Dynamic Mechanism Design:

Incentive Compatibility, Profit Maximization and Information Disclosure*

Alessandro Pavan
Northwestern University

Ilya Segal
Stanford University

Juuso Toikka
Stanford University

May 8, 2009

Abstract

We examine the design of incentive-compatible screening mechanisms for dynamic environments in which the agents' types follow a (possibly non-Markov) stochastic process, decisions may be made over time and may affect the type process, and payoffs need not be time-separable. We derive a formula for the derivative of an agent's equilibrium payoff with respect to his current type in an incentive-compatible mechanism, which summarizes all first-order conditions for incentive compatibility and generalizes Mirrlees's envelope formula of static mechanism design. We provide conditions on the environment under which this formula must hold in *any* incentive-compatible mechanism. When specialized to quasi-linear environments, this formula yields a dynamic "revenue-equivalence" result and an expression for dynamic virtual surplus, which is instrumental for the design of optimal mechanisms. We also provide some sufficient conditions for incentive compatibility, and for its robustness to an agent's observation of the other agents' past and future types. We apply these results to a number of novel settings, including the design of profit-maximizing auctions and durable-good selling mechanisms for buyers whose values follow an AR(k) process.

JEL Classification Numbers: D82, C73, L1.
Keywords: asymmetric information, stochastic processes, incentives

*For useful suggestions, we thank seminar participants at various institutions where this paper was presented. A special thank is to Li Hao and Narayana Kocherlakota for very detailed comments. Pavan also thanks the hospitality of Collegio Carlo Alberto where part of this work was completed. This paper supersedes 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. Earlier versions contained a section on the independent-shock approach which has now been incorporated in a companion paper (still work in progress). While waiting for a draft, the reader interested in this part can look at Section 3.3 in the November 2008 version posted on our webpages.

1 Introduction

We consider the design of 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 this information on past decisions. 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 to consumers who refine their valuations upon consumption, or the design of multi-period procurement auctions for bidders whose cost parameters evolve stochastically over time and may exhibit learning-by-doing effects.

A 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 misrepresentation 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 past decisions (from which he can make inferences about the other agents' types). Despite the resulting complications, we deliver some general necessary conditions for incentive compatibility and some sufficient conditions, and then use them to characterize optimal (e.g. profit-maximizing) mechanisms in several applications.

The cornerstone of our analysis is the derivation of a formula for the derivative of an agent's equilibrium 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 must hold in *any* incentive-compatible mechanism (not just in "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. In addition to the familiar direct effect of the current type on the agent's utility, the formula also accounts for the current type's impact on the type distributions in each of the future periods, which is both direct and indirect through its impact on the distribution of types in intermediate periods. All these stochastic effects are summarized with a function that can be interpreted as a (nonlinear) impulse response of the future type to the current type. Our dynamic payoff formula adds up the utility effects of all the future types weighted by their impulse responses to the current type. As for the current type's effects through the agent's messages to the mechanism, the formula ignores them, by the usual envelope theorem logic.

For ease of 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 and derive the dynamic payoff formula for this environment. In Section 4 we adapt the dynamic payoff formula to a multi-agent environment, in which an agent may observe only a part of the entire history generated by the mechanism, and must therefore form beliefs about the unobserved parts such as the types of the other agents as well as the unobserved past decisions made by the mechanism. We show that the single-agent analysis 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 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 use the dynamic payoff 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, for a single-agent deterministic mechanism, this theorem pins down, in each state, the total payment that is necessary to implement a given allocation rule, up to a state-independent constant. 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, which we call "Other-Ex Post Incentive Compatibility" (OEP-IC)—according to which the mechanism must remain incentive-compatible even if an agent is shown at the beginning of the game all the other agents' (future) types and randomization outcomes—then the theorem again pins down, for each agent and for each state, the total payment up to a state-independent constant (which may depend on the other agents' types and randomization outcomes).

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 agents' 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 (due to the unlimited bonding possibilities in the quasilinear environment with unbounded transfers), the distortions trade off efficiency for extraction of the agents' first-period information rents. However, due to informational linkages in the stochastic type processes, the principal distorts the agents' consumptions not only in period one, but also in any subsequent period in which his type is responsive to the first-period type, as measured by our impulse response function.

In particular, we find 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 agents' 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 typically non monotonic in the agents' types. 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 satisfactory unless one its solutions can be shown to satisfy all of the remaining incentive and participation constraints. We provide a few sufficient conditions for these constraints to be satisfied. 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 truthtelling), 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 is 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 the 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 show how the aforementioned results can be put to work in a couple of simple, yet illuminating, applications. The first application is a setting where the agents' types follow an autoregressive stochastic process of degree k (AR(k)) and where each agent's payoff is affine in his types (but not necessarily in his consumption). This specification can capture for example auctions with intertemporal capacity constraints, habit formation, and learning-by-doing. In this case, 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 function, which in the case of an AR(k) process is a constant that is independent of the realizations of the process.[1] Intuitively, the impulse response constant 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). 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 constructing an efficient mechanism may in general require solving an involved dynamic programming problem (due to possible intertemporal payoff interactions), once it is constructed it can be easily converted into a solution to the profit-maximizing 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 truthtelling 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 incentive compatibility in period one, we were only able to check it application-by-application, but we have been able to verify it in a few special settings.

The second application is an environment in which the agents' types continue to follow an AR(k) process, but where all agents' preferences are additively time-separable, with arbitrary flow payoffs. 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

[1] The term "handicapped auction" was first used in Eso and Szentes (2007), but in a more special setting (see Section 2).

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 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 rest of the paper is organized as follows. Section 2 briefly discusses some 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. All proofs are in the Appendix at the end of the manuscript.

2 Related Literature

The last few years have witnessed a fast-growing literature on dynamic mechanism design. A number of recent papers propose transfer schemes for implementing efficient (expected surplus-maximizing) mechanisms that generalize static VCG and expected-externality mechanisms (e.g., Bergemann and Välimäki (2007),Athey and Segal (2007), and references therein), but do not provide a general analysis of incentive compatibility in dynamic settings.

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 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 more general payoffs and stochastic structure. Contrary to these early papers, we also provide conditions on the primitive environment that validate the "first-order approach."

Battaglini (2005) derives the optimal selling mechanism for a monopolist facing a single consumer whose type follows a two-state Markov process. Our results for a model with continuous types indicate that many of his predictions are specific to his two-type setting (we discuss this in greater detail in subsection 4.6). Gershkov and Moldovanu (2008a) and Gershkov and Moldovanu (2008b) consider both efficient and profit maximizing mechanisms to allocate objects to buyers that arrive randomly over time. Since each agent in their model lives only instantaneously, their incentive constraints are static. The papers' payoff-equivalence can be viewed as a static result applied separately to each short-lived agent.[2]

[2] Other recent papers that study dynamic profit-maximizing mechanisms include Bognar, Borgers, and Meyer-ter Vehn (2008) and Zhang (2008). The key difference between these papers and ours is that these papers look at

Eso and Szentes (2007) consider a two-period model with many agents but with a single decision in the second period. They propose a novel approach to the characterization of optimal mechanisms, which uses the Probability Integral Transform Theorem (e.g., Angus, 1994) 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 Pavan, Segal, and Toikka (2009), we show how the independent-shock approach can be used to identify conditions for incentive compatibility in infinite-horizon settings. The results in that paper complement those in the current one in that, when applied to finite-horizon models, they permit one to validate the dynamic payoff formula identified in this paper under a different (and not nested) 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, even with many periods the optimal mechanism remains a "handicapped mechanism." The distinguishing feature of such mechanisms is that the allocation in a given period depends only on that period's reports and the first-period reports but not on any intermediate reports.[3]

The paper is also related to the macroeconomic literature on dynamic optimal 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 (2006), Tchistyi (2006), Biais, Mariotti, Plantin, and Rochet (2007), Zhang (2006), 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 some of the features of the optimal mechanism in the context of specific applications.[4]

Dynamic mechanism design is also related to the literature on multidimensional screening, as noted, e.g., in Rochet and Stole (2003). 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, and so he has fewer deviations

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.

[3] Also, while Eso and Szentes use their model to study primarily the effects of the seller's information disclosures on surplus extraction in a special setting, here we focus on the characterization of incentive compatible mechanisms in general dynamic settings.

[4] Some of this work limits its analysis to the characterization of first-order conditions for intetemporal consumption smoothing (the inverse Euler equation), either leaving the dynamics of the optimal mechanism unspecified or solving for it numerically.

available than in the corresponding static environment in which he observes all the dimensions at once. Because of this, the set of implementable allocation rules proves to be significantly larger in a dynamic environment than in the corresponding static multidimensional environment. For example, the profit-maximizing dynamic allocation rules obtained in our applications would not be 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 (2006a) and Calzolari and Pavan (2006b) 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. Pancs (2007) also studies the role of transparency in environments where agents take nonenforceable actions such as investment or information acquisition.

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 $t = 1, 2, \ldots, T$. 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 possible histories of feasible decisions is denoted by $Y \subset \prod_{\tau=1}^{T} Y_\tau$. That Y is a subset of $\prod_{\tau=1}^{T} Y_\tau$ captures the possibility that the decisions that are feasible in period t may depend on the decisions made in previous periods. Given Y, for any t we then let $Y^t \equiv \{y^t \in \prod_{\tau=1}^{t} Y_\tau : (y^t, y_{t+1}, ..., y_T) \in Y$ for some $(y_{t+1}, ..., y_T) \in \prod_{\tau=t+1}^{T} Y_\tau\}$ denote the set of feasible period-t histories of decisions.[5] 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 receives some private information $\theta_t \in \Theta_t \subset \mathbb{R}$. We implicitly endow the set Θ_t with the Borel sigma-algebra. We refer to θ_t as the agent's *current type*. The set of all possible type histories at period t is then 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 past types and on the history of past 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") $F_t\left(\cdot|\theta^{t-1}, y^{t-1}\right)$ on Θ_t,

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

such that $F_t(A|\cdot) : \Theta^{t-1} \times Y^{t-1} \to \mathbb{R}$ is measurable for all measurable $A \subset \Theta_t$.[6] Note that the distribution of θ_t depends only on variables observed by the agent. We denote the collection of kernels by

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

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 the cumulative distribution function (c.d.f.) corresponding to the measure $F_t(\theta^{t-1}, y^{t-1})$.

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

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

Although some form of time separability of U is typically assumed in applications, this is not needed for our results. What is essential is only that the agent's preferences be time consistent, which is captured here by the assumption that the agent is an expected-utility maximizer, with a Bernoulli function that is constant over time.

An often encountered special case in applications is one where private information evolves in a Markovian fashion, and where the agent's payoff is Markovian in the following sense.

Definition 1 *The environment is* Markov *if*

(M1) for all t, and all $(\theta^{t-1}, y^{t-1}) \in \Theta^{t-1} \times Y^{t-1}$, $F_t(\theta^{t-1}, y^{t-1})$ does not depend on θ^{t-2}, and

(M2) there exists functions $\left\langle A_t : \Theta_t \times Y^t \to \mathbb{R}_{++} \right\rangle_{t=1}^{T-1}$ and $\left\langle B_t : \Theta_t \times Y^t \to \mathbb{R} \right\rangle_{t=1}^{T}$ such that for all $(\theta, y) \in \Theta \times Y$,

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

Condition (M1) guarantees that the stochastic process governing the evolution of the agent's type is Markov, while Condition (M2) ensures that in any given period t, after observing history (θ^t, y^{t-1}), 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 $t < T$).

[6] Throughout, we adopt the convention that for any set A, $A^0 \equiv \{\varnothing\}$.

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 "indirect mechanism" can also be induced as an equilibrium outcome of a "direct mechanism" in which the agent is asked to report the current type, and where, in each period, he finds it optimal to report truthfully, conditional on having reported truthfully in the past.

Let $m_t \in \Theta_t$ denote the agent's period-t message, and $m^t \equiv (m_1, \ldots m_t) \in \Theta^t$ denote a collection of messages, from period one to period t.

Definition 2 *A* direct mechanism *is a collection*

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

such that (i) for all t*, all measurable* $A \subset Y_t$*,* $\Omega_t(A|\cdot) : \Theta^t \times Y^{t-1} \to [0,1]$ *is measurable, and (ii) for all* t*, all* $(\theta^t, y^{t-1}) \in \Theta^t \times Y^{t-1}$*,* $y_t \in Supp[\Omega_t(\cdot|\theta^t, y^{t-1})] \implies (y^{t-1}, y_t) \in Y^t$.

The notation $\Omega_t(A|m^t, y^{t-1})$ stands for the probability that the mechanism generates $y_t \in A \subset Y_t$ given history $(m^t, y^{t-1}) \in \Theta^t \times Y^{t-1}$, while $Supp[\Omega_t(\cdot|\theta^t, y^{t-1})]$ denotes the support of the distribution $\Omega_t(\cdot|\theta^t, y^{t-1})$: The requirement that, for any $y_t \in Supp[\Omega_t(\cdot|\theta^t, y^{t-1})]$, $(y_t, y^{t-1}) \in Y^t$ guarantees that the decisions taken in period t are feasible given past decisions 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 from $F_t\left(\cdot|\theta^{t-1}, y^{t-1}\right)$.

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

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

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 identifies a unique strategy; such a strategy has the property that the agent reports his current type truthfully after any history, including non-truthful ones. Note that we are not claiming it is safe to restrict attention to mechanisms with the property that the truthful strategy (as defined above) is optimal at all histories. As explained above, the Revelation Principle in fact simply guarantees that it is safe to restrict attention to mechanisms in which the agent finds it optimal to report truthfully conditional on having reported truthfully in the past; this is equivalent to requiring that the truthful strategy (as defined above) be optimal at all *truthful* histories.[7]

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

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

where by convention $H_0 = \{\varnothing\}$, 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 $m^t = \theta^t$.

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 $(\hat{\theta}, \hat{m}, \hat{y})$ such that $(\hat{\theta}^s, \hat{m}^t, \hat{y}^u) = (\theta^s, m^t, y^u)$. 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.

[7] One can safely restrict attention to mechanisms in which the agent finds it optimal to report truthfully at *any* history, provided in each period t, the agent is asked to report his *complete history* θ^t as opposed to the new information θ_t. This alternative class of direct mechanisms was proposed by Doepke and Townsend (2006). While these mechanisms permit one to restrict attention to one-stage deviations from truthtelling, the deviations that one must consider are multidimensional and contingent on possibly inconsistent reporting histories. Whether this alternative class of mechanisms facilitates the characterization of implementable outcomes is thus unclear.

Given this notation, $\mathbb{E}^{\mu[\Omega,\sigma]|h}[U(\tilde{\theta},\tilde{y})]$ denotes the agent's expected payoff at the history h when he plays according to the strategy σ in the future.[8]

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,[9]

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

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* $V^{\Omega}: H \to \mathbb{R}$ *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 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'*.*

[8] Throughout we use "tildes" to denote random variables with the same symbol without the tilde corresponding to a particular realization.

[9] Restricting attention to pure strategies is without loss: By the Revelation Principle the agent can be assumed to report truthfully on the equilibrium path. 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

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 expected utility from truthtelling following a truthful history $h = (\theta^t, \theta^{t-1}, y^{t-1})$ is

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

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 then replicates the argument from the static case. That is, consider the agent's problem of choosing a continuation strategy given the truthful history $(\theta^t, \theta^{t-1}, y^{t-1})$. Assuming that an envelope argument applies, one would then 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 affects 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 expected 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 differentiability of the value function at future histories can not be taken for granted so that the actual formal argument is more involved. (See the discussion after Proposition 1.) Furthermore, 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, in *any* IC mechanism, the value function is sufficiently well behaved.

Our derivation makes use of the following assumptions.

