) - c_i(\chi(\tilde{\theta})) \right] \right].$$

Note that the latter coincides with the expected total surplus in a model where the (gross) payoff to each agent i is $u_i(\theta, x)$ and where the (gross) payoff to the principal is

$$\hat{v}_0(\theta, x) \equiv v_0(x) - \sum_{i=1}^N \sum_{t=1}^T J_{i1}^t \eta_{i1}^{-1}(\theta_{i1}) x_{it}.$$

This implies that the solution to the Relaxed Program can be obtained by solving an efficient program where the principal has an extra marginal cost $J_{i1}^t \eta_{i1}^{-1}(\theta_{i1})$ of allocating a unit to agent i in period t . In general, this efficient program is a dynamic programming problem. However, in many applications, its solution can be readily found using existing methods. What is important to us is the following observation. Assuming the period-one types are reported truthfully, then any allocation rule that maximizes the expected dynamic virtual can be sustained through an "Handicapped" efficient mechanism. In period 1 each agent i sends a message m_{i1} determining his handicaps $J_{i1}^t \eta_{i1}^{-1}(m_{i1})$. The game that starts in period two then corresponds to a private-value environment where each agent i 's payoff, for $i = 1, \dots, N$, is as in (25), whereas agent 0's payoff (i.e. the principal's) is $\hat{v}_0(\theta, x)$. Because the decisions that are implemented are the efficient decisions for this environment and because this virtual environment is a private-value one, then incentives

at any period $t \geq 2$ can be provided using for example the “Team payments” (Athey and Segal, 2007) defined by

$$\psi_i(\theta) = \sum_{j \neq i} u_j(\theta, \chi(\theta)),$$

for all i , and all θ . We then have the following result (the proof follows directly from the arguments above).²⁷

Proposition 11 *In the environment with $AR(k)$ types described above, any allocation rule that maximizes the expected dynamic virtual surplus can be implemented in a mechanism that satisfies IC at all truthful histories in periods $t \geq 2$.*

Incentives in the first period must be checked application-by-application.²⁸ The environment considered in the next subsection provides an example where incentive compatibility obtains also at $t = 1$. In particular, incentive-compatibility in each period can be easily guaranteed in the environment considered above if the costs $c_i \equiv 0$ for all i (the environment then becomes a special case of the class considered in the next subsection.)

5.2 Time-Separable Environments

We now consider an environment in which the agents’ types continue to follow an $AR(k)$ process as in Example 4, but where payoffs separate over time. The set of possible decisions in each period t is $X_t \subset \mathbb{R}^{N+1}$. Each agent i (with the principal as agent 0) has an utility function of the form

$$u_i(\theta, x) = \sum_{t=1}^T u_{it}(\theta_{it}, x_{it}),$$

with the principal’s types θ_{0t} being common knowledge. As in the previous subsection, the support of the first period types is assumed to be bounded from below.

This model can fit many applications including sequential auctions (possibly with risk-averse bidders), procurement, and regulation.

Proposition 12 *Consider the separable environment with $AR(k)$ types described above. Suppose the assumptions of Proposition 2 hold for each agent $i = 1, \dots, N$. Suppose further that for all $i = 0, \dots, N$ and all periods t , the following are true: (1) the periodic utility function u_{it} has increasing differences in (θ_{it}, x_{it}) ; (2) the coefficient ϕ_{it} of the $AR(k)$ process is nonnegative; (3) the first-period*

²⁷What is important for the result in the next proposition is that (i) each agent’s payoff is affine in his own types and independent of the other agents’ types, (ii) that the principal’s payoff is independent of θ , and (iii) that the total information indexes are independent of θ .

²⁸At period 1 the model where the principal has extra costs is one with interdependent values since these costs depends on the agents’ true period-1 types through the hazard rates $\eta_{i1}(\theta_{i1})$.

hazard rate $\eta_{i1}(\theta_{i1})$ is nondecreasing; and (4) the partial derivative $\frac{\partial u_{it}(\theta_{it}, x_{it})}{\partial \theta_{it}}$ is nonnegative and submodular in (θ_{it}, x_{it}) . Then an allocation rule χ can be part of a profit-maximizing mechanism if and only if, for all t , and λ -almost all θ^t ,

$$\chi_t(\theta^t) \in \arg \max_{x_t} \left\{ u_{0t}(\theta_{0t}, x_{0t}) + \sum_{i=1}^N \left(u_{it}(\theta_{it}, x_{it}) - \frac{J_{i1}^t}{\eta_{i1}(\theta_{i1})} \frac{\partial u_{it}(\theta_{it}, x_{it})}{\partial \theta_{it}} \right) \right\}. \quad (26)$$

Furthermore, χ can be implemented in a mechanism that is OEP-IC at any history using payments constructed as follows. For any agent $i = 1, \dots, n$ and all θ ,

$$\psi_i(\theta) = \psi_{i1}(\theta_{i1}, \theta_{-i}^T) + \sum_{t=2}^T \psi_{it}(\theta_1, \theta_t), \quad (27)$$

where for all $t \geq 2$,

$$\psi_{it}(\theta_1, \theta_t) \equiv u_{it}(\theta_{it}, \chi_{it}(\theta_1, \theta_t)) - \int_{\underline{\theta}_{it}}^{\theta_{it}} \frac{\partial u_{it}(r, \chi_{it}(\theta_1, (r, \theta_{-i,t})))}{\partial \theta_{it}} dr, \quad (28)$$

and²⁹

$$\begin{aligned} \psi_{i1}(\theta_{i1}, \theta_{-i}^T) &\equiv \mathbb{E}^{\lambda_i | \theta_{i1}} \left[u_i((\tilde{\theta}_i^T, \theta_{-i}^T), \chi(\tilde{\theta}_i^T, \theta_{-i}^T)) - \sum_{t=2}^T \psi_{it}(\theta_1, (\tilde{\theta}_{it}, \theta_{-i,t})) \right] \\ &\quad - \int_{\underline{\theta}_{i1}}^{\theta_{i1}} \mathbb{E}^{\lambda_i | r} \left[\sum_{\tau=1}^T J_{i1}^\tau \frac{\partial u_i(\tilde{\theta}_i, \chi_i((r, \tilde{\theta}_{i,-1}), \theta_{-i}^T))}{\partial \theta_{i\tau}} \right] dr. \end{aligned} \quad (29)$$

The result in Proposition 12 follows essentially from Corollary 1 by observing that in this environment *incentives separate over time*. By inspecting (26) one can in fact see that the allocation rule that maximizes the expected dynamic virtual surplus has the property that in each period t the allocation $\chi_t(\theta^t)$ depends only on the current reports θ_t and on the agents' period-1 reports θ_1 . This in turn is a consequence of the following assumptions: (i) preferences are separable over time (USEP), (ii) decisions do not affect types (DNOT), and (iii) informational indexes $J_{i\tau}^t$ do not depend on the realized types, as it is the case with AR(k) processes. The problem that each agent i faces at any period $t \geq 2$ when he must choose which report to send then becomes a static problem (despite the informational channel that links types together over time). However, the fact that incentives separate over time alone does not guarantee implementability. As it is known from the literature on static mechanism design, one also needs “enough monotonicity” in the allocation rule $\chi_t(\theta_1, \theta_t)$ with respect to the reports of own types. Assumptions (1), (2) and (4) in the proposition (which

²⁹Recall that the notation $\lambda_i | \theta_i^t$ denotes the unique probability measure on Θ_i^T that corresponds to the stochastic process that starts in period one with θ_{i1} and whose transitions are given by the kernels of the AR(k) process.

imply SCP, FOSD, and PDPD) guarantee that each $\chi_{it}(\theta_1, \theta_t)$ is monotone in θ_{it} . Constructing payments that induce the agents to report truthfully at any period $t \geq 2$, even after they observe the other agents' types³⁰, is then particularly simple. It suffices to have each agent i pay in each period $t \geq 2$ a transfer $\psi_{it}(\theta_1, \theta_t)$ that, given (θ_1, θ_{-it}) , coincides with the payment type θ_{it} would made in a static mechanism implementing the monotone allocation rule $\chi_{it}(\cdot; \theta_1, \theta_{-it})$, as indicated in (28).

As for period one, in general providing incentives at $t = 1$ is difficult. However, note that assumptions (1)-(4) in the proposition guarantee that the allocation rule that maximizes the dynamic virtual surplus is *strongly monotone* in the sense of Corollary 1. Following the same steps as in the proof of Corollary 1, one can then add to the payments $\psi_{it}(\theta_1, \theta_t)$ —which for convenience can be assumed to be made in each of the corresponding periods— a final payment of $\psi_{i1}(\theta_{i1}, \theta_{-i}^T)$ to be made in period T , after all other agents' types θ_{-i}^T have been revealed. When the payments $\psi_{i1}(\theta_{i1}, \theta_{-i}^T)$ are as in (29) then incentives for truthtelling are guaranteed also in period one.

Finally consider possible implementations of the profit-maximizing rule. First, note that in the linear case (i.e., when $u_{it}(\theta_{it}, x_{it}) = \theta_{it}x_{it}$) the implementation is particularly simple. Suppose there is no allocation in the first period and assume the agents do not observe the other agents' types (both assumptions simplify the discussion but are clearly not essential for the argument). In period one, each agent i chooses from a menu of “handicaps” $J_{i1}^t \eta_{i1}^{-1}(\theta_{i1})$, indexed by θ_{i1} . Then in each period $t \geq 2$, a “handicapped” VCG mechanism is played with transfers as in (28). Lastly, in period $T + 1$, each agent is asked to make a final payment of $\psi_{i1}(\theta_{i1}, \tilde{\theta}_{-i}^T)$ (Eso and Szentes (2007) derive this result in the special case of a two-period model with allocation only in the second period.) This logic extends to nonlinear payoffs in the sense that in the first period the agents still choose from a menu of future plans (indexed by the first period type). In the subsequent periods the distortions now generally depend also on the current reports through the partial derivatives $\frac{\partial u_{it}(\theta_{it}, x_{it})}{\partial \theta_{it}}$. However intermediate reports (i.e., reports in periods $2, \dots, t - 1$) remain irrelevant both for the period- t allocation and for the period- t payments.

