Dynamic Mechanism Design:

Incentive Compatibility, Profit Maximization and Information Disclosure*

Alessandro Pavan — Northwestern University

Ilya Segal — Stanford University

Juuso Toikka — MIT

August 26, 2011

Abstract

We examine the design of incentive-compatible mechanisms in dynamic environments in which information arrives gradually over time and decisions are made over multiple periods. The model allows for multiple agents, for the horizon to be either finite or infinite, for the type processes to be non-Markov and controlled by past decisions, and for the payoffs to be non time-separable. This flexibility permits us to unify the existing literature by identifying general principles, while also favoring novel applications. Our first result is the derivation of an envelope formula for the derivative of an agent's equilibrium payoff with respect to his current type. We identify primitive assumptions that validate such formula as a necessary condition for incentive compatibility in all possible mechanisms. This formula combines the familiar marginal effects of types on payoffs with novel marginal effects of current types on future ones captured by the impulse response functions. We show how to construct transfers that guarantee that this formula is satisfied at all histories and qualify in what sense such transfers are pinned down by the allocation rule ("revenue equivalence"). Next, we show how this formula yields an expression for dynamic virtual surplus which is instrumental to the design of optimal mechanisms and to the study of the dynamics of distortions under such mechanisms. Lastly, we provide sufficient conditions for PBE-implementability of an allocation rule and for its robustness to an agent's observation of the other agents' types. We conclude by showing how the results can be put to work in applications such as the design of sequential auctions for durable-goods and for experience goods ("bandit auctions").

JEL Classification Numbers: D82, C73, L1.

Keywords: asymmetric information, stochastic processes, incentives, mechanism design, envelope theorems

*For useful comments and suggestions, we thank participants at various conferences, seminars, and institutions where this paper was presented. A special thank is to Thomas Chaney, Daniel Garrett, Li Hao, Narayana Kocherlakota, Andy Skrzypacz, and Xianwen Shi. 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, "Long-Term Contracting in a Changing World" by Pavan, and "Infinite-Horizon Mechanism Design" by Pavan, Segal, and Toikka.

1 Introduction

We consider the design of incentive-compatible mechanisms in a dynamic environment in which agents receive private information over time and decisions are made over multiple periods. The model allows for a finite or infinite horizon, for serial correlation of the agents' private information, as well as for the dependence of this information on past allocations. For example, it covers as special cases such problems as the allocation of private and public goods to agents whose valuations follow a stochastic process, the procedures for selling new experience goods to consumers who refine their valuations upon consumption, the design of multi-period procurement auctions for bidders whose costs evolve stochastically over time and may exhibit learning-by-doing effects, and the design of optimal dynamic taxes and of managerial compensation schemes.

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 his past allocations (from which he can make inferences about other agents' types and allocations). Despite the resulting complications, we deliver a general necessary condition for incentive compatibility and some sufficient conditions, and then show how to use these conditions to characterize optimal (profit-maximizing) mechanisms in applications.

We consider a multi-agent environment in which the stochastic processes governing the evolution of the agents' types are independent of one another, except through their dependence on the allocations observed by the agents. This assumption is a proper extension of the familiar "Independent Types" assumption for static mechanism design to the dynamic setting, and prevents the full extraction of the agents' surplus à la Cremer and McLean (1988).

The cornerstone of our analysis is the derivation (and validation) of an envelope formula for the derivative of an agent's equilibrium expected payoff with respect to his private information in any Bayesian incentive-compatible mechanism.[1] Similarly to Mirrlees's first-order approach for static environments (Mirrlees (1971)), this formula provides an envelope condition summarizing local incentive-compatibility constraints. Intuitively, the envelope formula represents the impact of an (infinitesimal) change in the agent's current type on his equilibrium expected payoff. It accounts both for the familiar direct effect of the current type on the agent's utility, as well as for the impact that all future types have on the utility, weighted by the effect that the current type has on the type distributions in each of the subsequent periods, which is both direct and indirect through its impact on the distribution of types in intermediate periods. All these stochastic effects are summarized in an *impulse response function* that describes the effect of the current type on all future ones by

[1] Even though this paper's focus is on mechanism design, the derived envelope formula may be useful for stochastic programming problems in different contexts.

representing future types as a combination of the current type, of the decisions taken over time, and of independent shocks. 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.