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

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

Assumption 3 *For all t, and all $(\theta^{t-1}, y^{t-1}) \in \Theta^{t-1} \times Y^{t-1}$, the c.d.f. $F_t(\cdot|\theta^{t-1}, y^{t-1})$ 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 an integrable function $B_t : \Theta_t \to \mathbb{R}_+ \cup \{+\infty\}$ such that for all $\tau < t$, and all $(\theta^t, y^{t-1}) \in \Theta^t \times Y^{t-1}$,*

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

Assumption 6 *For all t, and all $y^{t-1} \in Y^{t-1}$, the probability measure $F_t\left(\cdot|\theta^{t-1}, y^{t-1}\right)$ is continuous in θ^{t-1} in the total variation metric.*[10]

Assumptions 1 and 4 are familiar from static settings (see, e.g., Milgrom and Segal, 2002). The assumption that each Θ_t is open permits us (a) to accommodate the possibility of an unbounded

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

support (e.g. $\Theta_t = \mathbb{R}$), and (b) to avoid qualifications about left and right derivatives at the boundaries. All subsequent results extend to the case that Θ_t is closed. Assumptions 2 and 3 are also typically made in static models. Assumption 2 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 of *feasible* types be connected, Assumption 3 requires that the set of *relevant* types also be connected. The assumption that the support Θ_t is history-independent simplifies some of the proofs but is not essential; one may for example accommodate the case that θ_t follows an ARIMA process with bounded noise, in which case Assumption 3 is violated (see e.g. Pavan, Segal, and Toikka (2009)).[11]

Assumption 5 requires that the distribution of the current types depends 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. The combination of Assumptions 4 and 5 guarantees that the agent's expected payoff depends smoothly on types.[12] Since Assumption 5 does not have a 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, $\frac{\partial}{\partial \theta_\tau}\mathbb{E}[\tilde{\theta}_t|\theta^{t-1}, y^{t-1}]$, 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 settings Assumption 5 by itself does not impose enough smoothness on the dependence of the kernels on past types, which is why we impose also Assumption 6. Note that when the functions $F_t\left(\cdot|\theta^{t-1}, y^{t-1}\right)$ are absolutely continuous, this assumption is equivalent to the continuity of their densities in θ^{t-1} in the L^1 metric.

We are now ready to state our first main result.

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

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

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

[11] Depending on the notion of IC, the assumption of connectedness may also be dropped. This is the case, e.g., if one requires IC at all possible θ_t. Without connectedness, the interpretation however 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 continuum of types in which IC is imposed for all $\theta_1 \in [0, 1]$, or a model with two types with IC imposed only on $\theta_1 \in \{0, 1\}$?

[12] A weaker joint (or "reduced form") assumption imposing restrictions directly on the expected payoff as opposed to U and F separately would also suffice, although need not be easier to verify.

The recursive formula (1) pins down how the agent's equilibrium utility varies as a function of the current type θ_t. It is a dynamic generalization of Mirrlees's static envelope theorem formula (Mirrlees, 1971) (which obtains as a special case when $T = t = 1$). As suggested in the heuristic derivation preceding 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.

The proof, which is in the appendix, is by backward induction, rolling backwards an envelope result at each stage. The argument is not trivial because, contrary to the static case, the continuation payoff cannot be assumed to be differentiable in the current type. This is because the continuation payoff is itself a recursive formulation of the agent's future optimizations. Hence the standard version of the envelope theorem does not apply. However, we show that, at each history, the value function is Lipschitz continuous in the current type θ_t, and thus has left- and right-side derivatives with respect to θ_t everywhere. We then prove an envelope theorem similar to those in Milgrom and Segal (2002) and in Ely (2001) that relates the side derivatives of the value function to the side derivatives of the objective function (the continuation payoff) evaluated at the optimum. The backward induction in the proof uses this envelope theorem to roll backward the side derivatives. A final complication arises from the fact that, as noted by Ely (2001), an envelope theorem based on side derivatives yields payoff equivalence only if the possible "kinks" in the objective function are "convex" (i.e., open upwards). However, it turns out that whenever a maximizer exists, the value function has convex kinks. Thus, along the equilibrium path, where truthtelling is optimal by incentive compatibility (with probability one), any kinks in the value function must indeed be convex.

Remark 1 *While we have restricted the agent to observe a one-dimensional type in each period, the same necessary condition (1) for incentive compatibility can be applied to a model in which the agent observes a multidimensional type in each period, by restricting the agent to observe one dimension of his current type at a time and report it before observing the subsequent dimensions. Indeed, since this restriction only reduces the set of possible deviations, it preserves incentive compatibility, and so condition (1) must still hold. However, incentive compatibility is harder to ensure when the agent observes several dimensions at once (see Remark 2 for more detail).*

3.2.1 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 $\left(\theta^{t-1}, y^{t-1}\right) \in \Theta^{t-1} \times Y^{t-1}$, the function $F_t\left(\cdot|\theta^{t-1}, y^{t-1}\right)$ is absolutely continuous with density $f_t\left(\theta_t|\theta^{t-1}, y^{t-1}\right) > 0$ for a.e. $\theta_t \in \Theta_t$.*

The existence of a strictly positive density allows us to write the formula in (1) terms of expectation operators rather than integrals. Using iterated expectations then yields the following result.

Proposition 2 *Suppose Assumptions 1-7 hold. (If the environment is Markov, then 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*

$$\begin{aligned} &V^{\Omega}(\theta_t, h^{t-1}) \text{ is Lipschitz continuous in } \theta_t \text{, and for a.e. } \theta_t, \\ &\frac{\partial V^{\Omega}(\theta_t, h^{t-1})}{\partial \theta_t} = \mathbb{E}^{\mu[\Omega]|(\theta_t, h^{t-1})}\left[\sum_{\tau=t}^{T} J_t^{\tau}(\tilde{\theta}^{\tau}, \tilde{y}^{\tau-1}) \frac{\partial U(\tilde{\theta}, \tilde{y})}{\partial \theta_\tau}\right], \end{aligned} \tag{2}$$

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

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

with

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

The intuition for the formula in (2) is as follows. The term I_l^m can be interpreted as the "direct response" of signal θ_l to a small change in signal θ_m, $m > l$. The term J_t^{τ} can then be interpreted as the "total" impulse response of θ_t to θ_τ, $\tau > t$. It incorporates all the ways in which θ_t can affect the distribution of θ_τ, both directly and through its effect on the intermediate signals observed by the agent. The calculation of J_t^{τ} counts all possible chains of such effects. For example, in the Markov case, $I_l^m = 0$ for $l < m - 1$, hence we have $J_t^{\tau}\left(\theta^{\tau}, y^{\tau-1}\right) = \prod_{k=t+1}^{\tau} I_{k-1}^{k}(\theta^k, y^{k-1})$ – there is only one chain of effects, which goes through all the periods.

Example 1 *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 $j = 1, ..., k$, and where ε_t is the realization of the random variable $\tilde{\varepsilon}_t$ distributed according to some absolutely continuous c.d.f. G_t with strictly positive density over $\mathbb{R}$ if $t \geq 2$ and over a connected set $\Theta_1 \subset \mathbb{R}$ if $t = 1$, with $(\tilde{\varepsilon}_s)_{s=1}^{T}$ jointly independent. For convenience, hereafter we let $\phi_j \equiv 0$ for all $j > k$. Then

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

so that for any $\tau > t$,

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

and

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

Thus in this case the impulse response $J_t^\tau\left(\theta^\tau, y^{\tau-1}\right)$ is a constant that does not depend on the realizations of (θ, y) (but may still depend on t and τ). In the special case of an AR(1) process we have

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

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

3.3 Sufficient conditions for IC

While the formula in (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 a multidimensional mechanism design problem, for which the characterization of incentive compatibility is more difficult (see, e.g., Rochet and Stole, 2003). More precisely, there are two sources of difficulty in ensuring incentive compatibility in a dynamic setting: (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;[13] 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 *complete 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 truthtelling is an optimal continuation strategy even following deviations, and so incentive compatibility can be assured by checking one-stage 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 always necessary for IC, is sufficient and is easy to check in applications.

Proposition 3 *Suppose the environment satisfies the assumptions of Proposition 2. 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,

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

(ii) For any m_t, for a.e. θ_t,

$$\left[D^{\Omega}\left((\theta^{t-1},\theta_t),(\theta^{t-1},\theta_t),y^{t-1}\right) - D^{\Omega}\left((\theta^{t-1},\theta_t),(\theta^{t-1},m_t),y^{t-1}\right)\right]\cdot(\theta_t - m_t) \geq 0, \tag{3}$$

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

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

Proposition 2 implies that condition (i) in Proposition 3 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

[13] Note that the difficulty of controlling for multi-stage deviations is something one must deal with even if one considers the alternative class of direct mechanisms proposed by Doepke and Townsend (2006) in which the agent reports his complete history θ^t in each period t, as opposed to the new information θ_t.

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 is satisfied in some applications. As a 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 backward starting from period $T+1$ to establish IC in all periods and at all histories.

Proposition 3 can be generalized to establish the optimality of an arbitrary strategy σ at any arbitrary history h. To this purpose, it suffices to consider a fictitious mechanism $\hat{\Omega}$ that responds to the agent's reports about his type $\hat{\theta}_t$ with a recommendation $\sigma_t(h_{t-1}, \theta_t)$ about the message to send in the original mechanism Ω. Provided the support of $\sigma_t(h_{t-1}, \cdot)$ coincides with the entire set Θ_t—which guarantees that any deviation in period t from $\sigma_t(h_{t-1}, \theta_t)$ can be interpreted as a misrepresentation of θ_t to the new mechanism $\hat{\Omega}$—then the optimality of σ at $h_t = (h_{t-1}, \theta_t)$ can be verified by checking a single-crossing condition similar to the one that in (3) establishes optimality of a truthful strategy at a truthful history.

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 $i = 1, \ldots, N$. In each period $t = 1, \ldots, T$, each agent i is shown a nonmonetary "allocation" $x_{it} \in X_{it}$, where X_{it} is a measurable space, and asked to make a payment $p_{it} \in \mathbb{R}$. The set of all feasible histories of joint allocations is denoted by $X \subset \prod_{t=1}^{T} \prod_{i=1}^{N} X_{it}$. Given X, we then let $X_t \subset \prod_{i=1}^{N} X_{it}$ denote the set of period-t feasible allocations, $X_{-i,t} \subset \prod_{j \in \{1,\ldots,N\}, j \neq i} X_{jt}$ the set of period-t allocations for all agents other than i, and $X_i^t \subset \prod_{s=1}^t X_{is}$, $X^t \subset \prod_{s=1}^t \prod_{i=1}^N X_{is}$, and $X_{-i}^t \subset \prod_{s=1}^t \prod_{j \in \{1,\ldots,N\}, j \neq i} X_{js}$ the corresponding sets of period-t histories. This formulation allows for the possibility that the set of feasible allocations in each period depend on the allocations in the previous periods, and/or that the set of feasible decisions with each agent depends on the decisions taken with the other agents.[14,15]

Each agent i observes his own allocations x_i^T but not the other agents' allocations x_{-i}^T. The observability of x_{it} should be thought of as a technological restriction: in each period, a mechanism can reveal more information to agent i than x_{it}, but it cannot conceal x_{it}. As for the payments, because the results do not hinge 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, $\left(x^t, p^t\right)$ is an element of $X^t \times \mathbb{R}^{Nt}$.

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 $\theta_t \equiv (\theta_{1t}, \ldots, \theta_{Nt}) \in \Theta_t \equiv \prod_i \Theta_{it}$. 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 $x_i^T \equiv (x_{i1}, ..., x_{iT})$, $p_i^T \equiv (p_{i1}, ..., p_{iT})$, and $\theta_i^T \equiv (\theta_{i1}, ..., \theta_{iT})$ (and the sets they are elements of). This is to avoid confusion between, e.g., $x_t \equiv (x_{1t}, ..., x_{Nt})$ and $x_i \equiv (x_{i1}, ..., x_{iT})$.

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

We then define a quasi-linear environment as follows.

Definition 7 *The environment is* quasilinear *if the following hold:*

1. *There is a sequence of decisions $(x, p) \in X \times \mathbb{R}^{NT}$, where $x = \left(x_1^T, \ldots, x_N^T\right)$ 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 the kernels $\left\langle F_t : \Theta^{t-1} \times X^{t-1} \to \Delta(\Theta_t) \right\rangle_{t=1}^T$.*

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

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

[14] For example, the (intertemporal) allocation of a private good in fixed supply $\bar{x}$ can be modelled by letting $X = \{x \in \mathbb{R}_+^{tN} : \sum_{it} x_{it} \leq \bar{x}\}$, while the provision of a public good whose period-t production is independent of the level of production in any other period can be modelled by letting $X = \prod_{t=1}^T X_t$ with $X_t = \left\{x_t \in \mathbb{R}_+^N : x_{1t} = x_{2t} = ... = x_{Nt}\right\}$.

[15] This formulation does not explicitly allow for decisions that are not observed by any agent at the time they are made; however, such decisions can easily be accomodated by introducing a fictitious agent observing them. In this case, one can also interpret x_{it} as the "signal" that agent i receives in period t about the unobservable decision.

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

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.[16]

Definition 8 *We have* Independent Types *if, for all* t*, 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*, all* t*, all* $(\theta_i^{t-1}, x_i^{t-1}) \in \Theta_i^{t-1} \times X_i^{t-1}$*,* $F_{it}(\cdot|\theta_i^{t-1}, x_i^{t-1})$ *is a c.d.f. on* Θ_{it}*.*

This definition is the proper extension of the Independent-Type assumption of static mechanism design to the dynamic settings considered here; it permits us to extend such static results as revenue-equivalence and the 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 history of private decisions 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 assumptions (i) or (ii) are not satisfied, then a mechanism similar to the one proposed by Cremer and McLean (1988) could be used to fully extract the agents' surplus. It turns out that, if assumption (iii) is not satisfied, then a similar full surplus extraction is possible by using a randomized mechanism—see the discussion after Proposition 4 below.

Throughout this section we will maintain the assumptions that the environment is quasilinear and that types are independent. To highlight the role of the other assumptions, we will then dispense with such qualification in the subsequent results.

4.2 Multi-agent mechanisms

For most of the analysis we will focus on the Bayesian Nash Equilibria (BNE) of the 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

[16] Some readers may feel that an agent must always be able to observe his own final payoff (indeed, this was the case in the model in Section 3). This is compatible with an interdependent-valuation model in which agent i observes x_{-i}^T and θ_{-i}^T at the end of period T, provided that one assumes that at that moment the game is over, in the sense that the mechanism does not make payments contingent on reports about $u_i(\theta, x)$. If, instead, in an interdependent-valuation model, one were to allow agents to report their final payoffs and condition payments on such information, as in Mezzetti (2004), one would then effectively convert the model into one with correlated private observations, in which case full surplus extraction is possible.

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).[17] 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 9 *A* deterministic direct mechanism *is a pair* $\langle \chi, \psi \rangle$*, where* $\chi = \left\langle \chi_t : \Theta^t \to X_t \right\rangle_{t=1}^T$ *with* $\chi(\theta) \in X$ *for all* $\theta \in \Theta$ *is an* allocation rule*, and* $\psi : \Theta \to \mathbb{R}^N$ *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 $m^{t-1} = \left(m_1^{t-1}, \ldots, m_N^{t-1}\right)$ 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 $\chi_t\left(m^t\right)$.

4. Each agent i observes $\chi_{it}\left(m^t\right)$.

[17] In 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 $\psi\left(m^T\right)$.

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

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

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 $\left(\theta, m_i^T\right) \in \Theta \times \Theta_i^T$, 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 10 *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 of Section 3. To see this, fix an arbitrary agent i. Given any deterministic mechanism $\langle \chi, \psi \rangle$, when all agents other than i (henceforth denoted by $-i$) 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 $Y_i^t = X_i^t$, while for $t = T$, let $Y_{iT} = X_{iT} \times \prod_{j \neq i} \prod_{s=1}^T X_{js} \times \Theta_{-i}^T$ and then let $Y_i^T \subset \prod_{j=1}^N \prod_{s=1}^T X_{js} \times \Theta_{-i}^T$ be such that, for each $\left((x_i^T, x_{-i}^T), \theta_{-i}^T\right) \in Y_i^T$, $\left(x_i^T, x_{-i}^T\right)^{\nabla} \in X$, where $\left(x_i^T, x_{-i}^T\right)^{\nabla}$ is simply the reorganization of (x_i^T, x_{-i}^T)

given by $(x_{11},...,x_{N1},...,x_{1T},...,x_{NT})$. Finally, let $Y_{i,T+1}=\mathbb{R}$ and $Y_i^{T+1}=Y_i^T\times\mathbb{R}$. That is, in periods $t<T$, the decision $y_{it}=x_{it}$ consists of the part of the allocation observed by agent i and Y_i^T consists of the set of (feasible) histories of private decisions for 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); the set Y_i^T is then constructed to contain only histories of feasible decisions. Finally, 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 $F_i=\left\langle F_{it}:\Theta_i^{t-1}\times X_i^{t-1}\to\Delta(\Theta_{it})\right\rangle_{t=1}^T=\left\langle F_{it}:\Theta_i^{t-1}\times Y_i^{t-1}\to\Delta(\Theta_{it})\right\rangle_{t=1}^T$, where the equality is by definition of Y_i^t. There is no type in period $T+1$ so that, 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 $Y_i^{T+1}\subset\prod_{t=1}^{T+1}Y_{it}$. However, by construction $\Theta_i^T\times Y_i^{T+1}$ is simply a reordering of $\Theta\times X\times\mathbb{R}$, the domain of agent i's payoff function U_i in the multi-agent environment.