5.3 Learning

The last application we consider pertains the design of an optimal mechanism for a seller facing a buyer who learns his valuation over time through consumption. The problem arises for example in the markets for experience goods (such as prescription drugs) and in expert services (such as chiropractor's service). There is a seller who in each period $t = 1, \dots, T$ can produce a service at cost c . There is a single buyer whose valuation for the service is v . Payoffs are quasilinear and take the form of $\sum_t (p_t - x_t c_t)$ for the seller and $\sum_t (x_t v - p_t)$ for the buyer, where $x_t \in \{0, 1\} = X_t$.

³⁰In fact, due to time-separability, in periods $t \geq 2$ the mechanism is truly ex post IC in that it is robust also to the possibility that each agent i observes his own future types.

Neither the buyer nor the seller knows v . The buyer's prior belief is that $v \sim N(\theta_1, \frac{1}{\alpha})$, where θ_1 is the mean and α the *precision* (i.e., the inverse of the variance). The seller knows that the buyer's prior belief is Normal with precision α but does not know the mean θ_1 of the buyer's prior belief. The seller believes that θ_1 is distributed on $[\underline{\theta}_1, \bar{\theta}_1)$ according to some absolutely continuous cdf F_1 with $F_1' > 0$. We assume that the hazard rate $\eta_1(\theta_1)$ of F_1 is nondecreasing. If the buyer consumes the service in period t (i.e., if $x_t = 1$), he then receives a signal $s_t = v + \varepsilon_t$ where $\varepsilon_t \sim N(0, \frac{1}{\beta})$, where β is precision of the signal. The noises ε_t are i.i.d. and independent of v . If the buyer does not consume in period t , he does not receive any new information about v .³¹

Given the form of the buyer's payoff and the Gaussian structure of the underlying learning process, the relevant statistic for contracting in each period t is the buyer's posterior expectation of v , which we denote by θ_t . Using the properties of the Normal distribution, the evolution of θ_t can be expressed recursively as follows ((see, e.g., DeGroot, 2004)). For any x^t , let $\|x^t\| \equiv \sum_{\tau=1}^t x_\tau$ denote the number of times the buyer consumed the service in periods $1, \dots, t$. The buyer's posterior belief about v at the beginning of period $t = 1, \dots, T$ given x^{t-1} is then Normal with mean

$$\theta_t \equiv \frac{\alpha\theta_1 + \beta \sum_{j \in \{\tau: \tau < t \wedge x_\tau = 1\}} s_j}{\alpha + \|x^{t-1}\| \beta}$$

and precision $\alpha_t = \alpha + \|x^{t-1}\| \beta$. Depending on whether the buyer consumed or not the good in period $t-1$, we then have two cases. If $x_{t-1} = 0$, then $\theta_t = \theta_{t-1}$ and $\alpha_t = \alpha_{t-1}$. If instead $x_{t-1} = 1$, then

$$\theta_t = \frac{\alpha\theta_1 + \beta \sum_{j \in \{\tau: \tau < t-1 \wedge x_\tau = 1\}} s_j + \beta s_{t-1}}{\alpha + \|x^{t-1}\| \beta} = \frac{\alpha_{t-1}\theta_{t-1} + \beta s_{t-1}}{\alpha_{t-1} + \beta}$$

and $\alpha_t = \alpha_{t-1} + \beta = \alpha + \|x^{t-1}\| \beta$, where $\alpha_{t-1} = \alpha + \|x^{t-2}\| \beta$. Note that θ_t is a weighted average of the period $t-1$ posterior θ_{t-1} and the period $t-1$ signal s_{t-1} . Thus, before the signal s_{t-1} is realized, we have that

$$\theta_t | (\theta^{t-1}, x^{t-2}, x_{t-1} = 1) \sim N \left(\theta_{t-1}, \frac{\beta}{(\alpha + \|x^{t-1}\| \beta)(\alpha + \|x^{t-2}\| \beta)} \right),$$

and

$$\theta_t | (\theta^{t-1}, x^{t-2}, x_{t-1} = 0) = \theta_{t-1}.$$

These expressions define Markov kernels $F_t(\cdot | \theta_{t-1}, x^{t-1})$, where the sequence of past allocations determines the precision.

We first show that in terms of payoffs it is without loss to restrict attention to a subclass of allocation rules.

³¹See also Nazerzadeh, Saberi, and Vohra (2008) for a similar environment.

Definition 12 An allocation rule χ is a stopping rule if, for all t , all $s > t$ and all $\theta \in \Theta$, $\chi_t(\theta^t) = 0$ implies $\chi_s(\theta^s) = 0$. The set of stopping rules is denoted \mathcal{X}^S .

Lemma 5 Consider the learning environment described above. If $\langle \chi, \psi \rangle$ is an ex-ante IC mechanism, then there exists an ex-ante IC mechanism $\langle \hat{\chi}, \hat{\psi} \rangle$ such that $\hat{\chi}$ is a stopping rule and the expected payoffs of both the buyer and the seller under $\langle \hat{\chi}, \hat{\psi} \rangle$ are the same as under $\langle \chi, \psi \rangle$.

The lemma is similar to the well-known result that in a two-armed bandit problem with one safe arm the optimal strategy is a stopping rule. Given this result, in what follows we restrict attention to stopping rules. Then the only relevant period- t histories are the ones in which the agent has consumed in all the preceding periods. Thus we can replace $\|x^t\|$ in all the formulas above by t . In particular, before stopping, we have that

$$\theta_{t+1}|\theta_t \sim N\left(\theta_t, \frac{\beta}{(\alpha + t\beta)[\alpha + (t-1)\beta]}\right).$$

Denoting the standard deviation of the period- $t+1$ posterior by $\rho_{t+1} \equiv \sqrt{\beta[(\alpha + t\beta)(\alpha + (t-1)\beta)]^{-1/2}}$ we can then express the kernels as $F_{t+1}(\theta_{t+1}|\theta^t, x^t) = \Phi\left(\frac{\theta_{t+1}-\theta_t}{\rho_{t+1}}\right)$, where Φ is the cdf of the standard normal distribution. Thus, before stopping, the model satisfies the assumptions of Proposition 2 and the direct information index between any two adjacent periods is simply

$$I_t^{t+1}(\theta^{t+1}) = -\frac{\partial F_{t+1}(\theta_{t+1}|\theta^t, x^t) / \partial \theta_t}{f_{t+1}(\theta_{t+1}|\theta^t, x^t)} = -\frac{\partial \Phi\left(\frac{\theta_{t+1}-\theta_t}{\rho_{t+1}}\right) / \partial \theta_t}{\frac{1}{\rho_{t+1}} \phi\left(\frac{\theta_{t+1}-\theta_t}{\rho_{t+1}}\right)} = 1,$$

where ϕ is the density of the standard normal distribution. Since the model is Markovian, $I_t^\tau \equiv 0$ for $\tau > t + 1$. Hence, before stopping, we have $J_t^\tau \equiv 1$ for all τ and t . The maximization of the dynamic virtual surplus then takes the form

$$\max_{\chi \in \mathcal{X}^S} \mathbb{E}^{\lambda[\chi]} \left[\sum_{t=1}^T \chi_t(\tilde{\theta}^t) \left(\tilde{\theta}_t - c - \frac{1}{\eta_1(\theta_1)} \right) \right],$$

where the maximization is over the set of stopping rules \mathcal{X}^S .

The dynamic virtual surplus for this model is thus analogous to the one for the separable environment of the previous subsection. However, unlike in that model, we cannot solve here the maximization of the dynamic virtual surplus by pointwise maximization because it is a stopping problem. Instead, we proceed by backward induction. While it is difficult to get a close-form solution for the optimal allocation rule, it is possible to characterize it partially and get a clean comparison to the efficient allocation rule.

Definition 13 A stopping rule $\chi \in \mathcal{X}^S$ is a cutoff rule if for all t and all θ^{t-1} , $\chi_t(\theta^{t-1}, \theta_t)$ is nondecreasing in θ_t . The cutoffs are given by $z_t(\theta^{t-1}) \equiv \inf \{ \theta_t \in \Theta_t : \chi_t(\theta^{t-1}, \theta_t) = 1 \}$.

Proposition 13 Consider the learning environment described above. The following are true:

(1) The efficient allocation rule χ^* is a cutoff rule where for all t and all θ^{t-1} , the cutoff $z_t^*(\theta^{t-1}) \equiv z_t^*$ is independent of θ^{t-1} and nondecreasing in t .

(2) The allocation rule that solves the Relaxed Program is a cutoff rule χ where for all t and all θ^{t-1} , the cutoff $z_t(\theta^{t-1}) \equiv z_t(\theta_1)$ is independent of θ_{-1}^{t-1} , nondecreasing in t , and nonincreasing in θ_1 .

(3) For all t and all θ_1 , $z_t(\theta_1) \geq z_t^*$. In particular, together with (4) this implies that a profit maximizing monopoly experiments less than what is socially desirable.

(4) Both the allocation rule that solves the Relaxed Program χ and the efficient rule χ^* are implementable.

That the profit-maximizing cutoffs are increasing in t is due to the fact that the option value of learning is decreasing in the number of times the service has been provided. First, the impact of each new signal on the buyer's posterior belief declines with the number of signals received in the past. Second, as the remaining horizon gets shorter, the seller will reap the benefits from high valuations in fewer periods.

Perhaps more interestingly, the cutoffs in the profit-maximizing allocation rule depend on the buyer's first-period type. This implies that the optimal selling mechanism *cannot* be implemented with a sequence of prices. Actually, even history-dependent prices fail to implement the optimal mechanism. In fact, what is essential is to condition the prices not on the purchase history x^{t-1} but on the first period type θ_1 . This can be done by offering the buyer a *menu* of contracts, where each contract corresponds to a different price path. Because the optimal cutoffs are increasing in time, so are the prices in each path. To build demand, the monopolist thus optimally offers “introductory rates,” or “discounts,” that expire after the service has been provided for a few periods.