The difficulty in establishing the envelope formula comes from the fact that the usual conditions for an envelope theorem (such as those in Milgrom and Segal (2002)) cannot be taken for granted in our dynamic setting. To see the substantive problem, consider a deviation in which the agent misreports his current type and then reports truthfully from the next period onwards. Equivalently, we can view the future types as being observed and reported by different agents. But then, according to Cremer and McLean (1988), the correlation between current and future types would allow the principal to fully extract the agent's surplus, in which case the agent's payoff clearly would not satisfy the envelope formula. The technical manifestation of the problem is that in mechanisms of the Cremer-McLean kind, the agent's expected payoff from deviations such as the one just described is not well-behaved in the current type, even if his underlying utility function and the kernels are smooth in types.

We circumvent the problem by focusing on a carefully chosen subset of strategies that still includes truthtelling (which must be an optimal strategy in an incentive-compatible mechanism) and for which the expected payoff can be guaranteed to be well-behaved in the current type, under some appropriate primitive conditions. This subset is obtained by representing the type processes by means of serially independent shocks, which turns out to be always possible, and then restricting the agents to report the subsequent shocks truthfully (which amounts to a particular restriction on the set of strategies in the primitive representation).[2] We then derive our dynamic envelope formula by imposing, in addition to the usual assumptions of differentiability and equi-Lipschitz continuity of the utility functions in types, appropriate bounds on the impulse response functions (which bound the dependence of future types on the current type).

Next we use the dynamic envelope formula in a quasilinear environment to express the principal's expected payoff in any incentive-compatible and individually-rational mechanism as the expected "*virtual surplus*," appropriately defined for the dynamic setting. This derivation uses only the dynamic envelope formula and the participation constraints of the agents' lowest types in the initial period. The derivation yields the dynamic "*Relaxed Program*," which consists of maximizing the expected virtual surplus while ignoring all the other incentive and participation constraints. (In general, the Relaxed Program is a stochastic dynamic programming problem.) In particular, the Relaxed Program yields a simple intuition for the dynamics of distortions in profit-maximizing mechanisms: these distortions are introduced to reduce the agents' expected surplus, as computed at the time of contracting (i.e., when the participation constraints must be satisfied).[3] However, due to the serial

[2] This approach was first applied to a dynamic mechanism design problem by Eso and Szentes (2007).

[3] While we do not impose participation constraints in periods other than the initial one, in most applications of interest, the payments we propose guarantee participation also in subsequent periods.

correlation of the stochastic type processes, the principal distorts the agents' consumption choices not only in the first period, but also in any subsequent period in which the agent's type is responsive to the first-period type, as measured by the impulse response functions.

Our analysis demonstrates that the dynamics of distortions is determined by the dynamics of the impulse responses (and not by other properties of the type process, such as serial correlation). This insight sheds light on what drives the qualitative differences across the special cases studied in the previous literature. It also reveals that, contrary to what has been obtained in such special cases, in general distortions need not be monotone in either the current or past types and need not decline with time. Lastly, we identify conditions under which the profit-maximizing allocation rule entails downward distortions relative to the efficient rule.

Studying the Relaxed Program is clearly satisfactory only insofar as the allocation rule that solves this program can be shown to be implementable in a mechanism that guarantees participation and truthful reporting by all agents. Thus, we proceed to provide sufficient conditions for implementability of a given choice rule, not just in a Bayes-Nash equilibrium, but also in a Perfect-Bayesian Equilibrium (and in some cases in an equilibrium that is robust to the agents' observation of other agents' reports or types). For this purpose, we first construct payments that guarantee that the equilibrium payoffs satisfy the envelope formula in every period. In fact, we show that the net-present-value (NPV) of such payments is often "essentially" unique, which extends the static "*revenue-*, or more precisely, *payoff-equivalence*" result to the dynamic setting. Namely, with a single-agent, the allocation rule pins down the NPV of the payments in each state up to a state-independent constant. With many agents, the envelope conditions pin down the expectation (over the other agents' types) of the NPV of the payments to each agent as a function of his own types.

Next, we identify sufficient conditions for an allocation rule to be implementable in the PBE of a direct mechanism in which all agents follow a "*strongly truthful*" strategy. These strategies prescribe truthtelling at all histories, even those involving past lies. The focus on strongly truthful strategies is clearly restrictive, since in general an agent who lied in the past may find it optimal to keep lying. Yet, this focus is justified for Markov environments (where the agents' true past type history does not affect current and future incentives), and also for some specific non-Markov environments. For such environments, we derive a dynamic "single-crossing" condition on the allocation rule, which, together with the envelope formula, ensures that deviations from strong truthtelling are unprofitable at any history. The condition states that one-stage upward lies increase the partial derivative of the agent's expected payoff with respect to his true current type, while one-stage downward lies decrease this partial derivative (where the partial derivative can be calculated using the dynamic envelope formula). We first use this single-crossing condition to establish the suboptimality of one-stage deviations from strong truthtelling, and then extend this result to potentially infinitely lasting deviations by establishing an appropriate one-stage deviation principle for this setting.[4]