Agent i faces a randomized mechanism $\Omega_i=\Omega_i[\chi,\psi]\equiv\left\langle\Omega_{it}:\Theta_i^t\times Y_i^{t-1}\to\Delta(Y_{it})\right\rangle_{t=1}^{T+1}$ constructed as follows. Let $H_i\equiv\left\{(\theta_i^s,m_i^t,y_i^u)\in\Theta_i^s\times\Theta_i^t\times Y_i^u:T+1\geq s\geq t\geq u\geq s-1\geq-1\right\}$ denote the set of agent i's private histories and $H_i(\chi)\subset H_i$ denote the subset of H_i that is *feasible* given χ. Formally, $H_i(\chi)\equiv\left\{(\theta_i^s,m_i^t,y_i^u)\in H_i:y_i^u=\chi_i^u(m_i^u,m_{-i}^u)\text{ for some }m_{-i}^u\in\Theta_{-i}^u\right\}$ is simply the collection of private histories with the property that, given agent i's reports, the decisions observed by agent i can be obtained as the result of a report by the other agents.

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.[18] Fix $t\leq T$. Suppose that, for all periods $\tau\leq t$, all $m_i^{\tau-1}$, a rcpd $\Gamma_{i,\tau-1}^{\chi}(\cdot|\chi_i^{\tau-1}(m_i^{\tau-1},\tilde{\theta}_{-i}^{\tau-1}))$ on $\Theta_{-i}^{\tau-1}$ exists with the property that, for any $y_i^{\tau-1}=\chi_i^{\tau-1}(m_i^{\tau-1},\theta_{-i}^{\tau-1})$, $\theta_{-i}^{\tau-1}\in\Theta_{-i}^{\tau-1}$, $\Gamma_{i,\tau-1}^{\chi}(\cdot|\chi_i^{\tau-1}(m_i^{\tau-1},\tilde{\theta}_{-i}^{\tau-1})=y_i^{\tau-1})$ assigns measure one to the set $\{\theta_{-i}^{\tau-1}\in\Theta_{-i}^{\tau-1}:\chi_i^{\tau-1}(m_i^{\tau-1},\theta_{-i}^{\tau-1})=y_i^{\tau-1}\}$. Note that the conditioning here is on the random variable $\chi_i^{\tau-1}(m_i^{\tau-1},\tilde{\theta}_{-i}^{\tau-1})$ taking values in $Y_i^{\tau-1}$. The assumption clearly holds vacuously for $t=1$. Now, for all m_i^t, the rcpd $\Gamma_{i,t-1}^{\chi}(\cdot|\chi_i^{t-1}(m_i^{t-1},\tilde{\theta}_{-i}^{t-1}))$ 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.[19] Since $\Theta_{-i}^t\subset\mathbb{R}^{(N-1)t}$, a rcpd of $\tilde{\theta}_{-i}^t$ given $\Sigma(\chi_i^t(m_i^t,\tilde{\theta}_{-i}^t))$ always exists, 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). Furthermore, it is easy to see that there also always exists a rcpd with the property that, for any $y_i^\tau=\chi_i^\tau(m_i^\tau,\theta_{-i}^\tau)$, $\theta_{-i}^\tau\in\Theta_{-i}^\tau$, $\Gamma_{i,\tau}^{\chi}(\cdot|\chi_i^\tau(m_i^\tau,\tilde{\theta}_{-i}^\tau)=y_i^\tau)$ assigns measure one to the set $\{\theta_{-i}^\tau\in\Theta_{-i}^\tau:\chi_i^\tau(m_i^\tau,\theta_{-i}^\tau)=y_i^\tau\}$. We then let $\Gamma_{i,t}^{\chi}(\cdot|\chi_i^t(m_i^t,\tilde{\theta}_{-i}^t))$ be such 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 $T+1$.

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

[19] Formally, $F_{-i,t}(\cdot|\theta_{-i}^{t-1},\chi_{-i}^{t-1}(m_i^{t-1},\theta_{-i}^{t-1}))\equiv\prod_{j\neq i}F_{jt}(\cdot|\theta_j^{t-1},\chi_j^{t-1}(m_i^{t-1},\theta_{-i}^{t-1}))$.

Given the system of rcpd $\Gamma_i^\chi \equiv \langle \Gamma_{i,t-1}^\chi : \Theta_i^{t-1} \times Y_i^{t-1} \to \Delta(\Theta_{-i}^{t-1}) \rangle_{t=2}^T$, in each period t, agent i's beliefs over $(\theta_{-i}^t, m_{-i}^t, x_{-i}^{t-1})$ at any feasible private history $h_{it} = (\theta_i^t, m_i^t, y_i^t)$ assign probability one to $(\theta_{-i}^t, m_{-i}^t, x_{-i}^t) \in \Theta_{-i}^t \times \Theta_{-i}^t \times X_{-i}^t$ such that $m_{-i}^t = \theta_{-i}^t$ and $x_{-i}^t = \chi_{-i}^t(m_i^t, \theta_{-i}^t)$.

The randomized mechanism $\Omega_i = \Omega_i[\chi, \psi]$ is then derived as follows. Let $t < T$ and fix a history $(\theta_i^t, m_i^t, y_i^{t-1})$. Then for any measurable $A \subset Y_{it}$, the probability that $y_{it} \in A$ is

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

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

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

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 $\Omega_i = \Omega_i[\chi, \psi]$.

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.

Now, a strategy σ_i together with a private history $h_i \in H_i$ induce a probability measure $\mu_i[\Omega_i, \sigma_i]|h_i$ on $\Theta_i^T \times \Theta_i^T \times Y_i^{T+1}$. 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 $\mathbb{E}^{\mu_i[\chi,\psi]|h_i}[U_i(\tilde{\theta}_i^T, \tilde{y}_i^{T+1})] = \mathbb{E}^{\mu_i[\chi,\psi]|h_i}[U_i(\tilde{\theta}, \tilde{x}, \tilde{p}_i^T)]$, where the equality is by definition of y_i^{T+1}. We can then define the value function $V_i^{\Omega_i[\chi,\psi]} : H_i \to \mathbb{R}$ and incentive compatibility at a private history h_i analogously to the single-agent case.

It will be convenient to let $\mu_i^T[\chi, \sigma_i]|h_i$ denote the marginal of $\mu_i[\langle \chi, \psi \rangle, \sigma_i]|h_i$ on $\Theta_i^T \times \Theta_i^T \times Y_i^T$ given private history h_i, with $\mu_i^T[\chi]|h_i$ in case σ_i is the truthful strategy. Thus, $\mu_i^T[\chi, \sigma_i]|h_i$ is a probability measure 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 under a truthful strategy 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 hold for any i. We then have that in any mechanism that is IC for agent i at a feasible truthful private history $h_i^{t-1} = \left(\theta_i^{t-1}, \theta_i^{t-1}, x_i^{t-1}\right)$, the derivative of the value function with respect to θ_{it} does not depend on the payment scheme. Under the assumptions of Proposition 1, this can be seen by iterating (1) backward starting from $t = T$.

In a quasi-linear environment, the aforementioned proposition thus implies that, in any ex-ante BIC mechanism, the value function of each agent i at $\lambda[\chi]$-almost every truthful private history $h_i^t = \left(\theta_{it}, h_i^{t-1}\right)$, $t \geq 1$, 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}. This in turn implies that the "innovation" $\mathbb{E}^{\lambda[\chi]}[\psi_i(\tilde{\theta})|\tilde{\theta}_{it}, \tilde{h}_i^{t-1}] - \mathbb{E}^{\lambda[\chi]}[\psi_i(\tilde{\theta})|\tilde{h}_i^{t-1}]$ in the expected transfer of each agent i due to his own type θ_{it} is the same in any two ex-ante BIC deterministic mechanisms $\langle \chi, \psi \rangle$ and $\langle \chi, \hat{\psi} \rangle$ implementing the same allocation rule.[20]

Using the law of iterated expectations, one can also get rid of the dependence of the constant $k_i(h_i^{t-1})$ on the history h_i^{t-1}. To see this, suppose there is a single agent i and assume, for simplicity, that there are only two periods. Now 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 in period two, for $\lambda[\chi]$-almost every truthful history $h_i^1 = (\theta_{i1}, \theta_{i1}, \chi(\theta_{i1}))$, there exists a scalar $\kappa_i(h_i^1) = K_i(\theta_{i1})$ such that, for any θ_{i2}, $V^{\Omega_i[\chi,\psi]}\left(\theta_{i2}, h_i^1\right) - V^{\Omega_i[\chi,\hat{\psi}]}\left(\theta_{i2}, h_i^1\right) = K_i(\theta_{i1})$. 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 $V^{\Omega_i[\chi,\psi]}\left(\theta_{i2}, h_i^1\right) - V^{\Omega_i[\chi,\hat{\psi}]}\left(\theta_{i2}, h_i^1\right)$, 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 $u(\tilde{\theta}^T, \chi(\tilde{\theta}^T))$ and the expectation of $\psi(\tilde{\theta}^T)$, we have that the entire payment scheme ψ is uniquely determined by the allocation rule χ up to a scalar.

Next, consider a setting with multiple agents. Provided that types are independent, then the total payment that each agent i 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 8 (DNOT) *Decisions do Not Affect Types: For all* $i = 1, ..., N$, $t = 2, ..., T$, $\theta_i^{t-1} \in \Theta_i^{t-1}$, *the distribution* $F_{it}\left(\cdot|\theta_i^{t-1}, x_i^{t-1}\right)$ *does not depend on* x_i^{t-1}.

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

We then have the following result.

Proposition 4 *Suppose that, for each $i = 1, ..., N$, the assumptions of Proposition 1 hold. Consider any two ex-ante BIC deterministic mechanisms $\langle \chi, \psi \rangle$ and $\langle \chi, \hat{\psi} \rangle$ implementing the same allocation rule χ.*

(i) Then for all i, there exists a $K_i \in \mathbb{R}$ such that

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

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

$$\mathbb{E}^{\lambda}[\psi_i(\tilde{\theta}) \mid \tilde{\theta}_i^t] - \mathbb{E}^{\lambda}[\hat{\psi}_i(\tilde{\theta}) \mid \tilde{\theta}_i^t] = \mathbb{E}^{\lambda}[\psi_i(\tilde{\theta}) \mid \tilde{\theta}_i^s] - \mathbb{E}^{\lambda}[\hat{\psi}_i(\tilde{\theta}) \mid \tilde{\theta}_i^s]. \tag{4}$$

The value of Proposition 4 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 the principal's expected payoff as expected virtual surplus, as illustrated below. Both (a) and (b) use the result of Proposition 4 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 (4) also turns useful in certain applications.

Note also that the result in Proposition 4 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 the game, before learning his period-one type, the entire profile of the other agents' types θ_{-i}^T. Then a simple iterated expectation argument similar to the one sketched above implies that, for each agent i, payments are uniquely determined not only in expectation but for each state $(\theta_i^T, \theta_{-i}^T)$: given any pair of ex-ante BIC deterministic mechanisms $\langle \chi, \psi \rangle$ and $\left\langle \chi, \hat{\psi} \right\rangle$ implementing the same allocation rule, for any i there exists a scalar $K_i(\theta_{-i}^T)$ such that $\psi_i(\theta_i^T, \theta_{-i}^T) - \hat{\psi}_i(\theta_i^T, \theta_{-i}^T) = K_i(\theta_{-i}^T)$ for any θ_{-i}^T. (We provide sufficient conditions for the resulting mechanism to satisfy this robustness to information leakage in Proposition 9 below.) This strengthening is the dynamic extension of the static result whereby going from BIC to dominant-strategy incentive compatibility pins down the payment of each type of each agent up to a function of the other agents' types.

Lastly, note that a key assumption in Proposition 4 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 things 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 that 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 optimal mechanisms

In a static setting, the envelope formula permits one to calculate the agents' information rents, providing a useful tool for designing optimal 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 optimal 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 $u_0 : \Theta \times X \to \mathbb{R}$. 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.[21] 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.[22]

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

[21] If 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.

[22] While our definition of IR-BIC mechanism requires that *almost all* types θ_1 find it optimal to participate and that for *almost all* θ_1 the value function coincides with the equilibrium payoff under truthtelling (by all agents), in many applications it is simple to guarantee that participation and truthful reporting be optimal for *all* θ_1.

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.[23] While the possibility of bonding clearly simplifies the analysis, in many applications of interest, participation can be guaranteed in each period even without asking the agent to post bonds, by choosing an appropriate distribution of the transfers over time.

Hereafter, we thus restrict attention to the participation constraints that each agent faces at the moment he is 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

$$V^{\Omega_i[\chi,\psi]}(\theta_{i1}) \geq 0 \quad \text{for all } i, \text{ almost all } \theta_{i1} \in \Theta_{i1}. \tag{5}$$

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

$$\begin{aligned} &V^{\Omega_i[\chi,\psi]}(\theta_{i1}) \text{ is Lipschitz continuous, and for a.e. } \theta_{i1}, \\ &\frac{\partial V^{\Omega_i[\chi,\psi]}(\theta_{i1})}{\partial \theta_{i1}} = \mathbb{E}^{\mu_i^T[\chi]|\theta_{i1}} \left[\sum_{\tau=1}^{T} J_{i1}^{\tau}(\tilde{\theta}_i^{\tau}, \tilde{x}_i^{\tau-1}) \frac{\partial u_i(\tilde{\theta}, \tilde{x})}{\partial \theta_{i\tau}} \right]. \end{aligned} \tag{6}$$

The requirement that $\langle \chi, \psi \rangle$ is ex-ante BIC then implies that, for each $i = 1, ..., N$, agent i's ex-ante equilibrium expected payoff coincides with the expectation of his period-1 value function. Condition (6) can then be used to calculate the agents' expected information rents. Letting

[23] The 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) utilities in the mechanism being bounded from below and 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 harder, but still doable in certain applications.

$\eta_{i1}(\theta_{i1}) \equiv f_{i1}(\theta_{i1})/(1-F_{i1}(\theta_{i1}))$ denote the hazard rate of the distribution F_{i1} and integrating by parts, then gives

$$\begin{aligned}\mathbb{E}^{\lambda[x]}[U_i(\tilde{\theta},\chi(\tilde{\theta}),\psi_i(\tilde{\theta}))] &= \mathbb{E}^{\lambda[x]}[V^{\Omega_i[\chi,\psi]}(\tilde{\theta}_{i1})] \\ &= \mathbb{E}^{\lambda[x]}\left[\frac{1}{\eta_{i1}(\tilde{\theta}_{i1})}\sum_{\tau=1}^{T} J_{i1}^{\tau}(\tilde{\theta}_i^{\tau},\chi_i^{\tau-1}(\tilde{\theta}^{\tau-1}))\frac{\partial u_i(\tilde{\theta},\chi(\tilde{\theta}))}{\partial\theta_{i\tau}}\right] + V^{\Omega_i[\chi,\psi]}(\underline{\theta}_{i1}).\end{aligned} \tag{7}$$

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

$$V^{\Omega_i[\chi,\psi]}(\underline{\theta}_{i1}) \geq 0 \tag{8}$$

Finally, the relaxed program treats the functions $V^{\Omega_i[\chi,\psi]}(\underline{\theta}_{i1})$ in (7) and (8) as control variables that can be chosen independently from (χ,ψ). Formally, the Relaxed Program can be stated as follows.

$$\mathcal{P}^r: \begin{cases} \max\limits_{\chi,\psi,(V^{\Omega_i[\chi,\psi]}(\underline{\theta}_{i1}))_{i=1}^N} & \mathbb{E}^{\lambda[\chi]}[U_0(\tilde{\theta},\chi(\tilde{\theta}),\psi(\tilde{\theta}))] \\ \text{s.t., for all } i=1,...,N, & \text{(7) and (8) hold}\end{cases}$$

Substituting (7) into the principal's payoff then gives the following result.