References

- ANGUS, J. E. (1994): “The Probability Integral Transform and Related Results,” *SIAM Review*, 36(4), 652–654.
- ATHEY, S., AND I. SEGAL (2007): “An Efficient Dynamic Mechanism,” 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.
- CALZOLARI, G., AND A. PAVAN (2006a): “Monopoly with Resale,” *Rand Journal of Economics*, 37(2), 362–375.
- (2006b): “On the Optimality of Privacy in Sequential Contracting,” *Journal of Economic Theory*, 130(1), 168–204.
- COURTY, P., AND H. LI (2000): “Sequential Screening,” *Review of Economic Studies*, 67(4), 697–717.
- CREMER, J., AND R. P. MCLEAN (1988): “Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions,” *Econometrica*, 56(6), 1247–1257.
- DEGROOT, M. H. (2004): *Optimal Statistical Decisions*. Wiley-IEEE.
- DOEPKE, M., AND R. M. TOWNSEND (2006): “Dynamic Mechanism Design with Hidden Income and Hidden Actions,” *Journal of Economic Theory*, 126(1), 235–285.
- DUDLEY, R. M. (2002): *Real Analysis and Probability*. Cambridge University Press.
- ELY, J. (2001): “Revenue Equivalence Without Differentiability Assumptions,” Mimeo, Northwestern University.
- ESO, P., AND B. SZENTES (2007): “Optimal Information Disclosure in Auctions and the Handicap Auction,” *Review of Economic Studies*, 74(3), 705–731.
- GARCIA, D. (2005): “Monotonicity in Direct Revelation Mechanisms,” *Economics Letters*, 88(1), 21–26.
- GERSHKOV, A., AND B. MOLDOVANU (2007): “The Dynamic Assignment of Heterogenous Objects: A Mechanism Design Approach,” Discussion Paper, University of Bonn.
- MATTHEWS, S., AND J. MOORE (1987): “Monopoly Provision of Quality and Warranties: An Exploration in the Theory of Multidimensional Screening,” *Econometrica*, 55(2), 441–467.
- MEZZETTI, C. (2004): “Mechanism Design with Interdependent Valuations: Efficiency,” *Econometrica*, 72(5), 1617–1626.
- MILGROM, P., AND I. SEGAL (2002): “Envelope Theorems for Arbitrary Choice Sets,” *Econometrica*, 70(2), 583–601.

- MIRPLEES, J. A. (1971): “An Exploration in the Theory of Optimum Income Taxation,” *Review of Economic Studies*, 38(114), 175–208.
- MYERSON, R. B. (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.
- RUDIN, W. (1976): *Principles of Mathematical Analysis*. McGraw-Hill, 3 edn.
- STOKEY, N. L., AND R. E. LUCAS, JR. (1989): *Recursive Methods in Economic Dynamics*. Harvard University Press.
- STRAUSZ, R. (2006): “Deterministic versus Stochastic Mechanisms in Principal-Agent Models,” *Journal of Economic Theory*, 127(1), 306–314.
- TOPKIS, D. M. (1998): *Supermodularity and Complementarity*. Princeton University Press.

Appendices

A Statement and proof of Lemma A.1

Lemma A.1. *Assume the environment satisfies Assumption 2. Then Assumption 5 implies that for any t , and any $\tau < t$*

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

Proof of Lemma A.1. Assumption 5 implies that

$$\begin{aligned} \left| \frac{\partial}{\partial \theta_\tau} \int \theta_t dF_t(\theta_t | \theta^{t-1}, y^{t-1}) \right| &= \left| \lim_{\theta'_\tau \rightarrow \theta_\tau} \frac{\int \theta_t d[F_t(\theta_t | \theta^{t-1}_{-\tau}, \theta'_\tau, y^{t-1}) - F_t(\theta_t | \theta^{t-1}_{-\tau}, \theta_\tau, y^{t-1})]}{\theta'_\tau - \theta_\tau} \right| \\ &= \left| - \lim_{\theta'_\tau \rightarrow \theta_\tau} \int \frac{F_t(\theta_t | \theta^{t-1}_{-\tau}, \theta'_\tau, y^{t-1}) - F_t(\theta_t | \theta^{t-1}_{-\tau}, \theta_\tau, y^{t-1})}{\theta'_\tau - \theta_\tau} d\theta_t \right| \\ &= \left| - \int \frac{\partial F_t(\theta_t | \theta^{t-1}, y^{t-1})}{\partial \theta_\tau} d\theta_t \right|, \end{aligned}$$

The second inequality follows by Lemma 6 below. The last equality follows by the dominated convergence theorem since the integrand is bounded for all θ_t by the integrable function $B_t(\theta_t)$. Furthermore,

$$\left| - \int \frac{\partial F_t(\theta_t | \theta^{t-1}, y^{t-1})}{\partial \theta_t} d\theta_t \right| \leq \int B(\theta_t) d\theta_t,$$

from which the claim follows by taking $B \equiv \int B(\theta_t) d\theta_t$. ■

B Proof of Proposition 1

Two kinds of period- t histories appear frequently in the proof. Those including the message m_t but excluding the realization of y_t , and those including the current type θ_t but excluding the message m_t . For expositional clarity we introduce notation to distinguish the value functions associated with these two types of histories. For the first kind, we let $\Psi_t(\theta^t, m^t, y^{t-1}) \equiv V^\Omega(\theta^t, m^t, y^{t-1})$ denote the the supremum continuation expected utility. For the second kind, we continue to use the value function V^Ω but in order to clarify notation further we drop the superscript Ω and add a time subscript. Thus we write $V_t(\theta^t, m^{t-1}, y^{t-1}) \equiv V^\Omega(\theta^t, m^{t-1}, y^{t-1})$. Also, it is convenient to introduce period $T+1$ as a notional device and then let $\Psi_{T+1}(\theta^{T+1}, m^{T+1}, y) = V_{T+1}(\theta^{T+1}, m, y) = U(\theta, y)$. Note that by definition

$$\begin{aligned} \Psi_t(\theta^t, m^t, y^{t-1}) &= \int V_{t+1}(\theta^{t+1}, m^t, y^t) dF_{t+1}(\theta_{t+1} | \theta^t, y^t) d\Omega_t(y_t | m^t, y^{t-1}), \\ V_{t+1}(\theta^{t+1}, m^t, y^t) &= \sup_{m_{t+1}} \Psi_{t+1}(\theta^{t+1}, (m^t, m_{t+1}), y^t). \end{aligned} \quad (30)$$

The proof proceeds in a series of Lemmas.

Lemma 6 *For any Lipschitz function $G : \Theta_t \rightarrow \mathbb{R}$,*

$$\begin{aligned} \int G(\theta_t) dF_t(\theta_t | \theta^{t-1}, y^{t-1}) - \int G(\theta_t) dF_t(\theta_t | \eta^{t-1}, y^{t-1}) \\ = - \int G'(\theta_t) (F_t(\theta_t | \theta^{t-1}, y^{t-1}) - F_t(\theta_t | \eta^{t-1}, y^{t-1})) d\theta_t, \end{aligned}$$

where all the integrals exist.

Proof. First note that the first two integrals exist, since letting M be the Lipschitz constant for G , and picking any $\hat{\theta}_t \in \Theta_t$, we can write $|G(\theta_t)| \leq |G(\hat{\theta}_t)| + M|\hat{\theta}_t| + M|\theta_t|$, and all terms have expectations with respect to the probability distribution $dF_t(\theta_t | \theta^{t-1}, y^{t-1})$, the last term by

Assumption 2. Thus, we can use integration by parts to write

$$\begin{aligned}
& \int G(\theta_t) dF_t(\theta_t|\theta^{t-1}, y^{t-1}) - \int G(\theta_t) dF_t(\theta_t|\eta^{t-1}, y^{t-1}) \\
&= \int G(\theta_t) d(F_t(\theta_t|\theta^{t-1}, y^{t-1}) - F_t(\theta_t|\eta^{t-1}, y^{t-1})) \\
&= - \int G'(\theta_t) (F_t(\theta_t|\theta^{t-1}, y^{t-1}) - F_t(\theta_t|\eta^{t-1}, y^{t-1})) d\theta_t \\
&\quad + [G(\theta_t) (F_t(\theta_t|\theta^{t-1}, y^{t-1}) - F_t(\theta_t|\eta^{t-1}, y^{t-1}))]_{\theta_t=\underline{\theta}_t}^{\theta_t=\bar{\theta}_t}.
\end{aligned}$$

When both $\bar{\theta}_t$ and $\underline{\theta}_t$ are finite, we have $F_t(\bar{\theta}_t|\theta^{t-1}, y^{t-1}) = F_t(\bar{\theta}_t|\eta^{t-1}, y^{t-1}) = 1$ and $F_t(\underline{\theta}_t|\theta^{t-1}, y^{t-1}) = F_t(\underline{\theta}_t|\eta^{t-1}, y^{t-1}) = 0$, and the Lemma follows. If $\underline{\theta}_t = -\infty$, then as $\theta_t \rightarrow -\infty$,

$$\begin{aligned}
& |G(\theta_t) (F_t(\theta_t|\theta^{t-1}, y^{t-1}) - F_t(\theta_t|\eta^{t-1}, y^{t-1}))| \\
&\leq (|G(\hat{\theta}_t)| + M|\hat{\theta}_t|) |F_t(\theta_t|\theta^{t-1}, y^{t-1}) - F_t(\theta_t|\eta^{t-1}, y^{t-1})| \\
&\quad + M|\theta_t| (F_t(\theta_t|\theta^{t-1}, y^{t-1}) + F_t(\theta_t|\eta^{t-1}, y^{t-1})) \\
&\leq (|G(\hat{\theta}_t)| + M|\hat{\theta}_t|) |F_t(\theta_t|\theta^{t-1}, y^{t-1}) - F_t(\theta_t|\eta^{t-1}, y^{t-1})| \\
&\quad + M \left(\int_{z \leq \theta_t} |z| dF_t(z|\theta^{t-1}, y^{t-1}) + \int_{z \leq \theta_t} |z| dF_t(z|\eta^{t-1}, y^{t-1}) \right) \\
&\rightarrow 0
\end{aligned}$$

by Assumptions 2 and 5. Similarly, if $\bar{\theta}_t = +\infty$, then as $\theta_t \rightarrow +\infty$,

$$\begin{aligned}
& |G(\theta_t) (F_t(\theta_t|\theta^{t-1}, y^{t-1}) - F_t(\theta_t|\eta^{t-1}, y^{t-1}))| \\
&\leq (|G(\hat{\theta}_t)| + M|\hat{\theta}_t|) |F_t(\theta_t|\theta^{t-1}, y^{t-1}) - F_t(\theta_t|\eta^{t-1}, y^{t-1})| \\
&\quad + M|\theta_t| [(1 - F_t(\theta_t|\theta^{t-1}, y^{t-1})) + (1 - F_t(\theta_t|\eta^{t-1}, y^{t-1}))] \\
&\leq (|G(\hat{\theta}_t)| + M|\hat{\theta}_t|) |F_t(\theta_t|\theta^{t-1}, y^{t-1}) - F_t(\theta_t|\eta^{t-1}, y^{t-1})| \\
&\quad + M \left(\int_{z \geq \theta_t} |z| dF_t(z|\theta^{t-1}, y^{t-1}) + \int_{z \geq \theta_t} |z| dF_t(z|\eta^{t-1}, y^{t-1}) \right) \\
&\rightarrow 0
\end{aligned}$$

by Assumptions 2 and 5. ■

For any function $G : \Theta \rightarrow \mathbb{R}$, let

$$\frac{\partial^- G(\theta)}{\partial \theta_t} = \limsup_{\theta'_t \uparrow \theta_t} \frac{G(\theta'_t, \theta_{-t}) - G(\theta)}{\theta'_t - \theta_t} \quad \text{and} \quad \frac{\partial^+ G(\theta)}{\partial \theta_t} = \liminf_{\theta'_t \downarrow \theta_t} \frac{G(\theta'_t, \theta_{-t}) - G(\theta)}{\theta'_t - \theta_t}.$$

The following Lemma is similar to Theorem 1 of Milgrom and Segal (2002) and Theorem 1 of Ely (2001).