The single-crossing condition plays a similar role to the monotonicity of the allocation rule in the static setting, except that it is not generally necessary for incentive compatibility, only sufficient. Furthermore, we state some simple conditions that guarantee that the single-crossing condition is satisfied. Namely, we show that when the kernels are not controlled by the allocations and satisfy first-order stochastic dominance and the payoffs are supermodular, the single-crossing condition is guaranteed by the monotonicity of the allocation rule in each of the reports (*strong monotonicity*) or by a weaker *ex-post monotonicity* property that requires that, in each state, the intertemporal weighted average of each agent's consumption choices be nondecreasing in the agent's reports. These monotonicity conditions also ensure that, for appropriately constructed payments, the agents have no incentives to lie even if they were allowed to observe the other agents' types – past, current, and even future ones. Finally, we identify sufficient conditions on the primitives that guarantee that the allocation rule solving the relaxed program is indeed strongly monotone, so that the solution to the relaxed program also solves the full program.

In Section 6 we show how the aforementioned results can be put to work in applications. The first application is an experimentation setting where multiple buyers compete in each period for the provision of a good in fixed supply and where valuations are refined endogenously upon consumption. We show how our results permit one to recast the design of a profit-maximizing mechanism for this environment as a virtual multi-arm bandit problem whose solution is a "bandit auction" that allocates the good in each period to the buyer with the highest virtual Gittins index. The second application is a durable-good problem where the monopolist must decide when to sell an indivisible good (say, a license or a house) to buyers whose valuations evolve exogenously with time. These two applications are non-time separable and require dynamic programming. In contrast, the last application considers a family of problems where payoffs separate over time (as in sequential procurement auctions and regulation) and where the flow payoffs are possibly governed by a non-Markov process.

The rest of the paper is organized as follows. We wrap up this section by briefly discussing the related literature. Section 2 describes the environment. Section 3 characterizes necessary conditions for the implementability of an allocation rule as a BNE of a mechanism and establishes our dynamic payoff formula. Section 4 contains the relaxed program, it shows how to use the dynamic payoff formula to represent the principal's payoff as dynamic virtual surplus, and examines the dynamics of distortions under profit-maximization. Section 5 shows how to construct payments that guarantee that the equilibrium payoffs satisfy the necessary envelope conditions after each history, it derives conditions under which such payments are essentially unique (revenue equivalence), and finally identifies sufficient conditions for PBE (and periodic ex-post) implementability. Section 6 presents a few

[4] The validity of the one-stage deviation principle in our environment does not follow from previous results because the payoffs need not be continuous at infinity (e.g., the flow payoffs need not be bounded).

applications, while Section 7 contains concluding remarks. All proofs omitted in the main text are in the Appendix at the end of the manuscript.

1.1 Related Literature

Some of the results in the paper have precedents in the existing literature, but for very specific settings. A key contribution of the present paper is a general approach that unifies the existing literature, helps explain what drives the differences in the preceding papers and facilitates novel applications. This general approach also clarifies which of the previous findings were general and which ones depended on the specific settings being considered.

The literature on dynamic mechanism design goes back to the pioneering work of Baron and Besanko (1984), who used the first-order approach in a two-period single-agent setting to derive an optimal mechanism for regulating a natural monopoly. They characterized optimal distortions using an "informativeness measure," which is a two-period version of our impulse response formula (see also Riordan and Sappington (1987)). More recently, Courty and Li (2000) considered a similar model to study optimal advanced ticket sales. They also provided some sufficient conditions for a dynamic (two-period) allocation rule to be implementable. Our paper builds on some of the ideas in these papers, extending them to a setting with an arbitrary (possibly infinite) number of periods, multiple agents, and more general payoffs and type processes.[5] Furthermore, in contrast to these early papers, we provide conditions on the primitive environment that validate the first-order approach.

Besanko (1985) and Battaglini (2005) characterize the optimal mechanism for a single agent whose type follows an infinite-horizon Markov process. While Besanko (1985) considers an AR(1) process with a continuum of states, Battaglini (2005) considers a two-state Markov process.[6] The qualitative results in these two papers are quite different: while in Besanko (1985) the allocation in each period depends only on the agent's initial and current type and is downward distorted at each finite history with probability one, in Battaglini (2005) once the agent's type turns high his consumption becomes efficient at any subsequent period, irrespective of the subsequent types. Our analysis clarifies what drives this difference and shows more generally how the dynamics of distortions can be understood in terms of the dynamics of the impulse responses of future types on the initial ones.