Lemma 2 *Suppose that, for each $i=1,...,N$, the assumptions of Proposition 2 hold, and $\underline{\theta}_{i1} > -\infty$. Then the principal's expected payoff in any IR-BIC mechanism $\langle\chi,\psi\rangle$ equals*

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

In what follows we will refer to the expression

$$\mathbb{E}^{\lambda[\chi]}\left[\sum_{i=0}^{N} u_i(\tilde{\theta},\chi(\tilde{\theta})) - \sum_{i=1}^{N}\frac{1}{\eta_{i1}(\tilde{\theta}_{i1})}\sum_{t=1}^{T} J_{i1}^{t}(\tilde{\theta},\chi(\tilde{\theta}))\frac{\partial u_i(\tilde{\theta},\chi(\tilde{\theta}))}{\partial\theta_{it}}\right], \tag{9}$$

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.

$$V^{\Omega_i[\chi,\psi]}(\underline{\theta}_{i1}) = 0 \quad \text{for all } i, \tag{10}$$

and that the payment function ψ satisfies (7). 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 5 *Suppose that, for each $i = 1, ..., N$, the assumptions of Proposition 2 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" (9), the lowest types' participation constraints (10) bind, and all the participation constraints (5) 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 (9);*

(iii) the principal's expected payoff cannot be increased using randomized mechanisms.

Of course, Proposition 5 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 first focus on the Relaxed Program and characterize the distortions in the optimal allocation rule relative to the efficient one.

4.6 Distortions

In this subsection we analyze the allocative distortions featured by the solution to the Relaxed Program. First, we show that the classical static result about downward distortions extends to the dynamic model under appropriate assumptions. Second, we show that the other classical static result, that the magnitude of distortion monotonically decreases with type, need not extend to dynamic settings, i.e. one may naturally obtain efficient decisions both for the highest and the lowest type and distortions for intermediate types.

To fix ideas, we start with a special class of environments in which the expected virtual surplus (9) can be maximized simultaneously for all periods and states, obviating the need to solve a dynamic programming problem. This occurs when, in addition to assumption DNOT (which ensures that the stochastic process λ over Θ is exogenous and does not depend on the mechanism), the set of feasible allocations in any period t is independent of the allocations in the previous periods, and payoffs separate over time.

Assumption 9 (DSEP) *Separability in decisions: In addition to DNOT, $X = \prod_{t=1}^{T} X_t$ and, for all $i = 0, ..., N$, all $(\theta, x) \in \Theta \times X$, $u_i(\theta, x) = \sum_{t=1}^{T} u_{it}(\theta^t, x_t)$.*

If DSEP holds, then the Relaxed Program is solved by requiring that for all periods t, λ–almost all θ^t,

$$\chi_t\left(\theta^t\right) \in \arg\max_{x_t \in X_t}\left[\sum_{i=0}^{N} u_{it}\left(\theta^t, x_t\right) - \sum_{i=1}^{N} \frac{1}{\eta_{i1}\left(\theta_{i1}\right)} J_{i1}^{t}\left(\theta_i^t\right) \frac{\partial u_{it}\left(\theta^t, x_t\right)}{\partial \theta_{it}}\right] \tag{11}$$

We compare allocation rules solving (11) to efficient allocation rules χ^*, which, for any period t and λ–almost all θ^t, solve

$$\chi_t^*\left(\theta^t\right) \in \arg\max_{x_t \in X_t}\left[\sum_{i=0}^{N} u_{it}\left(\theta^t, x_t\right)\right]. \tag{12}$$

Comparing the two programs, one can see that distortions are determined by the properties of the impulse responses $J_{i1}^t\left(\theta_i^t\right)$ and of the partial derivative of the flow utility $u_{it}\left(\theta^t, x_t\right)$ with respect to θ_{it}. For example, suppose that, in addition to the aforementioned assumptions, the following holds.

Assumption 10 (FOSD) *First-Order Stochastic Dominance: For all $i = 1, ..., N$, all $t = 2, ..., T$, all $\theta_{it} \in \Theta_{it}$, and all $x_i^{t-1} \in X_i^{t-1}$, $F_{it}\left(\theta_{it}|\theta_i^{t-1}, x_i^{t-1}\right)$ is nonincreasing in θ_i^{t-1}.*

Under this assumption, the impulse responses are nonnegative, i.e. $J_{it}^{\tau}\left(\theta_i^{\tau}, x_i^{t-1}\right) \geq 0$. Comparing programs (11) and (12) in the case of a single agent ($N = 1$) then suggests that in the Relaxed Program the principal distorts x_t to reduce the partial derivative $\partial u_{it}\left(\theta^t, x_t\right)/\partial \theta_{it}$. In the standard case in which x_t is one-dimensional and the agent's flow payoff $u_{it}\left(\theta^t, x_t\right)$ 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. 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 the period-t allocation, by reducing the agent's information rent in period t, then also reduces his information rent in period one, thus raising the principal's expected payoff.

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

Assumption 11 (SCP) *Single Crossing Property: X is a lattice. Furthermore, for each $i = 1, ..., N$, $u_i\left(\theta, x\right)$ has increasing differences in $\left(\theta_i^T, x\right)$.*

The assumption that X is a lattice is not innocuous when we have more than one agent: For example, it holds when each x_t describes the provision of a one-dimensional public good, but it need not hold if x_t describes the allocation of a private good (see footnote 14 above for both examples). The lattice structure of X induces a lattice structure on the set $\mathcal{X}$ of all (measurable) decision rules. We then have the following result.

Proposition 6 *Suppose that, for each $i = 1, ..., N$, the assumptions of Proposition 2 hold, and $\underline{\theta}_{i1} > -\infty$. 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 $i = 0, ..., N$, 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, for all t.

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

Conditions (i) and (ii) hold trivially when DSEP holds and each X_t is a chain (e.g., $X_t \subset \mathbb{R}$). The result in the proposition then means that if χ^0 solves the relaxed problem and χ^* is efficient, then the decision rule $\left(\chi^0 \vee \chi^*\right)_t(\theta) = \chi_t^0(\theta) \vee \chi_t^*(\theta)$ is efficient and the decision rule $\left(\chi^0 \wedge \chi^*\right)_t(\theta) = \chi_t^0(\theta) \wedge \chi_t^*(\theta)$ 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. For an example of an environment that does not satisfy DSEP but where the result nevertheless applies, see the section 5.1.1 on a durable good monopolist.

Next we consider how the magnitude of distortions depends on the types. In the static setting, the classical conclusion is that the optimal mechanism implements efficient decisions for the highest type and creates downward distortions for all the other types. Consider now a dynamic setting with a single agent ($N = 1$) satisfying DSEP, and consider the allocation rule in a period $t > 1$. Suppose that Θ_{1t} is bounded, and that at its bounds the density functions $f_{1t}\left(\underline{\theta}_{1t}|\theta_1^{t-1}\right), f_{1t}\left(\bar{\theta}_{1t}|\theta_1^{t-1}\right) > 0$ for any θ_1^{t-1}. In this case, the impulse response $J_{11}^t\left(\theta_1^t\right) = 0$ when $\theta_{1t} \in \left\{\underline{\theta}_{1t}, \bar{\theta}_{1t}\right\}$, and so programs (11) and (12) coincide.[24] Thus, the optimal mechanism implements an efficient decision not only for the highest but also for the lowest type. To develop some intuition for this result, note that 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 payoffs, distorting the allocations in period t is then useful only to the extent that the type in period t is responsive to the period-1 type. When the agent's type in period t is either the highest or the lowest possible type for that period, it becomes unresponsive to the period-1 type, which eliminates any reason to distort the period-t decision. In a Markov model, in which $J_{11}^t\left(\theta_1^t\right) = \Pi_{\tau=1}^{t-1} I_{1\tau}^{\tau+1}\left(\theta_1^{\tau+1}\right)$, following $\theta_{1t} = \underline{\theta}_{1t}$ or $\theta_{1t} = \overline{\theta}_{1t}$ distortions then vanish also in all subsequent periods, since the responsiveness to the period-1 type is severed.

As stated, the prediction of no distortion for the lowest and the highest type is however of no value for it refers to a zero-probability event, and the optimal allocation rule is defined only with probability 1. However, under appropriate continuity assumptions, the finding of no distortion at the bounds can be extended to show that distortions are small for types that are close enough

[24] The derivatives used in calculating $J_{11}^t\left(\theta_1^t\right)$ at the bounds must be interpreted as the appropriate side derivatives.

to the bounds. Namely, letting $S_t\left(x_t, \theta_1^t\right)$ denote the total surplus in period t from allocation x_t (i.e., the objective function in (12)), under the assumption that $\lim_{\theta_{1t} \to \underline{\theta}_{1t}} J_{11}^t\left(\theta_{1t}, \theta_1^{t-1}\right) = 0$ the optimal mechanism satisfies $\mathbf{1}_{\left(\underline{\theta}_{1t}, \hat{\theta}_{1t}\right)}\left(\theta_{1t}\right) \cdot \left[\sup_{x_t \in X_t} S_t\left(x_t, \theta_1^t\right) - S_t\left(\chi_t\left(\theta_1^t\right), \theta_1^t\right)\right] \to 0$ with λ-probability one as $\hat{\theta}_{1t} \to \underline{\theta}_{1t}$, where $\mathbf{1}_A(z)$ is the indicator function that takes value one when $z \in A$ and zero otherwise. Similarly, when $\lim_{\theta_{1t} \to \bar{\theta}_{1t}} J_{11}^t\left(\theta_{1t}, \theta_1^{t-1}\right) = 0$, one can show that $\mathbf{1}_{\left(\hat{\theta}_{1t}, \bar{\theta}_{1t}\right)}\left(\theta_{1t}\right) \cdot \left[\sup_{x_t \in X_t} S_t\left(x_t, \theta_i^t\right) - S_t\left(\chi_t\left(\theta_i^t\right), \theta_i^t\right)\right] \to 0$ with λ-probability one as $\hat{\theta}_{1t} \to \bar{\theta}_{1t}$. Thus, under reasonable assumptions about the stochastic process, the optimal mechanism comes close to maximizing the total surplus for types that are close to the highest or the lowest type.

It is interesting to contrast this finding with the conclusions of Battaglini (2005), who studies a single-agent model satisfying DSEP in which the agent's type space in each period has only two elements and the evolution of the agent's type is governed by a Markov process. In his model, from the first moment the agent's type turns out to be high, distortions disappears thereafter (this result is referred to as Generalized No Distortion at the Top, or GNDT). Furthermore, the distortions occurring when the agent's type remains low are monotonically decreasing in time and vanish in the limit as $T \to \infty$ (this result is referred to as Vanishing Distortions at the Bottom, or VDB). As our analysis suggests, while the result of GNDT is quite robust in single-agent Markov models satisfying DSEP, the result of VDB need not. In particular, distortions need not be monotonic neither in type nor in time and need not vanish in the long run.

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; we believe these conditions can help in applications.

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

Proposition 7 *Suppose that, for each $i = 1, ..., N$, the assumptions of Proposition 2 hold, and that $\underline{\theta}_{i1} > -\infty$. In addition, suppose that $u_i(\theta, x)$ is increasing in each θ_{it} and that assumption FOSD holds. Then any mechanism $\langle \chi, \psi \rangle$ satisfying the lowest types' participation constraints (10) and the dynamic payoff formula (6) for period one for all i satisfies all the participation constraints (5).*

The result follows directly from (6) by noting that, under the conditions in the proposition, the value function for each agent i is nondecreasing in θ_{i1}.

Next, consider incentive constraints. In what follows we provide conditions ensuring not only that a mechanism is ex-ante Bayesian incentive-compatible, but that it is also incentive compatible at all feasible histories on the equilibrium path. More precisely, the value function of each agent i at *any* of his *truthful* private history $h_i \in H_i(\chi)$ 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, starting from any mechanism $\langle \chi, \psi \rangle$ and any corresponding system of beliefs $\langle \mu_i[\chi,\psi]|(\cdot)\rangle_{i=1}^N$, one can construct a new payment scheme $\hat{\psi}$ and a system of beliefs $\langle \mu_i[\chi,\hat{\psi}]|(\cdot)\rangle_{i=1}^N$ such that the resulting expected utility that each agent obtains in $\left\langle \chi, \hat{\psi} \right\rangle$ in equilibrium (i.e., under truthtelling by all agents) satisfies the dynamic payoff formula at all truthful feasible histories: i.e., after any truthful feasible private history $h_{i,t-1} = \left(\theta_i^{t-1}, \theta_i^{t-1}, x_i^{t-1}\right) \in H_{i,t-1}(\chi)$, any $i = 1, ..., N$, any $t = 1, ..., T$,

$$\Phi_{it}(\theta_{it}, h_{i,t-1}) \equiv \mathbb{E}^{\mu_i[\chi,\hat{\psi}]|(\theta_{it},h_{i,t-1})}\left[u_i(\tilde{\theta},\tilde{x}) - \tilde{p}_i\right] \text{ is Lipschitz continuous in } \theta_{it}\text{, and for a.e. } \theta_{it},$$
$$\frac{\partial \Phi_{it}(\theta_{it}, h_{i,t-1})}{\partial \theta_{it}} = \mathbb{E}^{\mu_i^T[\chi]|(\theta_{it},h_{i,t-1})}\left[\sum_{\tau=t}^{T} J_{it}^{\tau}\left(\tilde{\theta}_i^{\tau}, \tilde{x}_i^{\tau-1}\right)\frac{\partial u_i\left(\tilde{\theta},\tilde{x}\right)}{\partial \theta_{i\tau}}\right]. \tag{13}$$

(Recall that $\mu_i[\chi,\hat{\psi}]|h_{it}$ and $\mu_i^T[\chi]|h_{it}$ denote, respectively, the probability distribution over $\Theta^T \times \Theta_i^T \times X \times \mathbb{R}$ and the corresponding marginal distribution over $\Theta^T \times \Theta_i^T \times X$, when all agents other than i play truthful strategies, agent i's private history is h_{it}, and starting from period t agent i reports truthfully at all future periods, as defined in Subsection 4.3). To construct these payments, for all i, all t, 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

$$D_i^{[\chi]}(\theta_i^t, (\theta_i^{t-1}, m_{it}), x_i^{t-1}) \equiv \mathbb{E}^{\mu_i^T[\chi]|(\theta_i^t,(\theta_i^{t-1},m_{it}),x_i^{t-1})}\left[\sum_{\tau=t}^{T} J_{it}^{\tau}(\tilde{\theta}_i^{\tau}, \tilde{x}_i^{\tau})\frac{\partial u_i(\tilde{\theta},\tilde{x})}{\partial \theta_{i\tau}}\right]. \tag{14}$$

This function measures how agent i's expected payoff in period t changes with θ_{it} when the agent reported truthfully at all preceding periods, he sends a (possibly untruthful) message m_{it} in period t and then reports truthfully at all subsequent periods. We then have the following result.

Lemma 3 *Suppose that, for each $i = 1, ..., N$, the assumptions of Proposition 2 hold. Let $\langle \chi, \psi \rangle$ be any deterministic direct mechanism and $\langle \mu_i[\chi,\psi]|(\cdot)\rangle_{i=1}^N$ a corresponding system of beliefs obtained from $\langle \chi, \psi \rangle$ using the rcpd of Subsection 4.3. Fix a period t. Consider the payment scheme $\hat{\psi}$*

obtained from $\langle\chi,\psi\rangle$ *and* $\langle\mu_i[\chi,\psi]|(\cdot)\rangle_{i=1}^N$ *by setting for all* i*, all* $\theta\in\Theta$*,*

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

where $\hat{\theta}_{it}\leq\theta_{it}$ *is an arbitrary finite value in the closure of* Θ_{it} *and where* $\langle\mu_i^T[\chi]|(\cdot)\rangle_{i=1}^N$ *are the marginal distributions over* $\Theta^T\times\Theta_i^T\times X$ *of the distributions* $\langle\mu_i[\chi,\psi]|(\cdot)\rangle_{i=1}^N$*. Now let* $\langle\mu_i[\chi,\hat{\psi}]|(\cdot)\rangle_{i=1}^N$ *denote any system of beliefs whose marginal distributions over* $\Theta^T\times\Theta_i^T\times X$ *are the same as for* $\langle\mu_i[\chi,\psi]|(\cdot)\rangle_{i=1}^N$*. Then for any* i*, any truthful feasible private history* $h_{i,t-1}=(\theta_i^{t-1},\theta_i^{t-1},x_i^{t-1})\in H_{i,t-1}(\chi)$*, in period* t*, under the beliefs* $\mu_i[\chi,\hat{\psi}]|(\cdot)$*, the mechanism* $\langle\chi,\hat{\psi}\rangle$ *satisfies condition (13).*

Note that the construction achieves the satisfaction of Condition (13) 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, and any system of beliefs $\langle\mu_i^T[\chi]|(\cdot)\rangle_{i=1}^N$, iterating the construction of the payments backward from period T to period one yields a mechanism that, in any period, after any truthful feasible private history satisfies Condition (13), all $i=1,...,N$.