Lemma 7 *In an ex ante IC mechanism Ω , for any integers $1 \leq t \leq \tau$ and for $\mu[\Omega]$ -almost all histories $(\theta^\tau, \theta^{\tau-1}, y^{\tau-1})$,*

$$\frac{\partial^- V_\tau(\theta^\tau, \theta^{\tau-1}, y^{\tau-1})}{\partial \theta_t} \leq \frac{\partial^- \Psi_\tau(\theta^\tau, \theta^\tau, y^{\tau-1})}{\partial \theta_t} \quad \text{and} \quad \frac{\partial_+ V_\tau(\theta^\tau, \theta^{\tau-1}, y^{\tau-1})}{\partial \theta_t} \geq \frac{\partial_+ \Psi_\tau(\theta^\tau, \theta^\tau, y^{\tau-1})}{\partial \theta_t}.$$

Proof. By ex ante IC we have for $\mu[\Omega]$ -almost all histories $(\theta^\tau, \theta^{\tau-1}, y^{\tau-1})$,

$$V_\tau(\theta^\tau, \theta^{\tau-1}, y^{\tau-1}) = \Psi_\tau(\theta^\tau, \theta^\tau, y^{\tau-1}).$$

By definition of V_τ and Ψ_τ , we have for all $(\theta^\tau, \theta^{\tau-1}, y^{\tau-1})$ and all θ'_t ,

$$V_\tau((\theta'_t, \theta_{-t}^\tau), \theta^{\tau-1}, y^{\tau-1}) \geq \Psi_\tau((\theta'_t, \theta_{-t}^\tau), \theta^\tau, y^{\tau-1}).$$

Combining the two we have for $\mu[\Omega]$ -almost all histories $(\theta^\tau, \theta^{\tau-1}, y^{\tau-1})$ and all θ'_t ,

$$V_\tau((\theta'_t, \theta_{-t}^\tau), \theta^{\tau-1}, y^{\tau-1}) - V_\tau(\theta^\tau, \theta^{\tau-1}, y^{\tau-1}) \geq \Psi_\tau((\theta'_t, \theta_{-t}^\tau), \theta^\tau, y^{\tau-1}) - \Psi_\tau(\theta^\tau, \theta^\tau, y^{\tau-1}).$$

Taking $\theta'_t > \theta_t$, dividing by $\theta'_t - \theta_t$, and then taking liminf as $\theta'_t \downarrow \theta_t$ yields the second inequality in the lemma. Taking $\theta'_t < \theta_t$, dividing by $\theta'_t - \theta_t$, and then taking limsup as $\theta'_t \uparrow \theta_t$ yields the first inequality in the lemma. ■

The next two lemmas don't rely on IC.

Lemma 8 *For each t , $\Psi_t(\theta^t, m^t, y^t)$ and $V_t(\theta^t, m^{t-1}, y^t)$ are equi-Lipschitz continuous in θ^t — i.e., there exists M such that for all $\theta^t, \eta^t, m^t, y^t$,*

$$\begin{aligned} |\Psi_t(\eta^t, m^t, y^t) - \Psi_t(\theta^t, m^t, y^t)| &\leq M \|\eta^t - \theta^t\|, \\ |V_t(\eta^t, m^{t-1}, y^t) - V_t(\theta^t, m^{t-1}, y^t)| &\leq M \|\eta^t - \theta^t\|. \end{aligned}$$

Proof. By backward induction on t . $\Psi_{T+1}(\theta^{T+1}, m^{T+1}, y^T) = U(\theta^T, y^T)$ is equi-Lipschitz continuous in θ^T by Assumption 4. Now we show that for any t , if $\Psi_t(\theta^t, m^t, y^{t-1})$ is equi-Lipschitz continuous in θ^t , then $V_t(\theta^t, m^{t-1}, y^{t-1})$ and $\Psi_{t-1}(\theta^{t-1}, m^{t-1}, y^{t-2})$ are equi-Lipschitz continuous in θ^t and θ^{t-1} , respectively.

Indeed, suppose $\Psi_t(\theta^t, m^t, y^{t-1})$ is equi-Lipschitz continuous in θ^t with a constant M . Then

$$\begin{aligned} |V_t(\eta^t, m^{t-1}, y^{t-1}) - V_t(\theta^t, m^{t-1}, y^{t-1})| &\leq \sup_{m_t} |\Psi_t(\eta^t, (m^{t-1}, m_t), y^{t-1}) - \Psi_t(\theta^t, (m^{t-1}, m_t), y^{t-1})| \\ &\leq M \|\eta^t - \theta^t\|, \end{aligned}$$

and so V_t is also equi-Lipschitz continuous in θ^t . But then, using (30),

$$\begin{aligned} &|\Psi_{t-1}(\eta^{t-1}, m^{t-1}, y^{t-2}) - \Psi_{t-1}(\theta^{t-1}, m^{t-1}, y^{t-2})| \\ &\leq \sup_{y_{t-1}} \left| \int V_t((\eta^{t-1}, \theta_t), m^{t-1}, y^{t-1}) dF_t(\theta_t | \eta^{t-1}, y^{t-1}) - \int V_t((\theta^{t-1}, \theta_t), m^{t-1}, y^{t-1}) dF_t(\theta_t | \theta^{t-1}, y^{t-1}) \right| \\ &\leq \sup_{y_{t-1}} \left| \int (V_t((\eta^{t-1}, \theta_t), m^{t-1}, y^{t-1}) - V_t((\theta^{t-1}, \theta_t), m^{t-1}, y^{t-1})) dF_t(\theta_t | \eta^{t-1}, y^{t-1}) \right| \\ &\quad + \sup_{y_{t-1}} \left| \int V_t(\theta^t, m^{t-1}, y^{t-1}) dF_t(\theta_t | \eta^{t-1}, y^{t-1}) - \int V_t(\theta^t, m^{t-1}, y^{t-1}) dF_t(\theta_t | \theta^{t-1}, y^{t-1}) \right| \\ &\leq \sup_{y_{t-1}} \int |V_t((\eta^{t-1}, \theta_t), m^{t-1}, y^{t-1}) - V_t((\theta^{t-1}, \theta_t), m^{t-1}, y^{t-1})| dF_t(\theta_t | \eta^{t-1}, y^{t-1}) \\ &\quad + \sup_{y_{t-1}} \int |F_t(\theta_t | \eta^{t-1}, y^{t-1}) - F_t(\theta_t | \theta^{t-1}, y^{t-1})| \left| \frac{\partial V_t(\theta^t, m^{t-1}, y^{t-1})}{\partial \theta_t} \right| d\theta_t \\ &\leq M \|\eta^{t-1} - \theta^{t-1}\| \left(1 + \int B_t(\theta_t) d\theta_t \right), \end{aligned}$$

where we used Lemma 6 and Assumption 5. This shows that Ψ_{t-1} is equi-Lipschitz continuous in θ^{t-1} . ■

Lemma 9 For any integers τ, t such that $1 \leq t < \tau \leq T$, and any $(\theta^{\tau-1}, m^{\tau-1}, y^{\tau-2})$,

$$\frac{\partial^- \Psi_{\tau-1}(\theta^{\tau-1}, m^{\tau-1}, y^{\tau-2})}{\partial \theta_t} \leq \int \frac{\partial^- V_\tau(\theta^\tau, m^{\tau-1}, y^{\tau-1})}{\partial \theta_t} dF_\tau(\theta_\tau | \theta^{\tau-1}, y^{\tau-1}) d\Omega_{\tau-1}(y_{\tau-1} | m^{\tau-1}, y^{\tau-2}) \quad (31)$$

$$\begin{aligned} &- \int \frac{\partial V_\tau(\theta^\tau, m^{\tau-1}, y^{\tau-1})}{\partial \theta_\tau} \frac{\partial F_\tau(\theta_\tau | \theta^{\tau-1}, y^{\tau-1})}{\partial \theta_t} d\theta_\tau d\Omega_{\tau-1}(y_{\tau-1} | m^{\tau-1}, y^{\tau-2}), \\ \frac{\partial_+ \Psi_{\tau-1}(\theta^{\tau-1}, m^{\tau-1}, y^{\tau-2})}{\partial \theta_t} &\geq \int \frac{\partial_+ V_\tau(\theta^\tau, m^{\tau-1}, y^{\tau-1})}{\partial \theta_t} dF_\tau(\theta_\tau | \theta^{\tau-1}, y^{\tau-1}) d\Omega_{\tau-1}(y_{\tau-1} | m^{\tau-1}, y^{\tau-2}) \quad (32) \end{aligned}$$

$$- \int \frac{\partial V_\tau(\theta^\tau, m^{\tau-1}, y^{\tau-1})}{\partial \theta_\tau} \frac{\partial F_\tau(\theta_\tau | \theta^{\tau-1}, y^{\tau-1})}{\partial \theta_t} d\theta_\tau d\Omega_{\tau-1}(y_{\tau-1} | m^{\tau-1}, y^{\tau-2}).$$