[5] Allowing for an infinite horizon is important: Some applications naturally have an infinite horizon; others are facilitated by abstracting from terminal dates, for example because they permit a recursive representation. This, however, introduces complications stemming for example from the need to spread the payments appropriately over time and from the impossibility of using backward induction to establish incentive compatibility. Likewise, allowing for multiple agents not only permits novel applications, it introduces novel effects stemming for example from the inferences that the agents make over time about the other agents' types. Allowing for general payoffs and type processes confers flexibility to the model and also permits us to identify general principles.

[6] Although not meant to be methodological, these papers, together with Courty and Li (2000) and Eso and Szentes (2007), have been inspirational to many of us working on dynamic mechanism design.

Eso and Szentes (2007) consider a two-period model with many agents but with a single allocation decision, to be made in the second period. They 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, and use this representation to study the effects of the seller's disclosure of information on surplus extraction. Our analysis extends their independent-shock approach to incentive compatibility to infinite-horizon settings, allowing for many decisions and for decision-controlled processes, with the goal of studying the dynamics of decisions in optimal mechanisms. (Interestingly, the independent-shock approach can be avoided in finite-horizon models, in which backward induction coupled with integration by parts can be used to obtain an alternative derivation of the envelope formula. This backward-induction approach, first used by Baron and Besanko (1984), was generalized in an earlier version of this paper.)

A few recent papers propose transfer schemes for implementing efficient (i.e., expected surplus-maximizing) dynamic mechanisms that generalize static VCG and expected-externality mechanisms (see, e.g., Bergemann and Välimäki (2010), Athey and Segal (2007), and the references therein). These papers, however, do not provide a general analysis of incentive compatibility in dynamic settings.[7,8]

A number of independent papers (Board (2008), Gershkov and Moldovanu (2009a,b,c 2010a,b,c), Board and Skrzypacz (2010), Dizdar et al., (2011), Pai and Vohra (2008), Said (2011)) consider efficient or profit-maximizing dynamic mechanisms in settings where each agent receives only one piece of private information but where the agents or objects arrive stochastically over time. The characterization of incentive compatibility in all these papers is static and interesting dynamics emerges from the optimal stopping problem faced by the designer.[9,10] Building on this literature and on the results in the current paper, Garrett (2011) recently shows that novel interesting effects emerge when private information about arrival dates is combined with private information about time-varying valuations.

The paper is also related to the 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

[7] An alternative characterization of dynamic incentive compatibility is developed by Rahman (2010) who extends Rochet's (1987) cyclical monotonicity to a dynamic environment. The applicability of this approach to the design of optimal mechanisms is, however, yet to be explored.

[8] In a recent paper, Kakade et al. (2011) build on our first-order approach to characterize the properties of revenue-maximizing allocation rules in environments satisfying a certain separability-in-the-first-component condition. They then show that these rules can be implemented through the virtual analog of the dynamic pivot payments of Bergemann and Välimäki (2010).

[9] Board (2007) considers a setting where the agent's private information evolves over time and consists of the initial type along with a sequence of independent *transitory* shocks. As in the papers cited above, the characterization of incentive-compatibility in this setting is however static.

[10] For a survey of these papers see Bergemann and Said (2011).

(1992)), the more recent literature considers the case of persistent private information (e.g., Fernandes and Phelan (2000), Golosov et al. (2003), Kocherlakota (2005), Albanesi and Sleet (2006), Golosov and Tstvinski (2006), Battaglini and Coate (2008), Zhang (2009), Kapicka (2010)).[11] Particularly related are Farhi and Werning (2010) and Golosov et al. (2010), who build on our first-order approach to characterize the properties of optimal dynamic tax codes. Related are also the recent literature on financial contracting with persistent private information on income dynamics (e.g., Tchistyi (2006), Biais et al. (2007), Williams (2011)) and the recent literature on dynamic managerial compensation (Garrett and Pavan (2011a,b), Edmans and Gabaix (2011) and Edmans et al., (2011)) that applies our first-order approach to a setting where the manager's productivity changes with time.

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 easier to ensure in a dynamic setting than in a static multidimensional setting. This is because in a dynamic setting an agent is asked to report each dimension of his private information before learning the subsequent dimensions, and so has fewer deviations available than in the corresponding static setting in which he observes all the dimensions at once. Because of this, the set of implementable allocation rules is larger in a dynamic setting than in the corresponding static multidimensional setting. 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. On the other hand, our necessary conditions for incentive compatibility are valid also for multidimensional screening problems.