Now, using the payments constructed in Lemma 3, we provide a sufficient condition for the allocation rule χ to be implementable, which is obtained by specializing Proposition 3 to quasilinear environments.

Proposition 8 *Suppose that, for each* $i=1,...,N$*, the assumptions of Proposition 2 hold. Let* $\langle\chi,\psi\rangle$ *be any deterministic direct mechanism and* $\langle\mu_i[\langle\chi,\psi\rangle,\cdot]|(\cdot)\rangle_{i=1}^N$ *a corresponding system of beliefs (as described in Section 4.3). Suppose that, for any* $i=1,...,N$ *any* $\tau\geq t+1$*, given* $\mu_i[\langle\chi,\psi\rangle,\cdot]|(\cdot)$*, the mechanism* $\langle\chi,\psi\rangle$ *is IC for agent* i *at any (possibly non-truthful) period-*τ *feasible private history* $h_{i\tau}\in H_{i\tau}(\chi)$*. If for all* i*, all* (θ_i^t,x_i^{t-1}) *such that* $x_i^{t-1}=\chi_i^{t-1}(\theta_i^{t-1},\theta_{-i}^{t-1})$*,* $\theta_{-i}^{t-1}\in\Theta_{-i}^{t-1}$*,*

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

then there exists a payment scheme $\hat{\psi}$ *and a system of beliefs* $\langle\mu_i[\left\langle\chi,\hat{\psi}\right\rangle,\cdot]|(\cdot)\rangle_{i=1}^N$ *such that, for each agent* $i=1,...,N$*, given the beliefs* $\mu_i[\left\langle\chi,\hat{\psi}\right\rangle,\cdot]|(\cdot)$*, the mechanism* $\langle\chi,\hat{\psi}\rangle$ *is IC at (i) any truthful period-t feasible private history* $h_{it}\in H_{i\tau}(\chi)$*, and (ii) at any (possibly non-truthful) period-*τ *feasible private history* $h_{i\tau}\in H_{i\tau}(\chi)$*, any* $\tau\geq t+1$*. For any* i*, any* σ_i*, any* h_{it}*, the marginal distribution of* $\mu_i[\left\langle\chi,\hat{\psi}\right\rangle,\sigma_i]|h_{it}$ *over* $\Theta^T\times\Theta_i^T\times X$ *is the same as that of* $\mu_i[\langle\chi,\psi\rangle,\sigma_i]|h_{it}$*.*

To understand this result intuitively, fix a truthful feasible history $\left(\theta_i^{t-1},\theta_i^{t-1},x_i^{t-1}\right)$ and then 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} when starting from period $t+1$ the agent reports truthfully. One can think of m_{it} as a one-dimensional "allocation" chosen by agent i in period t. Note that $\partial\Psi_t\left(\theta_{it}, m_{it}\right)/\partial\theta_{it} = D_i^{[\chi]}(\theta_i^t, (\theta_i^{t-1}, m_{it}), x_i^{t-1})$; because, under the beliefs $\mu_i[\langle\chi,\psi\rangle, \cdot]|(\cdot)$, the mechanism $\langle\chi,\psi\rangle$ is IC at any (possibly non-truthful) feasible period-τ history, $\tau \geq t+1$, 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" $m_{it}\left(\theta_{it}\right) = \theta_{it}$ is then implementable (using payments constructed from the dynamic payoff formula using the construction in Lemma 3).

The proposition cannot in general be iterated backward, since it assumes IC at all feasible period-τ, $\tau \geq t+1$, histories but then derives IC only at truthful feasible 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 feasible histories implies IC at *all* feasible histories. The proposition can then be rolled backward to derive a mechanism that is IC at all feasible histories. In other words, starting from any system of rcpd $\langle\Gamma_i^{\chi}\rangle_{i=1}^{N}$ obtained from χ using the construction of Subsection 4.3, one may construct a transfers scheme ψ and a system of beliefs such that, the truthful strategies, together with such beliefs constitute a *weak PBE* of the mechanism $\langle\chi,\psi\rangle$.

The result in Proposition 8 may also be useful in certain non-Markov environments, as illustrated in subsection 5.2 below.

The monotonicity of $D_i^{[\chi]}(\theta_i^t, (\theta_i^{t-1}, m_{it}), x_i^{t-1})$ in m_{it} can be interpreted as a *weak monotonicity condition* of the allocation rule χ. This is reminiscent of familiar results from static mechanism design. In particular, when u_i satisfies the SCP and $N = T = 1$, the result in the proposition coincides with the familiar monotonicity condition that $\chi(m_{i1})$ be nondecreasing in m_{i1}. However, while in those environments, this condition is also necessary, this is not necessarily the case in the more general environments considered here.

Finally note that Proposition 8 can also be used to analyze the effects of disclosing information to the agents in the course of the mechanism 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 $\hat{x}_{it} = \left(x_{it}, x_{it}^d\right)$. 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. The result in Proposition 8 can then be extended to this environment by redefining $D_i^{[\chi]}$ so that the expectation in (14) is now made conditional on $\hat{x}_{it} = \left(x_{it}, x_{it}^d\right)$ 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 agent i somehow learns in period t all the other agents' types θ_{-i}^T (note that this includes past, current and future ones). Formally, this can be captured through a disclosure $x_{it}^d = \theta_{-i}^T$. We then say that the mechanism is Other-Ex-Post incentive compatible for agent i at the feasible private history h_{it} if truthtelling remains an optimal strategy for agent i in this mechanism when after history h_{it} he is shown $x_{it}^d = \theta_{-i}^T$, and θ_{-i}^T is consistent with the history of private decisions observed by agent i. Formally, suppose that assumption DNOT holds. With an abuse of notation, then let $\mu_i[\langle\chi,\psi\rangle,\sigma_i]|h_{it},\theta_{-i}^T$ denote the unique probability measure over $\Theta_i^T \times \Theta_i^T \times Y_i^{T+1}$ that corresponds to the process that starts after the history $h_{it} = (\theta_i^t, m_i^s, y_i^u) \in H_{it}(\chi)$, $t \geq s \geq u \geq t-1$, when agent i follows the strategy σ_i, all other agents follow a truthful strategy, and in period t the agent is shown $x_{it}^d = \theta_{-i}^T$, with θ_{-i}^T such that $y_i^u = \chi_i^u(m_i^u, \theta_{-i}^u)$ where $\theta_{-i}^u = (\theta_{-i,1}, ..., \theta_{-i,u}) \in \Theta_{-i}^u$ denotes the first $u \leq T$ components of θ_{-i}^T,. Note that, as in the single-agent case, this measure is now *uniquely* pinned down by the kernels F_i and the strategy σ_i.

Definition 11 *(i) The mechanism $\langle\chi,\psi\rangle$ is* Other-Ex-Post IC (OEP-IC) *for agent i at the feasible private history $h_{it} = (\theta_i^t, m_i^s, y_i^u) \in H_{it}(\chi)$, $t \geq s \geq u \geq t-1$, if, for any $x_{it}^d = \theta_{-i}^T$ such that $y_i^u = \chi_i^u(m_i^u, \theta_{-i}^u)$, any σ_i,*

$$\mathbb{E}^{\mu_i[\chi,\psi]|h_{it},\theta_{-i}^T}[u_i(\tilde{\theta},\tilde{x}) - \tilde{p}_i] \geq \mathbb{E}^{\mu_i[\langle\chi,\psi\rangle,\sigma_i]|h_{it},\theta_{-i}^T}[u_i(\tilde{\theta},\tilde{x}) - \tilde{p}_i].$$

(ii) The mechanism $\langle\chi,\psi\rangle$ is OEP-IC at period t if it is OEP-IC for each agent $i = 1, ..., N$ at any feasible private history $h_{it} \in H_{it}(\chi)$.

(iii) The mechanism is OEP-IC if it is OEP-IC at all $t = 1, ..., T$.

An appealing property of OEP-IC is that it guarantees that an agent who expects all other agents to follow a truthful strategy finds it optimal to report truthfully after any feasible history, irrespective of whether or not he has been truthful in the past and irrespective of the particular beliefs he may have about the other agents' types. In other words, truthful strategies, together with *any* system of beliefs obtained from $\langle\chi,\psi\rangle$ using *any* profile of rcpd (as described in Subsection 4.3) form a *weak PBE* of the mechanism.

It turns out that some allocation rules can be implemented in an OEP-IC mechanism, under some additional assumptions.

Assumption 12 (PDPD) *Payoffs Depend on Private Decisions: for each* $i = 1, ..., N$, $u_i(\theta, x)$ *depends on* x *only through* x_i^T.

Proposition 9 *Suppose that, for each* $i = 1, ..., N$, *the assumption of Proposition 2 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 period* $\tau \geq t + 1$. *If for all* i *and all* $\tau \geq t$,

$$\chi_{i\tau}(\theta^\tau) \text{ is nondecreasing in } (\theta_{it}, \ldots, \theta_{i\tau}) \text{ for all } \theta_i^{t-1}, \theta_{-i}^\tau, \tag{15}$$

then there exists a payment rule $\hat{\psi}$ *such that the mechanism* $\left\langle \chi, \hat{\psi} \right\rangle$ *is (i) OEP-IC for each agent* i *at any truthful feasible private history* $h_{it} \in H_{it}(\chi)$, *and (ii) OEP-IC at any period* $\tau \geq t + 1$.

For example, in a Markov environment, backward iteration of the result in the proposition implies that, under its assumptions, any allocation rule that is "strongly monotone" in the sense that each $\chi_{it}\left(\theta_i^t, \theta_{-i}^t\right)$ 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 a BIC mechanism *under any possible disclosure* policy.[25]

While it should be clear from Proposition 8 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. Section 5.2 provides 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"*

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

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

[25] By the definition of OEP-IC, any rule that is implementable in an OEP-IC mechanism is then also implementable as a weak PBE, under any disclosure policy.

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 natural conditions on the payoffs and the kernels that ensure that the allocation rule solving the Relaxed Program is strongly monotone. Recall from Subsection 4.6 that in a separable environment (i.e. under DSEP) at any period $t > 1$, the distortion in x_{it} is determined by the impulse response $J_{i1}^{t}(\theta_i^t)$ which need not be monotonic in θ_{it}; in particular, when Θ_{it} is bounded, densities are strictly positive everywhere and the functions $J_{i1}^{t}(\cdot)$ are continuous, the distortion is zero at both $\theta_{it} = \underline{\theta}_{it}$ and $\theta_{it} = \bar{\theta}_{it}$ and, under the additional assumptions of FOSD and SCP, x_{it} is distorted downward for all intermediate θ_{it}. Because of this nonmonotonicity in the distortion, we can have $\chi_{it}\left(\theta_{it}, \theta_i^{t-1}, \theta_{-i}^{t}\right) < \chi\left(\underline{\theta}_{it}, \theta_i^{t-1}, \theta_{-i}^{t}\right)$ 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} and a strictly positive density, this assumption can be satisfied only when the impulse response is identically equal to 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 impulse responses 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 in which the optimal 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 (and incentives) separate over time as it is often assumed in applications.

5.1 Handicapped mechanisms

Consider an environment where the set of feasible allocations is $X \subset \mathbb{R}^{(N+1)T}$. Note that this formulation allows for any possible dependence of the set of feasible decisions in each period on the history of past decisions.

The utility of each agent $i = 1, \ldots, N$ (gross of payments) is

$$u_i(\theta, x) = \sum_{t=1}^{T} \theta_{it} x_{it} - c_i(x), \tag{16}$$

where $c_i : \mathbb{R}^{(N+1)T} \to \mathbb{R}$ can be interpreted an intertemporal cost function. The principal's (gross) payoff 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, which permits us to accommodate such dynamic aspects as intertemporal capacity constraints, habit formation, and learning-by-doing. The private information of each agent $i = 1, \ldots, N$ is assumed to evolve according to a linear AR(k) process, as in Example 1, which ensures that the impulse responses are constants J_{i1}^t. 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

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

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 efficiency-maximization 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 program can be a fairly complex 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 surplus can be implemented through a *"Handicapped" efficient mechanism.* In period 1 each agent i sends a message m_{i1} determining his (time-varying) 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 $i = 1, ..., N$, is as in (16), 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, incentives at any period $t \geq 2$ can be provided using for example the "Team payments" (Athey and Segal, 2007) defined, for all θ, by

$$\psi_i(\theta) = \sum_{j \neq i} u_j(\theta, \chi(\theta)), \tag{17}$$

for all $i = 1, ...n$, where $j \neq i$ includes also $j = 0$. We then have the following result (the proof follows directly from the arguments above).[26]

Proposition 10 *Consider the environment with AR(k) types described above. Given any allocation rule χ that maximizes the expected dynamic virtual surplus, there exists a payment scheme ψ such that, irrespective of the beliefs, the mechanism $\langle \chi, \psi \rangle$ is IC for all i at all truthful feasible private histories h_{it}, any $t \geq 2$.*

That, under the "Team payments" of (17), each agent finds it optimal to report truthfully at any of his truthful history, irrespective of his beliefs about the other agents' types, messages and decisions, follows from the same arguments that establish dominance-strategy implementation of efficient rules under VCG payments.

Incentives in the first period must be checked application-by-application.[27] For example, incentive-compatibility in period one can be easily guaranteed if the costs c_i are identically equal to zero for all i and if $v_0(x)$ is time-separable—the environment then becomes a special case of the class considered in subsection 5.2.

5.1.1 Durable-good monopolist

As an example of a non-time-separable environment where the optimal mechanism takes the form of a handicapped mechanism and where incentives can be guaranteed also in period one, consider the problem of a monopolist selling a durable good. Let $N = 1$, $X_t = \{0, 1\}$ for $t = 1, \ldots, T$, where $x_t = 1$ if and only if the buyer owns the good in period t. The fact that the seller is restricted to sales contracts (rather than rental contracts) is captured in the feasible decision set $X = \left\{ x \in \prod_{\tau=1}^{T} X_\tau : x_t \geq x_{t-1} \ \forall t > 1 \right\}$. Thus, the seller chooses only the timing of the sale - the first period t in which $x_t = 1$.

While production costs could be normalized to zero, we can add generality by letting the seller incur a flow cost c_t in each period t in which the buyer owns the good (c_t could be interpreted as

[26] What is important for the result in Proposition 10 is that (i) the payoff of each agent i depends only on θ_i^T and that the derivatives of u_i with respect to each θ_{it} are independent of θ_i^T; (ii) that the principal's payoff is independent of θ; and that (iii) the total information indexes are independent of θ.

[27] In period 1, the model where the principal has payoff $\hat{v}_0(\theta, x)$ is one with interdependent values since $\hat{v}_0(\theta, x)$ depends on the agents' true period-1 types through the hazard rates $\eta_{i1}(\theta_{i1})$. Hence, the implementability of a virtually efficient allocation rule cannot be guaranteed directly by using Team payments, for the latter induce truthtelling only with private values.

service, or as an opportunity, cost). Thus the principal's (gross) payoff is

$$v(x) = -\sum_{t=1}^{T} c_t x_t.$$

The agent's (gross) payoff is

$$u(\theta, x) = \sum_{t=1}^{T} \theta_t x_t,$$

where the current type θ_t is interpreted as the flow utility of owning the durable good in period t.[28] The type θ_t evolves according to an AR(1) process with impulse-response coefficient $\phi \in (0,1)$. The first-period type θ_1 is distributed on $\Theta_1 = (\underline{\theta}_1, \bar{\theta}_1)$ according to an absolutely continuous c.d.f. F_1 with density $f_1(\theta)$ strictly positive on Θ_1 and hazard rate $\eta_1(\theta_1)$ nondecreasing in θ_1.

The dynamic virtual surplus takes the form

$$\mathbb{E}^{\lambda}\left[\sum_{t=1}^{T} \chi_t(\tilde{\theta}^t)(\tilde{\theta}_t - \phi^{t-1}\eta_1^{-1}(\tilde{\theta}_1) - c_t)\right]. \tag{18}$$

Definition 12 *An allocation rule χ is a* handicapped cut-off rule *if there exist a constant z_1 in the closure of Θ_1, and nonincreasing functions $z_t : \Theta_1 \to \mathbb{R} \cup \{-\infty, +\infty\}$, $t = 2, \dots, T$, such that for all $\theta \in \Theta$,*

$$\chi_1(\theta_1) = \begin{cases} 1 & \text{if } \theta_1 > z_1, \\ 0 & \text{if } \theta_1 \leq z_1, \end{cases}$$

and for $t > 1$,

$$\chi_t(\theta^t) = \begin{cases} 1 & \text{if } \theta_t > z_t(\theta_1) \text{ or } \chi_{t-1}(\theta^{t-1}) = 1, \\ 0 & \text{otherwise.} \end{cases}$$

Note that handicapped cut-off rules are strongly monotone. It is also easy to see that the environment satisfies the assumptions of Proposition 9. Furthermore, the environment is Markov. Thus we can iterate backwards the result of Proposition 9 to show that handicapped cut-off rules are implementable in a mechanism that is IC at all histories.