Proof. Using (30), write for any $\theta'_t \neq \theta_t$

$$\begin{aligned} & \frac{\Psi_{\tau-1}((\theta'_t, \theta_{-t}^{\tau-1}), m^{\tau-1}, y^{\tau-2}) - \Psi_{\tau-1}(\theta^{\tau-1}, m^{\tau-1}, y^{\tau-2})}{\theta'_t - \theta_t} \\ &= \int \frac{V_\tau((\theta'_t, \theta_{-t}^\tau), m^{\tau-1}, y^{\tau-1}) - V_\tau(\theta^\tau, m^{\tau-1}, y^{\tau-1})}{\theta'_t - \theta_t} dF_\tau(\theta_\tau | \theta^{\tau-1}, y^{\tau-1}) d\Omega_{\tau-1}(y_{\tau-1} | m^{\tau-1}, y^{\tau-2}) \end{aligned} \quad (33)$$

$$+ \int V_\tau(\theta^\tau, m^{\tau-1}, y^{\tau-1}) d \left[\frac{F_\tau(\theta_\tau | (\theta'_t, \theta_{-t}^{\tau-1}), y^{\tau-1}) - F_\tau(\theta_\tau | \theta^{\tau-1}, y^{\tau-1})}{\theta'_t - \theta_t} \right] d\Omega_{\tau-1}(y_{\tau-1} | m^{\tau-1}, y^{\tau-2}) \quad (34)$$

$$+ \int \frac{V_\tau((\theta'_t, \theta_{-t}^\tau), m^{\tau-1}, y^{\tau-1}) - V_\tau(\theta^\tau, m^{\tau-1}, y^{\tau-1})}{\theta'_t - \theta_t} d [F_\tau(\theta_\tau | (\theta'_t, \theta_{-t}^{\tau-1}), y^{\tau-1}) - F_\tau(\theta_\tau | \theta^{\tau-1}, y^{\tau-1})] \times \quad (35)$$

$$d\Omega_{\tau-1}(y_{\tau-1} | m^{\tau-1}, y^{\tau-2}).$$

We examine separately the behavior of each of the three integrals as $\theta'_t \rightarrow \theta_t$:

- (35): Note that for any $y^{\tau-1}$,

$$\begin{aligned} & \int \frac{V_\tau((\theta'_t, \theta_{-t}^\tau), m^{\tau-1}, y^{\tau-1}) - V_\tau(\theta^\tau, m^{\tau-1}, y^{\tau-1})}{\theta'_t - \theta_t} d [F_\tau(\theta_\tau | (\theta'_t, \theta_{-t}^{\tau-1}), y^{\tau-1}) - F_\tau(\theta_\tau | \theta^{\tau-1}, y^{\tau-1})] \\ & \qquad \qquad \qquad \rightarrow 0 \text{ as } \theta'_t \rightarrow \theta_t, \end{aligned}$$

since the integrand is bounded by Lemma 8, and the total variation of the measure

$$d [F_\tau(\theta_\tau | (\theta'_t, \theta_{-t}^{\tau-1}), y^{\tau-1}) - F_\tau(\theta_\tau | \theta^{\tau-1}, y^{\tau-1})]$$

converges to zero by Assumption 6. Thus, (35) is bounded in absolute value by a term that converges to zero as $\theta'_t \rightarrow \theta_t$.

(Note that in the Markov case, $V_\tau((\theta'_t, \theta_{-t}^\tau), m^{\tau-1}, y^{\tau-1}) - V_\tau(\theta^\tau, m^{\tau-1}, y^{\tau-1}) = U_t(\theta'_t, y^t) - U_t(\theta_t, y^t)$ does not depend on θ_τ so (35) equals zero without imposing Assumption 6.)

- (34): Using Lemma 8 and Lemma 6 it can be expressed as

$$- \int \frac{F_\tau(\theta_\tau | \theta_{-t}^{\tau-1}, \theta'_t, y^{\tau-1}) - F_\tau(\theta_\tau | \theta^{\tau-1}, y^{\tau-1})}{\theta'_t - \theta_t} \frac{\partial V_\tau(\theta^\tau, m^{\tau-1}, y^{\tau-1})}{\partial \theta_\tau} d\theta_\tau d\Omega_{\tau-1}(y_{\tau-1} | m^{\tau-1}, y^{\tau-2}).$$

Using in addition Assumption 5, the Dominated Convergence Theorem establishes that as $\theta'_t \rightarrow \theta_t$, the integral converges to the 2nd integral in (32) and (31).

- (33) Taking its limsup as $\theta'_t \uparrow \theta_t$ and using Fatou's Lemma,³² we see that the limsup is bounded above by the 1st integral in (31). Thus, we obtain (31). Similarly, taking the liminf of (33) as $\theta'_t \downarrow \theta_t$ and using Fatou's Lemma, we see that the liminf of this term is bounded below by the 1st integral in (32), so we obtain (32).

■

Now combining the inequalities in Lemma 9 for $m^\tau = \theta^\tau$ and the inequalities in Lemma 7 we obtain for $\mu[\Omega]$ -almost all histories $(\theta^{\tau-1}, \theta^{\tau-2}, y^{\tau-2})$,

$$\begin{aligned} \frac{\partial^- V_{\tau-1}(\theta^{\tau-1}, \theta^{\tau-2}, y^{\tau-2})}{\partial \theta_t} &\leq \int \frac{\partial^- V_\tau(\theta^\tau, \theta^{\tau-1}, y^{\tau-1})}{\partial \theta_t} dF_\tau(\theta_\tau | \theta^{\tau-1}, y^{\tau-1}) d\Omega_{\tau-1}(y_{\tau-1} | \theta^{\tau-1}, y^{\tau-2}) \\ &\quad - \int \frac{\partial V_\tau(\theta^\tau, \theta^{\tau-1}, y^{\tau-1})}{\partial \theta_\tau} \frac{\partial F_\tau(\theta_\tau | \theta^{\tau-1}, y^{\tau-1})}{\partial \theta_t} d\theta_\tau d\Omega_{\tau-1}(y_{\tau-1} | \theta^{\tau-1}, y^{\tau-2}), \\ \frac{\partial_+ V_{\tau-1}(\theta^{\tau-1}, \theta^{\tau-2}, y^{\tau-2})}{\partial \theta_t} &\geq \int \frac{\partial_+ V_\tau(\theta^\tau, \theta^{\tau-1}, y^{\tau-1})}{\partial \theta_t} dF_\tau(\theta_\tau | \theta^{\tau-1}, y^{\tau-1}) d\Omega_{\tau-1}(y_{\tau-1} | \theta^{\tau-1}, y^{\tau-2}) \\ &\quad - \int \frac{\partial V_\tau(\theta^\tau, \theta^{\tau-1}, y^{\tau-1})}{\partial \theta_\tau} \frac{\partial F_\tau(\theta_\tau | \theta^{\tau-1}, y^{\tau-1})}{\partial \theta_t} d\theta_\tau d\Omega_{\tau-1}(y_{\tau-1} | \theta^{\tau-1}, y^{\tau-2}). \end{aligned}$$

Furthermore, we have by definition of V_{T+1} ,

$$\frac{\partial^- V_{T+1}(\theta^{T+1}, \theta^T, y^T)}{\partial \theta_t} = \frac{\partial_+ V_{T+1}(\theta^{T+1}, \theta^T, y^T)}{\partial \theta_t} = \frac{\partial V_{T+1}(\theta^{T+1}, \theta^T, y^T)}{\partial \theta_t} = \frac{\partial U(\theta^T, y^T)}{\partial \theta_t}.$$

So iterating the above inequalities forward for $\tau = t+1, t+2, \dots, T+1$ yields for $\mu[\Omega]$ -almost all $(\theta^t, \theta^{t-1}, y^{t-1})$ the double inequality

$$\begin{aligned} \frac{\partial^- V_t(\theta^t, \theta^{t-1}, y^{t-1})}{\partial \theta_t} &\leq \\ \mathbb{E}^{\mu[\Omega] | (\theta^t, \theta^{t-1}, y^{t-1})} &\left[\frac{\partial U(\tilde{\theta}^T, \tilde{y}^T)}{\partial \theta_t} - \sum_{\tau=t+1}^T \int \frac{\partial V_\tau(\tilde{\theta}^{\tau-1}, \theta_\tau, \tilde{\theta}^{\tau-1}, \tilde{y}^{\tau-1})}{\partial \theta_\tau} \frac{\partial F_\tau(\theta_\tau | \tilde{\theta}^{\tau-1}, \tilde{y}^{\tau-1})}{\partial \theta_t} d\theta_\tau \right] \\ &\leq \frac{\partial_+ V_t(\theta^t, \theta^{t-1}, y^{t-1})}{\partial \theta_t}. \end{aligned}$$

To complete the proof of the proposition, recall that by definition,

$$V_t(\theta^t, \theta^{t-1}, y^{t-1}) = V^\Omega(\theta^t, \theta^{t-1}, y^{t-1}).$$

³²Note that even though the integrand need not be nonnegative, it is bounded in absolute value by the lipschitz constant M . Thus, in general we may have to add and subtract M from the integrand before applying Fatou's lemma.

So by Lemma 8 $V^\Omega(\theta^t, \theta^{t-1}, y^{t-1})$ is Lipschitz continuous in θ_t for all $(\theta^{t-1}, \theta^{t-1}, y^{t-1})$. Thus, given any $(\theta^{t-1}, \theta^{t-1}, y^{t-1})$, the partial derivative $\frac{\partial V^\Omega(\theta^t, \theta^{t-1}, y^{t-1})}{\partial \theta_t}$ exists for almost every θ_t . Whenever it does, it equals to both ends of the above double inequality and so (IC-FOC) obtains.

C Other Proofs Omitted in the Main Text

Proof of Proposition 4. The initial steps of the proof are in the main text. Here we simply prove that, under the assumptions in the proposition, the formula in (10) reduces to the one in (2).