The literature on dynamic mechanism design is also related to the literature on dynamic contracting with adverse selection and limited commitment (see, e.g., Laffont and Tirole (1988, 1990), and more recently Skreta (2006) and the references therein). While this literature typically assumes constant types and generates interesting dynamics through the lack of commitment, the literature on dynamic mechanism design assumes full commitment (on the principal's side) and generates interesting dynamics either through variations in types (as in this paper), or through population dynamics (as in the literature discussed above).[12]

Finally, we also touch here upon the issue of transparency in mechanisms. Calzolari and Pavan (2006a,b) 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.

[11] Some of this work limits the 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.

[12] A couple of recent papers that combine lack of commitment with change in types are Battaglini (2007) and Strulovici (2011).

2 The Environment

Decisions. Time is discrete and indexed by $t = 0, 1, \dots, \infty$. There are $N \geq 1$ agents. In each period t, each agent i observes a signal $\theta_{it} \in \Theta_{it} = (\underline{\theta}_{it}, \bar{\theta}_{it}) \subset \mathbb{R}$ for some $-\infty \leq \underline{\theta}_{it} \leq \bar{\theta}_{it} \leq +\infty$ and then sends a message to a mechanism which leads to a decision $x_{it} \in X_{it}$ and a payment p_{it} for each agent i.[13] Each X_{it} is assumed to be a measurable space (with the sigma-algebra left implicit). The set of feasible decision sequences is $X \subset \prod_{t=0}^{\infty} \prod_{i=1}^{N} X_{it}$. This formulation allows for the possibility that the set of feasible allocations in a given period depends on the allocations in the previous periods, or that the set of feasible decisions with each agent depends on the decisions taken with the other agents. We then let $X_t \equiv \prod_{i=1}^{N} X_{it}$, $X_i^t \equiv \prod_{s=0}^{t} X_{is}$, and $X^t \equiv \prod_{s=0}^{t} X_s$.[14]

Each agent i in each period t observes his own allocations x_{it} but not the other agents' allocations $x_{-i,t}$.[15] The observability of x_{it} should be thought of as a constraint: in each period, a mechanism can reveal more information to agent i than x_{it}, but it cannot conceal x_{it}. Our necessary conditions for incentive compatibility will not depend on what additional information is disclosed to the agent by the mechanism. As for sufficient conditions, we will provide conditions under which extra information can be disclosed to the agents without violating incentive compatibility (e.g., we will construct payments that can be disclosed in each period and identify conditions under which the other agents' reports and allocations can also be disclosed).

Stochastic Processes. The evolution of each agent i's information is described by the collection of kernels[16,17]

$$F_i \equiv \left\langle F_{it} : \Theta_i^{t-1} \times X_i^{t-1} \to \Delta(\Theta_{it}) \right\rangle_{t=0}^{\infty}$$

where $\Theta_i^{\infty} = \prod_{t=0}^{\infty} \Theta_{it}$ and $\Theta = \prod_{i=1}^{N} \Theta_i^{\infty}$. Thus, $F_{it}\left(\theta_i^{t-1}, x_i^{t-1}\right)$ denotes the history-dependent distribution (aka "kernel") of the random variable $\tilde{\theta}_{it}$, given the history of past signals $\theta_i^{t-1} \in \Theta_i^{t-1}$ in period $t-1$ and the history of past allocations $x_i^{t-1} \in X_i^{t-1}$.[18] The dependence on past decisions can capture, e.g., learning-by-doing or experimentation. The time-t signals of different agents are drawn independently of each other. We slightly abuse notation by using $F_{it}(\cdot|\theta_i^{t-1}, x_i^{t-1})$ to denote the cumulative distribution function (c.d.f.) corresponding to the measure $F_{it}(\theta_i^{t-1}, x_i^{t-1})$.

[13] Hereafter we will refer to each θ_{it} interchangeably as agent i's period-t signal or period-t type.

[14] For example, the (intertemporal) allocation of a private good in fixed supply $\bar{x}$ can be modelled by letting $X = \{x \in \mathbb{R}_+^{\infty N} : \sum_{it} x_{it} \leq \bar{x}\}$, 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 $X = \prod_{t=0}^{\infty} X_t$ with $X_t = \left\{x_t \in \mathbb{R}_+^N : x_{1t} = x_{2t} = \dots = x_{Nt}\right\}$.

[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 be accomodated by introducing a fictitious agent observing them.

[16] For any measurable set B, $\Delta(B)$ denotes the set of probability measures over B.