Proposition 11 *Consider the environment described above. There exists a handicapped cutoff rule that solves the relaxed problem and is part of an optimal mechanism.*

The following payment scheme implements the handicapped cutoff rule that maximizes the dynamic virtual surplus (18). In the first period the buyer is offered a menu of contracts, indexed by θ_1. Each contract entails an up-front payment $P(\theta_1)$ together with an additional payment

[28] Because there is no risk of confusion, in this example we simplify notation by dropping the subscripts $i = 1$ from all variables.

$p_t(\theta_1) = \sum_{\tau=t}^{T} \left(c_\tau + \phi^{\tau-1}\eta_1^{-1}(\theta_1)\right)$ to be made at the time the good is purchased. The buyer chooses a contract and then decides when to buy. The up-front payment $P(\theta_1)$ is computed using Lemma 3; it guarantees that each type θ_1 (weakly) prefers the contract $(P_1(\theta_1), (p_t(\theta_1))_{t=1}^T)$ to any contract $(P_1(\hat{\theta}_1), (p_t(\hat{\theta}_1))_{t=1}^T)$, $\hat{\theta}_1 \neq \theta_1$. Given the contract $(P_1(\theta_1), (p_t(\theta_1))_{t=1}^T)$, it is then immediate that the buyer has the incentives to choose a purchasing strategy that maximizes the dynamic virtual surplus (18).

Next, consider the efficient rule—i.e. the policy that maximizes $\mathbb{E}^\lambda \left[\sum_{t=1}^T \chi_t(\tilde{\theta}^t)(\tilde{\theta}_t - c_t)\right]$. This rule can be implemented, for example, using the team payments defined in (**??**), which amounts to setting $P(\theta_1) \equiv 0$ and $p_t(\theta_1) \equiv \sum_{\tau=t}^T c_\tau$. Comparing these prices to the above payment scheme for the profit-maximizing rule we see that the monopolist optimally reduces the provision of the good by selling to higher types than what is socially efficient. As a result, the sale may happen later than in the efficient plan. This permits the monopolist to limit the expected surplus left to the buyer. Note that, because X is a lattice, this result follows directly from the downward distortions result of Proposition 6. In fact, by that Proposition, the solution of the relaxed problem continues to have downward distortions even if we replace the AR(1) process with an arbitrary (possibly non-Markov) process where decisions do not affect types and where an increase in any past type has a first-order-stochastic-dominance effect on the distribution of the current type (i.e., the process satisfies assumptions DNOT and FSOD).

5.2 Time-Separable Environments

We now consider environments in which the agents' types continue to follow an AR(k) process as in Example 1, but where payoffs and decisions separate over time. The set of possible decisions in each period t is $X_t \subset \mathbb{R}^{N+1}$ with $X = \prod_{t=1}^T X_t$. 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, 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 $i = 1, \ldots, N$. Suppose further that for all $i = 0, ..., N$ 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 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, λ-almost all θ^t,*

$$\chi_t(\theta^t) \in \arg\max_{x_t \in X_t} \left\{ u_{0t}(\theta_{0t}, x_{0t}) + \sum_{i=1}^{N} \left(u_{it}(\theta_{it}, x_{it}) - \frac{J_{i1}^t}{\eta_{i1}(\theta_{i1})} \frac{\partial u_{it}(\theta_{it}, x_{it})}{\partial \theta_{it}} \right) \right\}. \tag{19}$$

Furthermore, χ can be implemented in an OEP-IC mechanism using payments constructed as follows. For any agent $i = 1, ..., n$ and all θ,

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

where for all $t \geq 2$,

$$\psi_{it}(\theta_1, \theta_t) \equiv u_{it}(\theta_{it}, \chi_{it}(\theta_1, \theta_t)) - \int_{\underline{\theta}_{it}}^{\theta_{it}} \frac{\partial u_{it}(r, \chi_{it}(\theta_1, (r, \theta_{-i,t})))}{\partial \theta_{it}} dr, \tag{20}$$

and[29]

$$\begin{aligned} \psi_{i1}(\theta_{i1}, \theta_{-i}^T) \quad \equiv \quad & \mathbb{E}^{\lambda_i|\theta_{i1}} \left[u_i((\tilde{\theta}_i^T, \theta_{-i}^T), \chi(\tilde{\theta}_i^T, \theta_{-i}^T)) - \sum_{t=2}^{T} \psi_{it}(\theta_1, (\tilde{\theta}_{it}, \theta_{-i,t})) \right] \\ & - \int_{\underline{\theta}_{i1}}^{\theta_{i1}} \mathbb{E}^{\lambda_i|r} \left[\sum_{\tau=1}^{T} J_{i1}^{\tau} \frac{\partial u_i((\tilde{\theta}_i^T, \theta_{-i}^T), \chi(\tilde{\theta}_i^T, \theta_{-i}^T))}{\partial \theta_{i\tau}} \right] dr. \end{aligned} \tag{21}$$

The result in Proposition 12 follows essentially from Proposition 9 by observing that, in this environment, *incentives separate over time*. That is, the profit-maximizing allocation rule can be implemented using a mechanism where in periods $t \geq 2$, given any history, each agent simply chooses his current message so as to maximize his flow payoff $u_{it} - \psi_{it}$. This can be seen as follows. First note, by inspection of (19), that the allocation rule χ that maximizes the expected dynamic virtual surplus has the property that, in each period t, $\chi_t(\theta^t)$ depends only on the current reports θ_t and on the agents' period-1 reports θ_1. This in turn follows from the fact that (i) preferences are separable over time, (ii) decisions do not affect types (DNOT), and (iii) the impulse responses $J_{i\tau}^t$ for the AR(k) processes do not depend on the realized types, and (iv) the set of feasible decisions in each period t is history-independent, that is, $X = \prod_{t=1}^{T} X_t$. Thus, for any $t \geq 2$, agent i's message m_{it} has no direct effect on the allocations in periods $\tau > t$ and, because the environment is time-separable (i.e., because (i) and (ii) hold), the allocation in period t has no direct effect on the future allocations.

[29] Recall 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.

Now to see that the allocation rule χ that maximizes the expected dynamic virtual surplus is implementable in an OEP-IC mechanism, fix a period $t \geq 2$ and a vector of (not necessarily truthful) first-period reports θ_1. It is useful to think of period t as a static problem indexed by the first-period reports θ_1. Assumptions (1), (2) and (4) in the proposition (which imply SCP, FOSD, and PDPD) guarantee that agent i's allocation $\chi_{it}(\theta_1, \theta_t)$ in this static problem is monotone in θ_{it} for all $(\theta_1, \theta_{-i,t})$. As is well known, a monotone allocation can be implemented using a transfer $\psi_{it}(\theta_1, \theta_t)$ that takes the form specified in (20). Furthermore, because the flow payoffs u_{it} depend only on own types (i.e. this is a private-value environment), then in a static setting, such allocation rule can be implement in dominant strategies. In our dynamic setting, this means that the period-t allocation can be implemented even if the agents were able to observe the other agents' true types as well as their messages. In other words, the agents' beliefs are irrelevant. The transfers $\psi_{it}(\theta_1, \theta_t)$ thus guarantee that, in each period $t \geq 2$, each agent i indeed faces a static problem and finds it optimal to report truthfully, regardless of the history, and even if he were able to observe all other agents' types and reports.[30]

As for period one, in general providing incentives at $t = 1$ is more involved. 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 Proposition 9. Following the same steps as in the proof of Proposition 9, 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 (21) 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 not essential for the argument). In period one, each agent i chooses from a menu of "handicaps" $(J_{i1}^t \eta_{i1}^{-1}(\theta_{i1}))_{t=1}^T$, indexed by θ_{i1}. Then in each period $t \geq 2$, a "handicapped" VCG mechanism is played with transfers as in (20). 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

[30] In fact, due to time-separability, in periods $t \geq 2$ the mechanism constructed above is not only OEP IC but truly ex-post IC (in the sense that, each agent i would find it optimal to report truthfully even if, in addition to learning all other agents' types, he were able to learn also all his own future types).

$\frac{\partial u_{it}(\theta_{it}, x_{it})}{\partial \theta_{it}}$. However intermediate reports (i.e., reports in periods $2, \ldots, t-1$) remain irrelevant both for the period-t allocation and for the period-t payments.

6 Conclusions

We showed how the first-order approach to the characterization of necessary conditions for incentive-compatibility can be adapted to a dynamic setting with general payoffs, many agents, a continuum of types, and decision-controlled (and possibly non-Markov) processes.

We then derived sufficient conditions for incentive compatibility that permit one to verify the optimality of truthful reporting at a truthful history once the optimality of truthful reporting at future histories has been verified. This backward-induction approach to the characterization of incentive-compatibility is based on the idea that, once the optimal continuation strategy is known, the analysis of incentive-compatibility at any point in time can be reconducted to a static problem with unidimensional types and multi-dimensional decisions.

We then specialized the analysis to multi-agent quasi-linear settings with independent types (across agents). We first qualified in what sense the celebrated revenue equivalence result from static mechanism design extends to dynamic environments. We then showed how optimal mechanisms can be obtained by first solving a relaxed program that consists in searching for an allocation rule that maximizes the expected dynamic virtual surplus (the latter is obtained by considering only the lowest period-1 types participation constraints and applying the dynamic payoff formula to express the agents' intertemporal rents in terms of the allocation rule). The analysis then proceeds by verifying that the allocation rule that solves the relaxed program induces a certain dynamic single-crossing condition on the agents' payoffs which is the analog of the familiar monotonicity condition from static mechanism design. Finally, the characterization is completed by iterating backwards the dynamic payoff formula to construct the supporting transfers. We illustrated how this approach can be put to work in a variety of applications including the design of revenue-maximizing auctions for buyers whose valuations change over time.

Throughout, we maintained two key assumptions. The first one is that of a finite horizon; the second is time-consistency of the agents' preferences. As mentioned in the introduction, extending the analysis to infinite-horizon settings requires a different approach. This alternative approach (which we explore in Pavan, Segal, and Toikka (2009)) complements the one in the present paper in that, when applied to a finite-horizon model, it permits one to validate the dynamic payoff formula under a different (and not nested) set of conditions. Relaxing the time-consistency assumption is challenging and represents an interesting line for future research.

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.

ATKESON, A., AND J. LUCAS, ROBERT E. (1992): "On Efficient Distribution with Private Information," *The Review of Economic Studies*, 59(3), 427–453.

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

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

BERGEMANN, D., AND J. VÄLIMÄKI (2007): "Dynamic Marginal Contribution Mechanism," Discussion Paper, Cowless Foundation.

BIAIS, B., T. MARIOTTI, G. PLANTIN, AND J.-C. ROCHET (2007): "Dynamic Security Design: Convergence to Continuous Time and Asset Pricing Implications," *Review of Economic Studies*, 74(2), 345–390.

BOGNAR, K., T. BORGERS, AND M. MEYER-TER VEHN (2008): "Dynamic Mechanism Design with Costly Participation," Discussion Paper, University of Michigan and UCLA.

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.

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.

Fernandes, A., and C. Phelan (2000): "A Recursive Formulation for Repeated Agency with History Dependence," *Journal of Economic Theory*, 91(2), 223–247.

Garcia, D. (2005): "Monotonicity in Direct Revelation Mechanisms," *Economics Letters*, 88(1), 21–26.

Gershkov, A., and B. Moldovanu (2008a): "Dynamic Revenue Maximization with Heterogenous Objects: A Mechanism Design Approach," University of Bonn.

——— (2008b): "Efficient Sequential Assignment with Incomplete Information," University of Bonn.

Golosov, M., N. Kocherlakota, and A. Tsyvinski (2003): "Optimal Indirect and Capital Taxation," *Review of Economic Studies*, 70(3), 569–587.

Golosov, M., and A. Tsyvinski (2006): "Designing Optimal Disability Insurance: A Case for Asset Testing," *Journal of Political Economy*, 114(2), 257–279.

Green, E. J. (1987): "Lending and the Smoothing of Uninsurable Income," in *Contractual Agreements for Intertemporal Trade*, ed. by P. E., and W. N. University of Minnesota Press.

Kapicka, M. (2006): "Efficient Allocations in Dynamic Private Information Economies with Persistent Shocks: A First Order Approach," University of California, Santa Barbara.

Kocherlakota, N. R. (2005): "Zero Expected Wealth Taxes: A Mirrlees Approach to Dynamic Optimal Taxation," *Econometrica*, 73(5), 1587–1621.

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.

MIRRLEES, 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.

PAVAN, A., I. SEGAL, AND J. TOIKKA (2009): "Infinite-Horizon Mechanism Design: The Independent-Shock Approach," Mimeo, Northwestern and 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.

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.

TCHISTYI, A. (2006): "Security Design with Correlated Hidden Cash Flows: The Optimality of Performance Pricing," Mimeo, NYU Stern.

THOMAS, J., AND T. WORRALL (1990): "Income fluctuation and asymmetric information: An example of a repeated principal-agent problem," *Journal of Economic Theory*, 51(2), 367–390.

TOPKIS, D. M. (1998): *Supermodularity and Complementarity.* Princeton University Press.

WILLIAMS, N. (2008): "Persistent Private Information," Mimeo, Princeton University.

ZHANG, H. (2008): "Solving an Infinite-Horizon Principal-Agent Problem by Direct Construction," Mimeo, University of Southern California.

ZHANG, Y. (2006): "Dynamic Contracting, Persistent Shocks and Optimal Taxation," Mimeo, University of Minnesota.

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* $\tau < t$

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

Proof of Lemma A.1. Note that by Assumption 2 the expectation $\mathbb{E}[\theta_t|\theta^{t-1}, y^{t-1}]$ exists. Taking its derivative we have

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

where the first equality follows by the definition of a derivative, the second equality follows by Lemma 4 below, and 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,

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

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 $\Psi_t\left(\theta^t, m^t, y^{t-1}\right) \equiv V^\Omega(\theta^t, m^t, y^{t-1})$ 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 $\Psi_{T+1}\left(\theta^{T+1}, m^{T+1}, y\right) = V_{T+1}\left(\theta^{T+1}, m, y\right) = U\left(\theta, y\right)$.

Note that, by definition,

$$\Psi_t\left(\theta^t, m^t, y^{t-1}\right) = \int V_{t+1}\left(\theta^{t+1}, m^t, y^t\right) dF_{t+1}\left(\theta_{t+1}|\theta^t, y^t\right) d\Omega_t\left(y_t|m^t, y^{t-1}\right), \tag{22}$$
$$V_{t+1}\left(\theta^{t+1}, m^t, y^t\right) = \sup_{m_{t+1}} \Psi_{t+1}\left(\theta^{t+1}, \left(m^t, m_{t+1}\right), y^t\right).$$

The proof proceeds in a series of Lemmas.