Differentiating the identity³³

$$F_s(F_s^{-1}(\varepsilon_s | \theta^{s-1}, y^{s-1}) | \theta^{s-1}, y^{s-1}) = \varepsilon_s.$$

with respect to θ_t , $t < s$, we have that for a.e. ε_s ,

$$0 = f_s(\theta_s | \theta^{s-1}, y^{s-1}) \Big|_{\theta_s = F_s^{-1}(\varepsilon_s | \theta^{s-1}, y^{s-1})} \cdot \frac{\partial F_s^{-1}(\varepsilon_s | \theta^{s-1}, y^{s-1})}{\partial \theta_t} + \frac{\partial F_s(\theta_s | \theta^{s-1}, y^{s-1})}{\partial \theta_t} \Big|_{\theta_s = F_s^{-1}(\varepsilon_s | \theta^{s-1}, y^{s-1})},$$

from which we obtain that

$$\frac{\partial F_s^{-1}(\varepsilon_s | \theta^{s-1}, y^s)}{\partial \theta_t} = - \frac{\frac{\partial F_s(\theta_s | \theta^{s-1}, y^{s-1})}{\partial \theta_t} \Big|_{\theta_s = F_s^{-1}(\varepsilon_s | \theta^{s-1}, y^{s-1})}}{f_s(\theta_s | \theta^{s-1}, y^{s-1}) \Big|_{\theta_s = F_s^{-1}(\varepsilon_s | \theta^{s-1}, y^{s-1})}}.$$

It follows that

$$A_j^s(\varepsilon^s, y^{s-1}) = - \frac{\partial F_s(\theta_s | \theta^{s-1}, y^{s-1}) / \partial \theta_j}{f_s(\theta_s | \theta^{s-1}, y^{s-1})} \Big|_{\theta^s = z^s(\varepsilon^s; y^{s-1})} \equiv I_j^s(\theta_s | \theta^{s-1}, y^{s-1}) \Big|_{\theta^s = z^s(\varepsilon^s; y^{s-1})}.$$

We conclude that

$$\begin{aligned} \frac{\frac{\partial z_s(\varepsilon^s; y^{s-1})}{\partial \varepsilon_t}}{\frac{\partial z_t(\varepsilon^t; y^{t-1})}{\partial \varepsilon_t}} &= I_t^s(\theta_s | \theta^{s-1}, y^{s-1}) \Big|_{\theta^s = z^s(\varepsilon^s; y^{s-1})} \\ &\quad + I_{t+1}^s(\theta_s | \theta^{s-1}, y^{s-1}) \Big|_{\theta^s = z^s(\varepsilon^s; y^{s-1})} I_t^{t+1}(\theta_{t+1} | \theta^t, y^t) \Big|_{\theta^{t+1} = z^{t+1}(\varepsilon^{t+1}; y^t)} + \dots \\ &= J_t^s(z^s(\varepsilon^s; y^{s-1}), y^{s-1}). \end{aligned}$$

³³Note that the differentiability of $F_s(\theta_s | \theta^{s-1}, y^{s-1})$ with respect to θ_t , $t < s$, follows from the assumptions in the proposition. This can be seen from the implicit function theorem applied to the identity $F_s^{-1}(F_s(\theta_s | \theta^{s-1}, y^{s-1}) | \theta^{s-1}, y^{s-1}) = \theta_s$.

Rewriting (10) as

$$\frac{\partial V^\Omega(z^t(\varepsilon^t; y^{t-1}), z^{t-1}(\varepsilon^{t-1}; y^{t-2}), y^{t-1})}{\partial \theta_t} = \mathbb{E}^{\hat{\mu}_{[\hat{\Omega}]|\varepsilon_t, \hat{h}^{t-1}}} \left[\frac{\partial U(z^T(\tilde{\varepsilon}^T; \tilde{y}^{T-1}), \tilde{y}^T)}{\partial \theta_t} + \sum_{s=t+1}^T \left(\frac{\frac{\partial z_s(\tilde{\varepsilon}^s; \tilde{y}^{s-1})}{\partial \varepsilon_t}}{\frac{\partial z_t(\varepsilon^t; y^{t-1})}{\partial \varepsilon_t}} \right) \frac{\partial U(z^T(\tilde{\varepsilon}^T; \tilde{y}^{T-1}), \tilde{y}^T)}{\partial \theta_s} \right]$$

we then have that

$$\begin{aligned} \frac{\partial V^\Omega(z^t(\varepsilon^t; y^{t-1}), z^{t-1}(\varepsilon^{t-1}; y^{t-2}), y^{t-1})}{\partial \theta_t} &= \mathbb{E}^{\hat{\mu}_{[\hat{\Omega}]|\varepsilon_t, \hat{h}^{t-1}}} \left[\frac{\partial U(z^T(\tilde{\varepsilon}^T; \tilde{y}^{T-1}), \tilde{y}^T)}{\partial \theta_t} + \sum_{s=t+1}^T J_t^s(z^s(\tilde{\varepsilon}^s; y^{s-1}), y^{s-1}) \frac{\partial U(z^T(\tilde{\varepsilon}^T; \tilde{y}^{T-1}), \tilde{y}^T)}{\partial \theta_s} \right] \\ &= \mathbb{E}^{\mu[\Omega]|z^t(\varepsilon^t; y^{t-1}), z^{t-1}(\varepsilon^{t-1}; y^{t-2}), y^{t-1}} \left[\frac{\partial U(\tilde{\theta}, \tilde{y})}{\partial \theta_t} + \sum_{s=t+1}^T J_t^s(\tilde{\theta}^s, y^{s-1}) \frac{\partial U(\tilde{\theta}, \tilde{y})}{\partial \theta_s} \right], \end{aligned}$$

which is the same formula as in (2). ■

Proof of Proposition 5.

By (iii), it suffices to consider only single-stage deviations in period t , i.e., deviations to some report m_t followed by truthtelling from $t+1$ onward. Thus, it suffices to verify that the agent's period- t payoff expectation from such a deviation at any truthful history $(\theta^{t-1}, \theta^{t-1}, y^{t-1})$ and at any current type θ_t , which is given by

$$\Psi(\theta_t, m_t; \theta^{t-1}, y^{t-1}) \equiv \mathbb{E}^{\mu[\Omega]||(\theta^{t-1}, \theta_t), (\theta^{t-1}, m_t), y^{t-1}} [U(\tilde{y}, \tilde{\theta})],$$

is maximized by reporting $m_t = \theta_t$. For this purpose, the following lemma is useful. (A similar approach has been applied to static mechanism design with one-dimensional type and multidimensional decisions but under stronger assumptions—see Garcia, 2005.)

Lemma 10 *Consider a function $\Psi : (\underline{\theta}, \bar{\theta})^2 \rightarrow \mathbb{R}$. Suppose that (a) $\Psi(\theta, m)$ is Lipschitz continuous in θ for all m , (b) $\Phi(\theta) \equiv \Psi(\theta, \theta)$ is Lipschitz continuous in θ , and (c) for any m , for a.e. θ , $(\Phi'(\theta) - \partial\Psi(\theta, m)/\partial\theta) \cdot (\theta - m) \geq 0$. Then $\Phi(\theta) \geq \Psi(\theta, m)$ for all (θ, m) .*

Proof of the Lemma: Let $g(\theta, m) \equiv \Phi(\theta) - \Psi(\theta, m)$. For any fixed m , $g(\cdot, m)$ is Lipschitz continuous in θ by (a) and (b). Hence, it is differentiable a.e. in θ , and

$$g(\theta, m) = \int_m^\theta \frac{\partial g(z, m)}{\partial \theta} dz = \int_m^\theta \left[\Phi'(z) - \frac{\partial \Psi(z, m)}{\partial \theta} \right] dz.$$

By (c), the integrand is nonnegative for a.e. $z \geq m$ and nonpositive for a.e. $z \leq m$. Therefore, $g(\theta, m) \geq 0$ for both $\theta \geq m$ and $\theta \leq m$. ■

Now, to apply the Lemma, we interpret $\Psi(\theta_t, m_t; \theta^{t-1}, y^{t-1})$ as the agent's expected utility from truthtelling in the mechanism $\hat{\Omega}$ constructed from Ω by ignoring the agent's report in period t and substituting m_t instead. Assumption (iii) means that the mechanism $\hat{\Omega}$ is IC at *any* history in period t , and therefore $\Psi(\theta_t, m_t; \theta^{t-1}, y^{t-1})$ is the agent's value function in the mechanism. Applying to $\hat{\Omega}$ the result in Proposition 2, or equivalently in Proposition 4, we have that, for any m_t , $\Psi(\theta_t, m_t; \theta^{t-1}, y^{t-1})$ is Lipschitz continuous in θ_t and $\partial\Psi(\theta_t, m_t; \theta^{t-1}, y^{t-1})/\partial\theta_t = D^\Omega(\theta^t, (\theta^{t-1}, m_t), y^{t-1})$ a.e. θ_t . The former property establishes assumption (a) in the Lemma. Assumption (i) in the proposition establishes assumption (b) in the Lemma and, together with assumption (ii) in the proposition, it establishes assumption (c) in the Lemma. The Lemma then implies that $\Psi(\theta_t, m_t; \theta^{t-1}, y^{t-1})$ is indeed maximized by reporting $m_t = \theta_t$ which implies that Ω is IC at any truthful period- t history. ■

Proof of Proposition 6. Let $\Omega_i[\chi, \psi]$ and $\Omega_i[\chi, \hat{\psi}]$ denote the randomized direct mechanisms that agent i faces respectively under $\langle \chi, \psi \rangle$ and $\langle \chi, \hat{\psi} \rangle$, as defined above. Let $V^{\Omega_i[\chi, \psi]} : H_i \rightarrow \mathbb{R}$ and $V^{\Omega_i[\chi, \hat{\psi}]} : H_i \rightarrow \mathbb{R}$ denote the corresponding value functions.

We first establish the following result.