[17] All functions are assumed to be measurable throughout the paper.

[18] Throughout, tildes will be used to denote random variables, whereas the same variable without tilde will denote a realization. Also note that the sets Θ_{it} denote the sets over which the supports of the kernels are defined as opposed to the supports the kernels themselves.

Note that we build in the assumption of "*independent types*" in the sense of Athey and Segal (2007): in addition to independence of agents' signals within any period t, we require that the distribution of an agent's private signal be determined by things he observed $\left(\theta_i^{t-1}, x_i^{t-1}\right)$. Without these restrictions, the first-order approach would not be valid, as it would be possible to extract the agents' information rents as in Cremer and McLean (1988). On the other hand, dependence on other agents' past signals through the implemented observable decisions x_i^{t-1} is allowed.

Preferences. Each agent i has von Neumann-Morgenstern preferences over lotteries over $\Theta \times X \times \mathbb{R}$, described by a Bernoulli utility function of the quasilinear form $U_i(\theta, x) + P_i$, where $U_i : \Theta \times X \to \mathbb{R}$ and where P_i can be interpreted as the Net Present Value (NPV) of the payments received by agent i.[19] (In some applications P_i could have a different interpretation – see, e.g., Garrett and Pavan (2011b)). The special case of "finite horizon" arises when each $U_i(\theta, x)$ depends only on $\left(\theta^T, x^T\right)$ for some finite T.[20]

We impose the following technical conditions on the utility functions.[21]

Condition 1 (U-TD) *Utility Totally Differentiable: For each* $i = 1, ..., N$, $t \geq 0$, $x \in X$, *and* $\theta \in \Theta$, $U_i(\theta_i, \theta_{-i}, x)$ *is totally differentiable in* $\theta_i^t \in \Theta_i^t$.

Observe that with a finite horizon T, this condition simply means that $U_i\left(\theta_i^T, \theta_{-i}^T, x^T\right)$ is totally differentiable in θ_i^T.

We fix a discount factor $\delta \in (0,1)$ and define the norm $\|\theta_i\| \equiv \sum_{t=0}^{\infty} \delta^t |\theta_{it}|$, and let $\Theta_{i\delta} = \{\theta_i \in \Theta_i^\infty : \|\theta_i\| < \infty\}$.[22] With this norm, we then assume that the following condition holds.

Condition 2 (U-ELC) *Utility Equi-Lipschitz Continuous: For each* $i = 1, ..., N$, *the function family* $\{U_i(\cdot, \theta_{-i}, x)\}_{\theta_{-i} \in \Theta_{-i}, x \in X}$ *is equi-Lipschitz continuous on* $\Theta_{i\delta}$. *In other words, there exists* $A_i < \infty$ *such that for all* $\theta_i, \theta_i' \in \Theta_{i\delta}$, $\theta_{-i} \in \Theta_{-i}, x \in X$,

$$\left|U_i\left(\theta_i', \theta_{-i}, x\right) - U_i\left(\theta_i, \theta_{-i}, x\right)\right| \leq A_i \left\|\theta_i' - \theta_i\right\|.$$

[19] Most of our characterizations of incentive compatibility also apply to non-quasilinear environments, simply by letting $P_i \equiv 0$. The only exceptions concern results in which satisfying incentive compatibility involves constructing appropriate payments.

[20] With finite horizon, the payments can be made all at once after the mechanism is over. With infinite horizon, however, the payments have to be spread over time. We will ensure that this can be done in Section 5 below.

[21] To facilitate the reading, throughout the entire exposition, we will denote by U- conditions referring to the utility functions and by F- conditions referring to the processes.

[22] Observe that we could always rescale θ_{it} to normalize $\delta = 1$, making $\|\cdot\|$ coincide with the standard l_1 norm on sequences. However, we allow $\delta < 1$ to deal without rescaling with the standard economic applications with time discounting. Note as well that for a finite horizon, the norm $\|\cdot\|$ is equivalent to the Euclidean norm for any $\delta > 0$, and so the choice of δ is irrelevant. For infinite horizon, increasing δ weakens the conditions imposed below on the utility function while strengthening the conditions imposed on the kernels.