Lemma 4 *For any Lipschitz function $G : \Theta_t \to \mathbb{R}$,*

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

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 $|G(\theta_t)| \leq \left|G\left(\hat{\theta}_t\right)\right| + M\left|\hat{\theta}_t\right| + M|\theta_t|$, and all terms have finite expectations with respect to the probability distributions $F_t\left(\cdot|\theta^{t-1}, y^{t-1}\right)$ and $F_t(\cdot|\eta^{t-1}, y^{t-1})$, the last term by Assumption 2. Thus, we can use integration by parts to write

$$\int G(\theta_t) dF_t\left(\theta_t|\theta^{t-1}, y^{t-1}\right) - \int G(\theta_t) dF_t\left(\theta_t|\eta^{t-1}, y^{t-1}\right)$$
$$= \int G(\theta_t) d[F_t\left(\theta_t|\theta^{t-1}, y^{t-1}\right) - F_t\left(\theta_t|\eta^{t-1}, y^{t-1}\right)]$$
$$= -\int G'(\theta_t)[F_t\left(\theta_t|\theta^{t-1}, y^{t-1}\right) - F_t\left(\theta_t|\eta^{t-1}, y^{t-1}\right)]d\theta_t$$
$$+ \left[G(\theta_t)[F_t\left(\theta_t|\theta^{t-1}, y^{t-1}\right) - F_t\left(\theta_t|\eta^{t-1}, y^{t-1}\right)]\right]\Big|_{\theta_t=\underline{\theta}_t}^{\theta_t=\bar{\theta}_t}.$$

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

$$\left|G(\theta_t)[F_t\left(\theta_t|\theta^{t-1}, y^{t-1}\right) - F_t\left(\theta_t|\eta^{t-1}, y^{t-1}\right)]\right|$$
$$\leq \left(|G(\hat{\theta}_t)| + M|\hat{\theta}_t|\right)\left|F_t(\theta_t|\theta^{t-1}, y^{t-1}) - F_t(\theta_t|\eta^{t-1}, y^{t-1})\right|$$
$$+ M|\theta_t|\left(F_t\left(\theta_t|\theta^{t-1}, y^{t-1}\right) + F_t\left(\theta_t|\eta^{t-1}, y^{t-1}\right)\right)$$
$$\leq \left(|G(\hat{\theta}_t)| + M|\hat{\theta}_t|\right)\left|F_t\left(\theta_t|\theta^{t-1}, y^{t-1}\right) - F_t\left(\theta_t|\eta^{t-1}, y^{t-1}\right)\right|$$
$$+ M\left(\int_{z\leq\theta_t} |z|\, dF_t\left(z|\theta^{t-1}, y^{t-1}\right) + \int_{z\leq\theta_t} |z|\, dF_t\left(z|\eta^{t-1}, y^{t-1}\right)\right)$$
$$\to 0,$$

where the first term converges to zero since $F_t\left(\theta_t|\theta^{t-1}, y^{t-1}\right)$ is continuous in θ^{t-1} by Assumption 5, and the second term converges to zero by 2. The case where $\overline{\theta}_t = +\infty$ and $\theta_t \to +\infty$ is treated symmetrically. ■

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

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

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

Lemma 5 *In an ex-ante IC mechanism Ω, for any integers $1 \leq t \leq \tau$ and for $\mu[\Omega]$-almost all histories $\left(\theta^\tau, \theta^{\tau-1}, y^{\tau-1}\right)$,*

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

Proof. By ex-ante IC we have that, for $\mu[\Omega]$-almost all histories $\left(\theta^\tau, \theta^{\tau-1}, y^{\tau-1}\right)$, $V_\tau\left(\theta^\tau, \theta^{\tau-1}, y^{\tau-1}\right) = \Psi_\tau\left(\theta^\tau, \theta^\tau, y^{\tau-1}\right)$. By definition of V_τ and Ψ_τ, we have for all $\left(\theta^\tau, \theta^{\tau-1}, y^{\tau-1}\right)$ and all θ_t',

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

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

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

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

$$\left|\Psi_t\left(\eta^t, m^t, y^{t-1}\right) - \Psi_t\left(\theta^t, m^t, y^{t-1}\right)\right| \leq M\left\|\eta^t - \theta^t\right\|,$$
$$\left|V_t\left(\eta^t, m^{t-1}, y^{t-1}\right) - V_t\left(\theta^t, m^{t-1}, y^{t-1}\right)\right| \leq M\left\|\eta^t - \theta^t\right\|.$$

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

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

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

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

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

where the first inequality uses (22), the third inequality uses Lemma 4, and the last inequality uses Assumption 5. This shows that Ψ_{t-1} is equi-Lipschitz continuous in θ^{t-1}. ■

Lemma 7 *For any integers τ, t such that $1 \leq t < \tau \leq T$, and any $\left(\theta^{\tau-1}, m^{\tau-1}, y^{\tau-2}\right)$,*

$$\begin{aligned}\frac{\partial^- \Psi_{\tau-1}\left(\theta^{\tau-1}, m^{\tau-1}, y^{\tau-2}\right)}{\partial \theta_t} &\leq \int \frac{\partial^- V_\tau\left(\theta^\tau, m^{\tau-1}, y^{\tau-1}\right)}{\partial \theta_t} dF_\tau\left(\theta_\tau|\theta^{\tau-1}, y^{\tau-1}\right) d\Omega_{\tau-1}\left(y_{\tau-1}|m^{\tau-1}, y^{\tau-2}\right) \qquad (23) \\ &\quad - \int \frac{\partial V_\tau\left(\theta^\tau, m^{\tau-1}, y^{\tau-1}\right)}{\partial \theta_\tau} \frac{\partial F_\tau\left(\theta_\tau|\theta^{\tau-1}, y^{\tau-1}\right)}{\partial \theta_t} d\theta_\tau d\Omega_{\tau-1}\left(y_{\tau-1}|m^{\tau-1}, y^{\tau-2}\right), \\ \frac{\partial_+ \Psi_{\tau-1}\left(\theta^{\tau-1}, m^{\tau-1}, y^{\tau-2}\right)}{\partial \theta_t} &\geq \int \frac{\partial_+ V_\tau\left(\theta^\tau, m^{\tau-1}, y^{\tau-1}\right)}{\partial \theta_t} dF_\tau\left(\theta_\tau|\theta^{\tau-1}, y^{\tau-1}\right) d\Omega_{\tau-1}\left(y_{\tau-1}|m^{\tau-1}, y^{\tau-2}\right) \qquad (24) \\ &\quad - \int \frac{\partial V_\tau\left(\theta^\tau, m^{\tau-1}, y^{\tau-1}\right)}{\partial \theta_\tau} \frac{\partial F_\tau\left(\theta_\tau|\theta^{\tau-1}, y^{\tau-1}\right)}{\partial \theta_t} d\theta_\tau d\Omega_{\tau-1}\left(y_{\tau-1}|m^{\tau-1}, y^{\tau-2}\right).\end{aligned}$$

Proof. Using (22), write for any $\theta_t' \neq \theta_t$

$$\frac{\Psi_{\tau-1}\left(\left(\theta_t', \theta_{-t}^{\tau-1}\right), m^{\tau-1}, y^{\tau-2}\right) - \Psi_{\tau-1}\left(\theta^{\tau-1}, m^{\tau-1}, y^{\tau-2}\right)}{\theta_t' - \theta_t} \tag{25}$$

$$= \int \frac{V_\tau\left(\left(\theta_t', \theta_{-t}^{\tau}\right), m^{\tau-1}, y^{\tau-1}\right) - V_\tau\left(\theta^\tau, m^{\tau-1}, y^{\tau-1}\right)}{\theta_t' - \theta_t} dF_\tau\left(\theta_\tau|\theta^{\tau-1}, y^{\tau-1}\right) d\Omega_{\tau-1}\left(y_{\tau-1}|m^{\tau-1}, y^{\tau-2}\right)$$

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

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

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

We examine separately the behavior of each of the three integrals in (25) as $\theta_t' \to \theta_t$:

- Third integral: Note that for any $y^{\tau-1}$,

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

  $$\to 0 \text{ as } \theta_t' \to \theta_t,$$

  since the integrand is bounded by Lemma 6, and the total variation of the measure

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

  converges to zero by Assumption 6. Thus, the third integral is bounded in absolute value by a term that converges to zero as $\theta_t' \to \theta_t$. Now note that, in the Markov case,

  $$\begin{aligned} & V_\tau\left(\left(\theta_t', \theta_{-t}^{\tau}\right), m^{\tau-1}, y^{\tau-1}\right) - V_\tau\left(\theta^\tau, m^{\tau-1}, y^{\tau-1}\right) \\ = \; & \prod_{l=1}^{t-1} A_l\left(\theta_l, y^l\right)\left[B_t\left(\theta_t', y^t\right) - B_t\left(\theta_t, y^t\right)\right] \\ & + \left(\frac{A_t\left(\theta_t', y^t\right)}{A_t\left(\theta_t, y^t\right)} - 1\right)\left[V_\tau\left(\theta^\tau, m^{\tau-1}, y^{\tau-1}\right) - \sum_{s=1}^{t}\left(\prod_{l=1}^{s-1} A_l\left(\theta_l, y^l\right)\right) B_s\left(\theta_s, y^s\right)\right]. \end{aligned}$$

  Using Lemmas 4 and 6 the third integral then becomes

  $$-\left(\frac{A_t\left(\theta_t', y^t\right)}{A_t\left(\theta_t, y^t\right)} - 1\right)\int \frac{\partial V_\tau\left(\theta^\tau, m^{\tau-1}, y^{\tau-1}\right)}{\partial \theta_\tau} \frac{F_\tau\left(\theta_\tau|\left(\theta_t', \theta_{-t}^{\tau-1}\right), y^{\tau-1}\right) - F_\tau\left(\theta_\tau|\theta^{\tau-1}, y^{\tau-1}\right)}{\theta_t' - \theta_t} d\theta_\tau. \tag{26}$$

  By Lemma 6 and Assumption 5 the integral in (26) is bounded. Thus, the expression in (26) goes to zero as $\theta_t' \to \theta_t$ provided that $A_t\left(\theta_t, y^t\right)$ is continuous in θ_t, which in turn follows from Assumption 4 as long as $B_s\left(\theta_s, y^s\right)$ are not identically zero for $s > t$. When instead $B_s\left(\theta_s, y^s\right)$ are identically zero for all $s > t$, then

$$V_\tau\left(\left(\theta_t', \theta_{-t}^\tau\right), m^{\tau-1}, y^{\tau-1}\right) - V_\tau\left(\theta^\tau, m^{\tau-1}, y^{\tau-1}\right) = \prod_{l=1}^{t-1} A_l\left(\theta_l, y^l\right)\left[B_t\left(\theta_t', y^t\right) - B_t\left(\theta_t, y^t\right)\right]$$

in which case the third integral is identically equal to zero.

- Second integral: Using Lemma 6 and Lemma 4 it can be expressed as

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

  Using in addition Assumption 5, the Dominated Convergence Theorem establishes that as $\theta_t' \to \theta_t$, the second integral in (25) converges to the second integral in (24) and (23).

- First integral: Taking its limsup as $\theta_t' \uparrow \theta_t$ and using Fatou's Lemma,[31] we see that the limsup is bounded above by the first integral in (23). Thus, we obtain (23). Similarly, taking the liminf of the first integral in (25) as $\theta_t' \downarrow \theta_t$ and using Fatou's Lemma, we see that the liminf of this term is bounded below by the first integral in (24), so we obtain (24).

■

Now combining the inequalities in Lemma 7 for $m^\tau = \theta^\tau$ and the inequalities in Lemma 5 we obtain for $\mu[\Omega]$-almost all histories $\left(\theta^{\tau-1}, \theta^{\tau-2}, y^{\tau-2}\right)$,

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

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

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

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

$$\frac{\partial^- V_t\left(\theta^t, \theta^{t-1}, y^{t-1}\right)}{\partial \theta_t} \leq$$

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

$$\leq \frac{\partial_+ V_t\left(\theta^t, \theta^{t-1}, y^{t-1}\right)}{\partial \theta_t}.$$

To complete the proof of the proposition, recall that by definition, $V_t\left(\theta^t, \theta^{t-1}, y^{t-1}\right) = V^\Omega\left(\theta^t, \theta^{t-1}, y^{t-1}\right)$. So by Lemma 6 $V^\Omega\left(\theta^t, \theta^{t-1}, y^{t-1}\right)$ is Lipschitz continuous in θ_t for all $\left(\theta^{t-1}, \theta^{t-1}, y^{t-1}\right)$. Thus, given any $\left(\theta^{t-1}, \theta^{t-1}, y^{t-1}\right)$, the partial derivative $\partial V^\Omega\left(\theta^t, \theta^{t-1}, y^{t-1}\right)/\partial \theta_t$ exists for almost every θ_t. Whenever it does exist, it must be equal to both ends of the above double inequality, which establishes (1).

C Other Proofs Omitted in the Main Text

Proof of Proposition 2. 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 $t \in \{1, \ldots, T-1\}$. We will show that it holds also for t. Using iterated expectations and the induction hypothesis, (1) can be written as

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

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

Proof of Proposition 3.

By (iii), it suffices to consider one-stage deviations in period t. In other words, it suffices to verify that, at any truthful history $\left(\theta^{t-1}, \theta^{t-1}, y^{t-1}\right)$, and for any current type θ_t, the agent's period-t expected payoff from sending the message m_t in period t and then following the truthful strategy at any future history, 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 at $m_t = \theta_t$. 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 8 *Consider a function $\Psi : (\underline{\theta}, \overline{\theta})^2 \to \mathbb{R}$. 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 $\theta < m$. ■

Now, to apply the Lemma, we interpret $\Psi(\theta_t, m_t; \theta^{t-1}, y^{t-1})$ as the agent's expected payoff from truthtelling in the mechanism $\hat{\Omega}$ constructed from Ω by ignoring the agent's report in period t and substituting it with m_t. Assumption (iii) then implies that the mechanism $\hat{\Omega}$ is IC at *any* period-t history, and by implication, $\Psi(\theta_t, m_t; \theta^{t-1}, y^{t-1})$ coincides with the agent's value function in the new mechanism $\hat{\Omega}$. Applying to $\hat{\Omega}$ the result in Proposition 2, we then have that, for any m_t, $\Psi(\cdot, m_t; \theta^{t-1}, y^{t-1})$ is Lipschitz continuous in θ_t and $\partial\Psi(\theta_t, m_t; \theta^{t-1}, y^{t-1})/\partial\theta_t = D^\Omega\left(\theta^t, \left(\theta^{t-1}, m_t\right), y^{t-1}\right)$ 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, at any truthful history $(\theta^{t-1}, \theta^{t-1}, y^{t-1})$, and given any θ_t, the function $\Psi(\theta_t, \cdot; \theta^{t-1}, y^{t-1})$ is maximized at $m_t = \theta_t$ which in turn implies that the mechanism Ω is IC at *any* truthful period-t history. ■

Proof of Proposition 4. Given $\langle\chi, \psi\rangle$ and $\langle\chi, \hat{\psi}\rangle$, let $\Omega_i[\chi, \psi]$ and $\Omega_i[\chi, \hat{\psi}]$ denote any two randomized direct mechanisms that agent i faces respectively under $\langle\chi, \psi\rangle$ and $\langle\chi, \hat{\psi}\rangle$, as defined in subsection 4.3. Then let $V^{\Omega_i[\chi,\psi]} : H_i \to \mathbb{R}$ and $V^{\Omega_i[\chi,\hat{\psi}]} : H_i \to \mathbb{R}$ denote the corresponding value functions.

We first establish the following result.

Lemma 9 *Suppose the assumptions in Proposition 4 hold. Then, for any i, any t, $\lambda[\chi]$–almost all truthful private histories $h_i^{t-1} = (\theta_i^{t-1}, \theta_i^{t-1}, \chi_i^{t-1}(\theta_i^{t-1}, \theta_{-i}^{t-1}))$, there exists a scalar $K_{it}(h_i^{t-1})$ such that*

$$V^{\Omega_i[\chi,\psi]}\left(\theta_{it}, h_i^{t-1}\right) - V^{\Omega_i[\chi,\hat{\psi}]}\left(\theta_{it}, h_i^{t-1}\right) = K_{it}(h_i^{t-1}) \text{ for all } \theta_{it}. \tag{27}$$

Proof of the Lemma. Take any $i = 1, ..., N$. From Lemma 1, the fact that $\langle \chi, \psi \rangle$ and $\langle \chi, \hat{\psi} \rangle$ are ex-ante BIC implies that they are IC at $\mu_i^T[\chi]$-almost all truthful private histories $h_i^{t-1} \equiv (\theta_i^{t-1}, \theta_i^{t-1}, \chi_i^{t-1}(\theta_i^{t-1}, \theta_{-i}^{t-1}))$, for any $t \geq 1$ (recall that $\mu_i^T[\chi]$ is the *unique* probability measure describing agent i's beliefs over $\Theta^T \times \Theta_i^T \times X$ in any mechanism implementing the allocation rule χ).

Iterating (1) backward, 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}, the value functions $V^{\Omega_i[\chi,\psi]}\left(\cdot, h_i^{t-1}\right)$ and $V^{\Omega_i[\chi,\hat{\psi}]}\left(\cdot, h_i^{t-1}\right)$ are Lipschitz continuous in θ_{it} and

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

This also implies 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 (27) holds. ■

The result in part (i) 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, the value function coincides with the expected payoff under truthtelling with probability one.

The proof for part (ii) is by induction. Suppose there exists a $K_i \in \mathbb{R}$ such that

$$\mathbb{E}^{\lambda[\chi]}[V^{\Omega_i[\chi,\psi]}(\tilde{\theta}_{i\tau}, \tilde{h}_i^{\tau-1}) \mid \tilde{\theta}_i^{\tau}] - \mathbb{E}^{\lambda[\chi]}[V^{\Omega_i[\chi,\hat{\psi}]}(\tilde{\theta}_{i\tau}, \tilde{h}_i^{\tau-1}) \mid \tilde{\theta}_i^{\tau}] = K_i \tag{28}$$

when $\tau = t \geq 1$. We then show that (28) holds also $\tau = t + 1$.