Lemma 11 *Suppose the assumptions in Proposition 6 hold. Then, for $\lambda[\chi]$ -almost all truthful private histories h_i^{t-1} , there exists a scalar $K_{it}(h_i^{t-1})$ such that*

$$V^{\Omega_i[\chi, \psi]}(\tilde{\theta}_{it}, \tilde{h}_i^{t-1}) - V^{\Omega_i[\chi, \hat{\psi}]}(\tilde{\theta}_{it}, \tilde{h}_i^{t-1}) = K_{it}(h_i^{t-1}). \quad (36)$$

From Lemma 1, the fact that $\langle \chi, \psi \rangle$ and $\langle \chi, \hat{\psi} \rangle$ are ex-ante BIC implies that they are BIC at $\mu_i^T[\chi]$ -almost all truthful private histories $h_i^{t-1} \equiv (\theta_i^{t-1}, \theta_i^{t-1}, x_i^{t-1})$, for any i and any $t \geq 1$. Iterating (IC-FOC) backward (or alternatively using (6)) and using the result in Proposition 1 (alternatively, the result in Proposition 3), then implies that, under quasi-linearity, for any $t \geq 1$ and $\mu_i^T[\chi]$ -almost all truthful private histories $h_i^{t-1} \equiv (\theta_i^{t-1}, \theta_i^{t-1}, x_i^{t-1})$ the value functions $V^{\Omega_i[\chi, \psi]}(\cdot, h_i^{t-1})$ and $V^{\Omega_i[\chi, \hat{\psi}]}(\cdot, h_i^{t-1})$ are Lipschitz continuous in θ_{it} and

$$\frac{\partial V^{\Omega_i[\chi, \psi]}(\theta_{it}, h_i^{t-1})}{\partial\theta_{it}} = \frac{\partial V^{\Omega_i[\chi, \hat{\psi}]}(\theta_{it}, h_i^{t-1})}{\partial\theta_{it}} \text{ a.e. } \theta_{it}.$$

It follows that for $\lambda[\chi]$ -almost all truthful private histories h_i^{t-1} , there exists a scalar $K_{it}(h_i^{t-1})$ such that the condition in (36) holds. ■

The result for $t = 1$ then follows directly from this lemma by letting $K_i = K_{i1}(h^0)$, where h^0 is

the null history, and noting that, in any ex-ante BIC mechanism, with probability one, the value function coincides with the expected payoff under truth-telling.

The proof for the second result in the proposition is by induction. Suppose there exists a $K_i \in \mathbb{R}$ such that

$$\mathbb{E}^\lambda[V^{\Omega_i[\chi, \psi]}(\tilde{\theta}_{i\tau}, \tilde{h}_i^{\tau-1}) \mid \tilde{\theta}_i^\tau] - \mathbb{E}^\lambda[V^{\Omega_i[\chi, \hat{\psi}]}(\tilde{\theta}_{i\tau}, \tilde{h}_i^{\tau-1}) \mid \tilde{\theta}_i^\tau] = K_i \quad (37)$$

when $\tau = t \geq 1$. We then show that (37) holds also $\tau = t + 1$.

First note that for $\lambda[\chi]$ -almost all private histories (θ_{it}, h_i^{t-1}) ,

$$V^{\Omega_i[\chi, \psi]}(\theta_{it}, h_i^{t-1}) = \mathbb{E}^{\mu_i^T[\chi] \mid \theta_{it}, h_i^{t-1}}[V^{\Omega_i[\chi, \psi]}(\theta_{it+1}, h_i^t)].$$

By the law of iterated expectations, we then have that

$$\mathbb{E}^{\lambda[\chi]}[V^{\Omega_i[\chi, \psi]}(\tilde{\theta}_{it}, \tilde{h}_i^{t-1}) \mid \tilde{\theta}_i^t] = \mathbb{E}^{\lambda[\chi]}[V^{\Omega_i[\chi, \psi]}(\tilde{\theta}_{i,t+1}, \tilde{h}_i^t) \mid \tilde{\theta}_i^t]$$

It follows that

$$\begin{aligned} & \mathbb{E}^{\lambda[\chi]}[V^{\Omega_i[\chi, \psi]}(\tilde{\theta}_{it}, \tilde{h}_i^{t-1}) \mid \tilde{\theta}_i^t] - \mathbb{E}^{\lambda[\chi]}[V^{\Omega_i[\chi, \hat{\psi}]}(\tilde{\theta}_{it}, \tilde{h}_i^{t-1}) \mid \tilde{\theta}_i^t] \\ &= \mathbb{E}^{\lambda[\chi]}[V^{\Omega_i[\chi, \psi]}(\tilde{\theta}_{i,t+1}, \tilde{h}_i^t) \mid \tilde{\theta}_i^t] - \mathbb{E}^{\lambda[\chi]}[V^{\Omega_i[\chi, \hat{\psi}]}(\tilde{\theta}_{i,t+1}, \tilde{h}_i^t) \mid \tilde{\theta}_i^t] = \mathbb{E}^{\lambda[\chi]}[K_{i,t+1}(\tilde{h}_i^t) \mid \tilde{\theta}_i^t], \end{aligned} \quad (38)$$

where $\tilde{h}_i^t = (\tilde{h}_i^{t-1}, \tilde{\theta}_{it}, \tilde{m}_{it}, \chi_{it}(\tilde{m}_{it}, \tilde{\theta}_{-i,t}))$ with $\tilde{m}_{it} = \tilde{\theta}_{it}$.

Now note that, when assumption DNOT holds, the stochastic process $\lambda[\chi]$ over Θ does not depend on χ . Because any truthful private history \tilde{h}_i^t is a deterministic function of $\tilde{\theta}_i^t$ and $\tilde{\theta}_{-i}^t$ and because types are independent we then have that

$$\begin{aligned} \mathbb{E}^\lambda[K_{i,t+1}(\tilde{h}_i^t) \mid \tilde{\theta}_i^t] &= \mathbb{E}^\lambda[K_{i,t+1}(\tilde{h}_i^t) \mid \tilde{\theta}_i^{t+1}] \\ &= \mathbb{E}^\lambda[V^{\Omega_i[\chi, \psi]}(\tilde{\theta}_{i,t+1}, \tilde{h}_i^t) \mid \tilde{\theta}_i^{t+1}] - \mathbb{E}^\lambda[V^{\Omega_i[\chi, \hat{\psi}]}(\tilde{\theta}_{i,t+1}, \tilde{h}_i^t) \mid \tilde{\theta}_i^{t+1}]. \end{aligned} \quad (39)$$

Combining (38) with (39) then gives

$$\begin{aligned} & \mathbb{E}^\lambda[V^{\Omega_i[\chi, \psi]}(\tilde{\theta}_{i,t+1}, \tilde{h}_i^t) \mid \tilde{\theta}_i^{t+1}] - \mathbb{E}^\lambda[V^{\Omega_i[\chi, \hat{\psi}]}(\tilde{\theta}_{i,t+1}, \tilde{h}_i^t) \mid \tilde{\theta}_i^{t+1}] \\ &= \mathbb{E}^{\lambda[\chi]}[V^{\Omega_i[\chi, \psi]}(\tilde{\theta}_{it}, \tilde{h}_i^{t-1}) \mid \tilde{\theta}_i^t] - \mathbb{E}^{\lambda[\chi]}[V^{\Omega_i[\chi, \hat{\psi}]}(\tilde{\theta}_{it}, \tilde{h}_i^{t-1}) \mid \tilde{\theta}_i^t] \end{aligned}$$

Using again the fact that the value function coincides with the equilibrium payoff with probability one then gives the result.

Finally note that, when $N = 1$, \tilde{h}_1^t is a deterministic function of $\tilde{\theta}_1^t$ only. The result in (39) is thus always true when the allocation rule is deterministic. We conclude that, when $N = 1$, the

result in the second part of the proposition holds even if assumption DNOT is dispensed with. ■

Proof of Proposition 12. We show that, under conditions (1)–(4), any allocation rule that is part of a profit-maximizing mechanism must maximize the expected dynamic virtual surplus.

First note that, by Proposition 9, assumption (2) and (4) guarantee that the participation constraints for all types other than the lowest ones can be ignored.

Next note that, because the environment satisfies assumption USEP (i.e. payoffs are time-separable), then an allocation rule maximizes the expected dynamic virtual surplus if and only if, for all t λ -almost all θ^t , $\chi_t(\cdot)$ satisfies condition (26) in the proposition. To prove the result it thus suffices to show that any allocation rule that satisfies condition (26) is implementable in an OEP-IC mechanism that gives zero expected surplus to the lowest types. The result in Proposition 7 then implies that *any* allocation rule that is part of a profit-maximizing mechanism must necessarily maximize the expected dynamic virtual surplus.

As a preliminary step, note that, by inspection, the period- t allocation depends only on the current types θ_t and the first period types θ_1 . Assumptions (1), (3) and (4) then imply that the period- t -state- θ^t virtual surplus has increasing differences in (θ_{i1}, x_{it}) and in (θ_{it}, x_{it}) (for any fixed values of the other arguments). Thus any allocation rule that maximizes the expected dynamic virtual surplus has the property that $\chi_{it}(\cdot)$ is increasing in θ_i^t (in the product order) implying that χ is *strongly monotone*.

Assume now that all agents other than i are truthful. Suppose further that at each period t , before sending his message m_{it} , agent i has observed $(\theta_i^t, \theta_{-i}^T, m_i^{t-1}, x^{t-1})$ (because the other agents are assumed to be truthful we omit the specification of the other agents' messages.) Now consider the allocation rule $\chi_i(\cdot; \theta_{-i}^T)$ that is obtained from χ by fixing the type profile for all agents other than i to θ_{-i}^T . For all θ_{-i}^T , we first construct payments of the form $\psi_i(m_i^T; \theta_{-i}^T) = \sum_{t=2}^T \psi_{it}(m_{i1}, m_{it}; \theta_{-i,1}, \theta_{-i,t})$ that make truthtelling optimal for agent i in all periods $t \geq 2$ and for any period t history. Thus consider an arbitrary period $t \geq 2$. Because at any period $\tau > t$ both $\chi_{i\tau}(\cdot; \theta_{-i}^T)$ and $\psi_{i\tau}(\cdot)$ do not depend on agent i 's message m_{it} in period t and because assumptions DNOT, USEP, and PDPD hold in this environment, then the agent's incentives separate over time. That is, the choice of the optimal message m_{it} depends on $(\theta_i^t, \theta_{-i}^T, m_i^{t-1}, x^{t-1})$ only through $(\theta_{it}, m_{i1}, \theta_{-i,1}, \theta_{-i,t})$. Or, equivalently, agent i 's period- t problem is a static problem indexed by $(m_{i1}, \theta_{-i,1}, \theta_{-i,t})$. Now think of $\chi_{it}(\cdot; m_{i1}, \theta_{-i,1}, \theta_{-i,t})$ as a static allocation rule indexed by $(m_{i1}, \theta_{-i,1}, \theta_{-i,t})$. By strong monotonicity this allocation rule is nondecreasing in m_{it} . Standard results from static mechanism design then guarantee that, when assumption (1) holds, for each $(m_{i1}, \theta_{-i,1}, \theta_{-i,t}) \equiv k$, truthtelling can be made optimal for agent i using payments of the form

$$\psi_{it}(\theta_{it}, k) \equiv u_{it}(\theta_{it}, \chi_{it}(\theta_{it}, k)) - \int_{\theta_{it}}^{\theta_{it}} \frac{\partial u_{it}(r, \chi_{it}(r, k))}{\partial \theta_{it}} dr.$$

Repeating these steps for each period $t \geq 2$ and each agent i , then gives a mechanism $\langle \chi, \psi \rangle$, where ψ is as constructed above, that is OEP-IC at any period- t history, for any $t \geq 2$.