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

$$V^{\Omega_i[\chi,\psi]}(\theta_{it}, h_i^{t-1}) = \mathbb{E}^{\mu_i^T[\chi]|\theta_{it}, h_i^{t-1}}[V^{\Omega_i[\chi,\psi]}(\tilde{\theta}_{it+1}, \tilde{h}_i^t)].$$

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

$$\begin{aligned}
&\mathbb{E}^{\lambda[\chi]}[V^{\Omega_i[\chi,\psi]}(\tilde{\theta}_{it},\tilde{h}_i^{t-1}) \mid \tilde{\theta}_i^t] - \mathbb{E}^{\lambda[\chi]}[V^{\Omega_i[\chi,\hat{\psi}]}(\tilde{\theta}_{it},\tilde{h}_i^{t-1}) \mid \tilde{\theta}_i^t] \\
&= \mathbb{E}^{\lambda[\chi]}[V^{\Omega_i[\chi,\psi]}(\tilde{\theta}_{i,t+1},\tilde{h}_i^{t}) \mid \tilde{\theta}_i^t] - \mathbb{E}^{\lambda[\chi]}[V^{\Omega_i[\chi,\hat{\psi}]}(\tilde{\theta}_{i,t+1},\tilde{h}_i^{t}) \mid \tilde{\theta}_i^t] \\
&= \mathbb{E}^{\lambda[\chi]}[V^{\Omega_i[\chi,\psi]}(\tilde{\theta}_{i,t+1},\tilde{h}_i^{t}) - V^{\Omega_i[\chi,\hat{\psi}]}(\tilde{\theta}_{i,t+1},\tilde{h}_i^{t}) \mid \tilde{\theta}_i^t] \\
&= \mathbb{E}^{\lambda[\chi]}[K_{i,t+1}(\tilde{h}_i^t) \mid \tilde{\theta}_i^t],
\end{aligned} \tag{29}$$

where the last equality follows from Lemma 9.

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 then a deterministic function of $\tilde{\theta}_i^t$ and $\tilde{\theta}_{-i}^t$ and because types are independent we then have that

$$\begin{aligned}
\mathbb{E}^{\lambda}[K_{i,t+1}(\tilde{h}_i^t) \mid \tilde{\theta}_i^t] &= \mathbb{E}^{\lambda}[K_{i,t+1}(\tilde{h}_i^t) \mid \tilde{\theta}_i^{t+1}] \\
&= \mathbb{E}^{\lambda}[V^{\Omega_i[\chi,\psi]}(\tilde{\theta}_{i,t+1},\tilde{h}_i^{t}) \mid \tilde{\theta}_i^{t+1}] - \mathbb{E}^{\lambda}[V^{\Omega_i[\chi,\hat{\psi}]}(\tilde{\theta}_{i,t+1},\tilde{h}_i^{t}) \mid \tilde{\theta}_i^{t+1}],
\end{aligned} \tag{30}$$

where the last equality follows again from Lemma 9. Combining (29) with (30) 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 $\tilde{\theta}_1^t$. The result in (30) is thus always true when the allocation rule is deterministic. We conclude that, when $N = 1$, the result in part (ii) holds even if assumption DNOT is dispensed with. ■

Proof of Proposition 5. Parts (i) and (ii) follow directly from Lemma 2. 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 mechanisms 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). ■

Proof of Proposition 6. Define $g : \mathcal{X} \times \{-1, 0\} \to \mathbb{R}$ as

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

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 each $J_{i1}^t\left(\theta_i^t, x_i^{t-1}\right)$ does not depend on x_i^{t-1}, which is reflected in the formula. The assumption of FOSD ensures that each $J_{i1}^t\left(\tilde{\theta}_i^t\right) \geq 0$. 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). ■

Proof of Proposition 7. Under the assumptions in the proposition, $J_{i1}^t\left(\theta_i^t, x_i^{t-1}\right) \geq 0$ and $\partial u_i(\theta, x)/\partial\theta_{it} \geq 0$ all i, t, all (θ, x); hence, by (6), $V^{\Omega_i[\chi,\psi]}(\theta_{i1})$ is nondecreasing. ■

Proof of Lemma 3. Fix $i = 1, ..., N$. By construction, at any truthful feasible private histories $h_{i,t-1} = (\theta_i^{t-1}, \theta_i^{t-1}, x_i^{t-1}) \in H_{i,t-1}(\chi)$.

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

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 truthtelling (by all agents). The second equality follows directly from the definition of $\delta_i\left(\theta_i^t, x_i^{t-1}\right)$. Note that the function $D^{[\chi]}\left(\left(\theta_i^{t-1}, \cdot\right), \left(\theta_i^{t-1}, \cdot\right), x_i^{t-1}\right)$ is measurable and bounded and therefore integrable. Thus at any truthful feasible private history $h_{i,t-1} = (\theta_i^{t-1}, \theta_i^{t-1}, x_i^{t-1}) \in H_{i,t-1}(\chi)$, in period t, under the beliefs $\mu_i[\chi, \hat{\psi}]|(\cdot)$, the mechanism $\langle\chi, \hat{\psi}\rangle$ satisfies Condition (13). ■

Proof of Proposition 8. Let $\hat{\psi}$ be the payment rule that is obtained from $\langle\chi, \psi\rangle$ and $\langle\mu[\chi, \psi]|(\cdot)\rangle_{i=1}^N$ using the construction indicated in the proof of Lemma 3. Take any system of beliefs $\langle\mu[\langle\chi, \hat{\psi}\rangle, \cdot]|(\cdot)\rangle_{i=1}^N$ such that, for any i, any σ_i, any h_{it}, the marginal distribution of $\mu_i[\left\langle\chi, \hat{\psi}\right\rangle, \sigma_i]|h_{it}$ over $\Theta^T \times \Theta_i^T \times X$ is the same as that of $\mu_i[\langle\chi, \psi\rangle, \sigma_i]|h_{it}$. By construction, under such beliefs, the following are true, for any i. (a) The mechanism $\langle\chi, \hat{\psi}\rangle$ is IC at any period-τ, $\tau \geq t+1$ feasible private history; hence the mechanism $\left\langle\chi, \hat{\psi}\right\rangle$ satisfies Condition (iii) of Proposition 3. (b) After any truthful feasible private history $h_{i,t-1} = \left(\theta_i^{t-1}, \theta_i^{t-1}, x_i^{t-1}\right)$, Condition (13) in period t is satisfied; This establishes Condition (i) of Proposition 3 for period t. (c) $D_i^{[\chi]}(\theta_i^t, (\theta_i^{t-1}, m_{it}), x_i^{t-1})$ is nondecreasing in m_{it}; This implies that Condition (ii) of Proposition 3 is also verified. The result then follows from Proposition 3. ■

Proof of Proposition 9. 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, λ is the product of each agent i's stochastic process over Θ_i^T, 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^T given θ_i^t.

The payment rule $\hat{\psi}$ is obtained by adapting the construction of Lemma 3 to the situation where agent i is shown θ_{-i}^T in period t and faces a stochastic process λ_i over his own types (which is essentially a single-agent situation):[32]

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

where $\hat{\theta}_{it} \leq \theta_{it}$ is any arbitrary finite value in the closure of Θ_{it}, and where

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

Note that, under assumption DNOT, $J_{it}^{\tau}\left(\theta_i^{\tau}, x_i^{\tau-1}\right)$ does not depend on x_i^{τ}. By FOSD, $J_{it}^{\tau}\left(\theta_i^{\tau}\right) \geq 0$. By SCP, PDPD, and (15), $\partial u_i\left(\theta, \chi\left(\left(m_{it}, \theta_{i,-t}^T\right), \theta_{-i}^T\right)\right)/\partial \theta_{i\tau}$ is nondecreasing in m_{it} for all θ_{-i}^T. 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}^T. The result then follows from Proposition 8 applied to this setting. ■

Proof of Proposition 11. As argued in the main text before the statement of the proposition, that handicapped cut-off rules are implementable in a mechanism that is IC at all histories, follows from Proposition 9 . Hence it suffices to show that the expected dynamic virtual surplus in (18) is maximized by such a rule. It is immediate that the optimal policy depends on θ^t only through θ_1 and θ_t. What remains to show is that, for any t any θ^t such that $\sum_{s=1}^{t-1} \chi_s(\theta^s) = 0$, the optimal period-$t$ allocation $\chi_t(\theta^t)$ is non-decreasing in (θ_1, θ_t) (when $\Theta_1 \times \Theta_t$ is endowed with the product order).

Fix a history θ^t such that the good has not been sold in periods $1, \ldots, t-1$. Let $\varepsilon^{-t} \equiv (\varepsilon_\tau)_{\tau=t+1}^T$ denote the realizations of the shocks in periods $\tau = t+1, \ldots, T$ in the AR(1) process determining the types. Consider an arbitrary *continuation* policy to be followed in periods $\tau = t+1, ..., T$, conditional on not selling the good in period t. Since we are fixing θ^t, such a policy can always be described as a function of the shocks ε^{-t}. That is, let $\hat{\chi}_\tau((\varepsilon_s)_{s=t+1}^{\tau})$ be the allocation in period $\tau > t$. The principal's continuation payoff from *not* selling at t and following $\hat{\chi}$ in the future, given

[32] The specific moment in period t at which each agent i is shown θ_{-i}^T is irrelevant for the result.

ε^{-t}, is denoted by

$$b\left(\theta^t, \varepsilon^{-t}; \hat{\chi}\right) \equiv \sum_{\tau=t+1}^{T} \hat{\chi}_\tau((\varepsilon_s)_{s=t+1}^{\tau}) \left(\phi^{\tau-t}\theta_t + \sum_{s=t+1}^{\tau} \phi^{\tau-s}\varepsilon_s - \phi^{\tau-1}\eta_1^{-1}(\theta_1) - c_\tau\right).$$

Now let θ^t and $\left(\theta'\right)^t$ be any two type histories such that $\theta_1 \geq \theta_1'$ and $\theta_t \geq \theta_t'$. Fix ε^{-t}. For any continuation strategy $\hat{\chi}$,

$$\begin{aligned}
b(\theta^t, \varepsilon^{-t}; \hat{\chi}) - b((\theta')^t, \varepsilon^{-t}; \hat{\chi}) &= \sum_{\tau=t+1}^{T} \hat{\chi}_\tau((\varepsilon_s)_{s=t+1}^{\tau}) \left(\phi^{\tau-t}\theta_t + \sum_{s=t+1}^{\tau} \phi^{\tau-s}\varepsilon_s - \phi^{\tau-1}\eta_1^{-1}(\theta_1) - c_\tau\right) \\
&- \sum_{\tau=t+1}^{T} \hat{\chi}_\tau((\varepsilon_s)_{s=t+1}^{\tau}) \left(\phi^{\tau-t}\theta_t' + \sum_{s=t+1}^{\tau} \phi^{\tau-s}\varepsilon_s - \phi^{\tau-1}\eta_1^{-1}(\theta_1') - c_\tau\right) \\
&= \sum_{\tau=t+1}^{T} \hat{\chi}_\tau((\varepsilon_s)_{s=t+1}^{\tau}) \left[\phi^{\tau-t}(\theta_t - \theta_t') - \phi^{\tau-1}\left(\eta_1^{-1}(\theta_1) - \eta_1^{-1}(\theta_1')\right)\right] \\
&\leq \sum_{\tau=t+1}^{T} \left[\phi^{\tau-t}(\theta_t - \theta_t') - \phi^{\tau-1}\left(\eta_1^{-1}(\theta_1) - \eta_1^{-1}(\theta_1')\right)\right] \\
&\leq \sum_{\tau=t}^{T} \left(\phi^{\tau-t}\theta_t - \phi^{\tau-1}\eta_1^{-1}(\theta_1)\right) - \sum_{\tau=t}^{T} \left(\phi^{\tau-t}\theta_t' - \phi^{\tau-1}\eta_1^{-1}(\theta_1')\right),
\end{aligned}$$

where the last expression is the difference in the expected payoffs from selling at t given θ^t and selling at t given $\left(\theta'\right)^t$. Taking expectations over ε^{-t} on both sides and re-arranging gives

$$\sum_{\tau=t}^{T} \left(\phi^{\tau-t}\theta_t' - \phi^{\tau-1}\eta_1^{-1}(\theta_1')\right) - \mathbb{E}\left[b((\theta')^t, \tilde{\varepsilon}^{-t}; \hat{\chi})\right] \leq \sum_{\tau=t}^{T} \left(\phi^{\tau-t}\theta_t - \phi^{\tau-1}\eta_1^{-1}(\theta_1)\right) - \mathbb{E}\left[b(\theta^t, \tilde{\varepsilon}^{-t}; \hat{\chi})\right].$$

Hence for *any* continuation strategy $\hat{\chi}$, the payoff differential between selling in period t and not selling is higher for θ^t than for $\left(\theta'\right)^t$. Taking the infimum of both sides over all possible continuation policies, we then have that

$$\sum_{\tau=t}^{T} \left(\phi^{\tau-t}\theta_t' - \phi^{\tau-1}\eta_1^{-1}(\theta_1')\right) - \sup_{\hat{\chi}} \mathbb{E}\left[b((\theta')^t, \tilde{\varepsilon}^{-t}; \hat{\chi})\right] \leq \sum_{\tau=t}^{T} \left(\phi^{\tau-t}\theta_t - \phi^{\tau-1}\eta_1^{-1}(\theta_1)\right) - \sup_{\hat{\chi}} \mathbb{E}\left[b(\theta^t, \tilde{\varepsilon}^{-t}; \hat{\chi})\right].$$

This implies that the optimal allocation rule takes the form of a handicapped cutoff rule. ■

Proof of Proposition 12. We first show that, under conditions (1)–(4), any allocation rule χ that is part of a profit-maximizing mechanism must satisfy condition (19) in the proposition for all t, λ-almost all θ^t.

To this purpose, first note that, by Proposition 7, assumption (2) and (4) guarantee that the participation constraints for all types other than the lowest ones can be ignored.

Next note that, because payoffs and decisions are time-separable, then an allocation rule maximizes the expected dynamic virtual surplus (9) if and only if, for all t, λ-almost all θ^t, $\chi_t(\cdot)$ satisfies condition (19). To prove the result it then suffices to show that any allocation rule that satisfies condition (19) is implementable in an OEP-IC mechanism that gives zero expected surplus to the lowest types. The result in Proposition 5 then implies that *any* allocation rule that is part of a profit-maximizing mechanism must necessarily satisfy condition (19) for all t, λ-almost all θ^t.

As a preliminary step, note that, by inspection, the allocation rule that solves the relaxed program has the property that, in each period t, the period-t allocation $\chi_t(\cdot)$ depends only on the period-t 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 is shown $x_{it}^d = \theta_{-i}^T$. Because the other agents are assumed to be truthful, then necessarily θ_{-i}^T is consistent with x_i^{t-1}, in the sense of Definition 11; that is, given m_i^{t-1}, $x_i^{t-1} = \chi_i^{t-1}(m_i^{t-1}, \theta_{-1}^{t-1})$.

Now consider the allocation rule $\chi_i^T(\cdot; \theta_{-i}^T)$ 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 $t \geq 2$. The property that 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, together with the fact that payoffs are time-separable and independent of other agents' types and that assumptions DNOT, and PDPD hold in this environment, then implies that the agent's incentives separate over time. In particular, at any feasible history $(\theta_i^t, m_i^{t-1}, x_i^{t-1})$, given $x_{it} = \theta_{-i}^T$, the choice of the optimal message m_{it} depends on $(\theta_i^t, m_i^{t-1}, x_i^{t-1}, \theta_{-i}^T)$ 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

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

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

Next, consider period 1. We proved above that there exists a payment scheme ψ such that the mechanism $\langle \chi, \psi \rangle$ that is OEP-IC at any period $\tau \geq t+1$. Because assumptions DNOT, FOSD, SCP and PDPD hold in this environment, and because χ is strongly monotone, then Proposition 9 implies that there exists a payment rule $\hat{\psi}$ such that $\left\langle \chi, \hat{\psi} \right\rangle$ is OEP-IC at any history. The construction of the payments $\hat{\psi}$ then follows from the proof of that proposition and leads to the remaining terms $\psi_{i1}(\theta_{i1}, \theta_{-i}^{T})$. ■