Next, consider period 1. We proved above that there exists a mechanism $\langle \chi, \psi \rangle$ that is OEP-IC at any (possibly non-truthful) period $t \geq 2$ history. Because assumptions DNOT, FOSD, SCP and PDPD hold in this environment, and because χ is strongly monotone, then Corollary 1 implies that there exists a payment rule $\hat{\psi}$ such that $\langle \chi, \hat{\psi} \rangle$ is OEP-IC at any history. The construction of the payments then follows from the proof of that corollary. ■

Proof of Lemma 5. Fix an arbitrary history $\theta \in \Theta$ and let t be the first period such that $\chi_t(\theta^t) = 0$. Then let $s > t$ be the first period after t such that $\chi_s(\theta^s) = 1$. Because there is no learning in periods $t + 1, \dots, s$, the last $s - t$ components of θ^s are necessarily equal to θ_t , the last component of θ^t (that is, $\theta_\tau = \theta_t$ for $\tau = t, t + 1, \dots, s$). Now consider an allocation rule $\hat{\chi}$ such that (1) $\hat{\chi}_t(\theta^t) = \chi_s(\theta^s) = 1$, (2) for any successor θ^τ to θ^t , the behavior of $\hat{\chi}_\tau$ is defined by the behavior of $\chi_{s+(\tau-t)}$ for the analogous successor $\theta^{s+(\tau-t)}$ to θ^s , with $\hat{\chi}_\tau \equiv 0$ if $s + (\tau - t) > T$, and (3) $\hat{\chi}$ agrees with χ elsewhere. Next let $\hat{\psi}$ be the payment scheme that is obtained from ψ following the same construction as for $\hat{\chi}$.

Now note that, because there is no learning during periods of no sales and because there is no discounting, the mechanism $\langle \hat{\chi}, \hat{\psi} \rangle$ leads to the same payoffs as $\langle \chi, \psi \rangle$. Repeating the above construction for all possible histories $\theta \in \Theta$ gives rise to an IR-BIC mechanism $\langle \hat{\chi}, \hat{\psi} \rangle$ such that $\hat{\chi}$ is a stopping rule and the expected payoffs of both the buyer and the seller under $\langle \hat{\chi}, \hat{\psi} \rangle$ are the same as under $\langle \chi, \psi \rangle$. ■

Proof of Proposition 13. Part (1). Consider the efficient allocation rule χ^* . It solves a stopping problem where the period t payoff is $x_t(\theta_t - c)$ with θ_t distributed as above. Let $v_t^*(\theta_t)$ denote the continuation value from period t onwards, which depends only on the current type given the Markov structure. We have

$$v_t^*(\theta_t) = \max \left\{ 0, \theta_t - c + \mathbb{E} \left[v_{t+1}^*(\tilde{\theta}_{t+1}) | \theta_t \right] \right\}. \quad (40)$$

(We are using the conditional expectation notation for convenience; the expectation is actually taken with respect to the kernel identified above.) We proceed by backward induction. At T , for any θ , the efficient allocation $\chi_T^*(\theta)$ solves

$$v_T^*(\theta_T) = \max \{ 0, \theta_T - c \}.$$

Thus χ_T^* has cut-off $z_T^* = c$, which is independent of θ^{T-1} ; by implication, v_T^* is nondecreasing. Suppose then that the properties identified for period T are true for some period $t + 1$ (That is,

χ_{t+1}^* has cut-off z_{t+1}^* independent of θ^t and v_{t+1}^* is nondecreasing). We want to show that the same properties hold in period t .

Given θ^t , $\chi_t^*(\theta^t)$ solves the maximization problem in (40). Since v_{t+1}^* is nondecreasing by the induction hypothesis and we have FOSD, χ_t^* has a cutoff z_t^* which does not depend on θ^{t-1} . Furthermore, v_t^* is nondecreasing.

We conclude that the efficient rule is a cutoff rule where the cutoffs depend only on t . It remains to show that the cutoffs z_t^* are nondecreasing in t . By inspection of (40) it suffices to show that v_t^* is nonincreasing in t . To this end we first establish by backward induction that the functions v_t^* are convex. This is clearly true of v_T^* . Suppose then that v_{t+1}^* is convex. Note that the kernels F_t identified above imply that $\theta_{t+1} = \theta_t + \gamma_t$, where $\tilde{\gamma}_t \sim N(0, \rho_{t+1}^2)$. Note that the distribution of $\tilde{\gamma}_t$ is independent of θ_t . Thus for any θ_t, θ'_t and $\lambda \in [0, 1]$ we have

$$\begin{aligned}
v_t^*(\lambda\theta_t + (1-\lambda)\theta'_t) &= \max \left\{ 0, \lambda\theta_t + (1-\lambda)\theta'_t - c + \mathbb{E} \left[v_{t+1}^*(\tilde{\theta}_{t+1}) | \lambda\theta_t + (1-\lambda)\theta'_t \right] \right\} \\
&= \max \left\{ 0, \lambda\theta_t + (1-\lambda)\theta'_t - c + \mathbb{E} \left[v_{t+1}^*(\lambda\theta_t + (1-\lambda)\theta'_t + \tilde{\gamma}_t) \right] \right\} \\
&= \max \left\{ 0, \lambda\theta_t + (1-\lambda)\theta'_t - c + \mathbb{E} \left[v_{t+1}^*(\lambda(\theta_t + \tilde{\gamma}_t) + (1-\lambda)(\theta'_t + \tilde{\gamma}_t)) \right] \right\} \\
&\leq \max \left\{ 0, \lambda(\theta_t - c + \mathbb{E} [v_{t+1}^*(\theta_t + \tilde{\gamma}_t)]) + (1-\lambda)(\theta'_t - c + \mathbb{E} [v_{t+1}^*(\theta'_t + \tilde{\gamma}_t)]) \right\} \\
&\leq \lambda \max \left\{ 0, \theta_t - c + \mathbb{E} [v_{t+1}^*(\theta_t + \tilde{\gamma}_t)] \right\} \\
&\quad + (1-\lambda) \max \left\{ 0, \theta'_t - c + \mathbb{E} [v_{t+1}^*(\theta'_t + \tilde{\gamma}_t)] \right\} \\
&= \lambda v_t^*(\theta_t) + (1-\lambda)v_t^*(\theta'_t).
\end{aligned}$$

Thus v_t^* is convex. Suppose then that for some t , $v_t^* \geq v_\tau^*$ for all $\tau \geq t$. Note that this holds vacuously for $t = T$. Consider period $t - 1$. For any $a \in \mathbb{R}$,

$$\begin{aligned}
v_{t-1}^*(a) &= \max \left\{ 0, a - c + \mathbb{E} [v_t^*(a + \tilde{\gamma}_t)] \right\} \\
&\geq \max \left\{ 0, a - c + \mathbb{E} [v_{t+1}^*(a + \tilde{\gamma}_t)] \right\} \\
&\geq \max \left\{ 0, a - c + \mathbb{E} [v_{t+1}^*(a + \tilde{\gamma}_{t+1})] \right\} = v_t^*(a),
\end{aligned}$$

where the first equality follows by the induction hypothesis and the second by the convexity of v_{t+1}^* , since the distribution of $\tilde{\gamma}_{t+1}$ second order stochastically dominates that of $\tilde{\gamma}_t$.

Part (2). Next, consider the Relaxed Program. Let $v_t(\theta^t)$ denote the continuation value from period t onwards. We have

$$v_t(\theta^t) = \max \left\{ 0, \theta_t - c - \frac{1}{\eta_1(\theta_1)} + \mathbb{E} [v_{t+1}(\theta^{t+1}) | \theta^t] \right\}. \quad (41)$$

By backward induction one sees that $v_t(\theta^t)$ depends only on (θ_1, θ_t) . Thus the allocation rule χ

that solves the Relaxed Program is an efficient rule in the model parameterized by θ_1 where the seller's cost is $c - \frac{1}{\eta_1(\theta_1)}$. The result in part (1) then implies that χ is a cutoff rule, where the cutoffs $z_t(\theta_1)$ depend only on t and the parameter θ_1 , and are nondecreasing in t . Since the hazard rate $\eta_1(\theta_1)$ is assumed to be monotone, the second term on the right hand side is nondecreasing in θ_1 . This implies that $z_t(\theta_1)$ is nonincreasing in θ_1 .

Part (3). We prove the result by verifying the conditions of Proposition 8. Super- and submodularity (respectively of $u_i(\theta, x)$ and of $\partial u_i(\theta, x) / \partial \theta_{it}$) are satisfied since the payoffs are time-separable and the flow payoffs are linear. By inspection so is SCP. We also have FOSD since θ_t follows a nonstationary random walk. DNOT obtains since given the restriction to stopping rules, for any nontrivial history (i.e., where selling hasn't yet stopped) the distributions depend only on t . Finally, the set of stopping rules is seen to be a lattice as follows: Define the pointwise order on \mathcal{X}^S by setting $\chi \succeq \chi'$ if for all t , all θ^t , $\chi_t(\theta^t) \geq \chi'_t(\theta^t)$. It is then straightforward to verify that the meet and the join of any two stopping rules are stopping rules. The result then follows from Proposition 8.

Part (4). Implementability of each of the two rules follows from Proposition 9 and Corollary 1 since both rules are clearly strongly monotone. Other assumptions are verified as in the proof of part (3). ■