A feasible **allocation rule** is a mapping $\chi : \Theta \to X$ such that the allocation $\chi_t(\theta)$ implemented in each period t depends only on the history θ^t (and so will be written as $\chi_t(\theta^t)$). We denote the set of feasible allocation rules by $\mathcal{X}$. Given the kernels F, an allocation rule χ uniquely defines a stochastic process over Θ, which we denote by $\lambda[\chi]$.[23] Also, for any agent i and initial signal θ_{i0}, we let $\lambda_i[\chi]|\theta_{i0}$ denote the process where first θ_{j0} is drawn from F_{j0} for each $j \neq i$ and then $\lambda[\chi]$ is applied to the starting point $(\theta_{i0}, \theta_{-i,0})$. We assume that these processes have an Expected Present Value (primitive conditions on the kernels that ensure this assumption for every χ are given in Lemma 2 below):

Condition 3 (F-BE$_0$) *Process Bounded in Expectation at time* 0*: For each* $i = 1, ..., N$*,* $\theta_{i0} \in \Theta_{i0}$*, and* $\chi \in \mathcal{X}$*,* $\mathbb{E}^{\lambda_i[\chi]|\theta_{i0}}[||\tilde{\theta}_i||] < \infty$*.*

Condition F-BE$_0$ implies, in particular, that for any feasible allocation rule χ, for $\lambda_i[\chi]|\theta_{i0}$–almost all θ_i in Θ_i, $||\theta_i|| < \infty$. That is, with $\lambda_i[\chi]|\theta_{i0}$–probability one, $\tilde{\theta}_i \in \Theta_{i\delta}$.

3 First-Order Necessary Condition for Incentive Compatibility

A key contribution of the paper is in establishing a first-order necessary condition for the implementability of a dynamic **choice rule** $\langle \chi, \Psi \rangle$, which consists of a feasible allocation rule $\chi \in \mathcal{X}$ and a **payment rule** $\Psi : \Theta \to \mathbb{R}^N$.[24] To make the result as strong as possible, we will focus on the weakest solution concept, Bayes-Nash equilibrium. On the other hand, when providing sufficient conditions for implementability in Section 5 below, we will assume a stronger solution concept, Perfect Bayesian Equilibrium.

A general mechanism asks the agents to send messages in each period and then implements allocations and payments based on the agents' reports and discloses some information to the agents in each period. A Bayes-Nash equilibrium for such mechanism is a profile of contingent strategies in the mechanism such that each agent i's strategy maximizes the expected payoff for every starting type θ_{i0}, assuming that the other agents follow their equilibrium strategies. This solution concept satisfies the Revelation Principle: without loss of generality, we can restrict attention to implementing a given choice rule in a "direct" mechanism, in which each agent i's message space in each period t is his type space Θ_{it}, no additional information is disclosed to him (in addition to his observed allocation x_{it}), and in equilibrium the agents report their types truthfully. That is, suppose that agent i deviates to a reporting strategy $\sigma_i = \left\langle \sigma_{it} : \Theta_i^t \times \Theta_i^{t-1} \times X_i^{t-1} \to \Theta_{it} \right\rangle_{t=0}^{\infty}$, where

[23] This probability measure exists and is unique by the Tulcea extension theorem (see, e.g., Pollard (2002), Chapter 4, Theorem 49).

[24] While we do not model randomized choice rules explicitly, they could be incorporated by conditioning on the random types of a fictitious agent. Below, we will offer sufficient conditions for profit-maximization to be achieved with deterministic choice rules.

$\sigma_{it}(\theta_i^t, \hat{\theta}_i^{t-1}, x_i^{t-1}) \in \Theta_{it}$ represents the report in period t when the true type history is θ_i^t, the reported type history is $\hat{\theta}_i^{t-1}$, and the allocation history is x_i^{t-1}.[25] In the direct mechanism implementing the choice rule $\langle \chi, \Psi \rangle$, when the other agents report truthfully, this deviation results in the choice rule consisting of the allocation rule $\chi \circ \sigma_i$ defined iteratively by $(\chi \circ \sigma_i)_t(\theta^t) = \chi_t(\bar{\sigma}_{it}(\theta^t), \theta_{-i}^t)$ where $\bar{\sigma}_{it}(\theta^t) = \sigma_{it}(\theta_i^t, \bar{\sigma}_i^{t-1}(\theta^{t-1}), (\chi \circ \sigma_i)_i^{t-1}(\theta^{t-1}))$ and of the payment rule $(\Psi \circ \sigma_i)(\theta) = \Psi(\bar{\sigma}_i(\theta), \theta_{-i})$. Truthtelling being a Bayes-Nash equilibrium means that any such deviation cannot improve upon truthtelling:

Definition 1 (BNIC) *The choice rule $\langle \chi, \Psi \rangle$ is Bayes-Nash Incentive Compatible (BNIC) if for each agent i, each initial type θ_{i0}, agent i's equilibrium payoff*

$$V_i^{\langle \chi, \Psi \rangle}(\theta_{i0}) \equiv \mathbb{E}^{\lambda_i[\chi]|\theta_{i0}} \left[U_i(\tilde{\theta}, \chi(\tilde{\theta})) + \Psi_i(\tilde{\theta}) \right] \tag{1}$$

cannot be raised by deviating to any reporting strategy σ_i, i.e., for any such strategy we have that $V_i^{\langle \chi \circ \sigma_i, \Psi \circ \sigma_i \rangle}(\theta_{i0}) \leq V_i^{\langle \chi, \Psi \rangle}(\theta_{i0})$.[26]

Our first-order approach will apply an envelope theorem to the agent's problem of choosing the optimal reporting plan for the dynamic problem described above. Since the space of reporting plans is quite rich, we will apply an envelope theorem of the type of those in Milgrom and Segal (2002), which does not impose any restrictions on the choice space. This theorem requires that, for any fixed choice, the objective function be well-behaved (differentiable and equi-Lipschitz continuous) in the parameter. As anticipated in the Introduction, this is unfortunately not guaranteed in a dynamic setting. The reason is that the agent's expected payoff depends on the parameter θ_{i0} not just through the impact that the latter has on the utility function U_i but also through its effect on the stochastic process $\lambda_i[\chi]|\theta_{i0}$ that determines his future types. Because future allocations may depend on future types in arbitrary ways when arbitrary mechanisms and strategies are considered, there is no guarantee that the aforementioned properties hold.

To circumvent the problem, we focus on a carefully chosen subset of strategies, which still includes truthtelling, and then use this subset to identify primitive properties on the stochastic processes and on the utility functions that guarantee that the agent's expected payoff is well-behaved in the parameter θ_{i0}, for all possible strategies in this subset, and for all possible BNIC mechanisms.[27]

We first introduce an auxiliary representation of the stochastic processes where types are generated via independent shocks:

---

[25] Hereafter, we follow the convention that any set with superscripts -1 is simply an empty set.

[26] In particular, this inequality requires that $V_i^{\langle \chi, \Psi \rangle}(\theta_{i0})$ be well-defined and finite, while $V_i^{\langle \chi \circ \rho_i, \Psi \circ \rho_i \rangle}(\theta_{i0})$ be either finite or $-\infty$. The same applies to the definition of on-path IC below.

[27] Note that being able to establish that these properties hold for all possible mechansims is not trivial, for it is akin to requiring that, in an abstrac dynamic stochastic optimization problem, these properties hold for different objective functions. Indeed, holding constant the message spaces (and hence the sets of possible strategies) and changing the mechanism is equivalent to changing the objective function for the underlying optimization problem faced by the agent.

Definition 2 (S-representation) *A triple* $\langle \mathcal{E}, G, z \rangle$ *where* $\mathcal{E} = \left\langle \langle \mathcal{E}_{it} \rangle_{i=1}^N \right\rangle_{t=0}^\infty$ *is a collection of measurable spaces,* $G = \left\langle \langle G_{it} \rangle_{i=1}^N \right\rangle_{t=0}^\infty$ *is a collection of probability distributions* $G_{it} \in \Delta(\mathcal{E}_{it})$ *on these spaces, and*

$$z = \left\langle \left(z_{it} : \Theta_i^{t-1} \times X_i^{t-1} \times \mathcal{E}_{it} \to \Theta_{it} \right)_{i=1}^N \right\rangle_{t=0}^\infty$$

is a collection of functions, is a state representation *of the kernels* F *if for each* $i = 1, ..., N$, $t \geq 0$, $\left(\theta_i^{t-1}, x_i^{t-1} \right) \in \Theta_i^{t-1} \times X_i^{t-1}$, $z_{it} \left(\theta_i^{t-1}, x_i^{t-1}, \tilde{\varepsilon}_{it} \right)$ *is a random variable distributed according to* $F_{it} \left(\theta_i^{t-1}, x_i^{t-1} \right)$ *when* $\tilde{\varepsilon}_{it}$ *is a random variable distributed according to* G_{it}.

This definition means that the process $\lambda [\chi] | \theta_{i0}$ can be generated as follows: Let $\tilde{\varepsilon}$ be a random variable on $\Pi_{t=1}^\infty \Pi_{i=1}^N \mathcal{E}_{it}$ such that each component $\tilde{\varepsilon}_{it}$ is distributed according to G_{it} and the components are distributed independently of each other. Construct iteratively $\theta_{it} = z_{it} \left(\theta_i^{t-1}, x_i^{t-1}, \varepsilon_{it} \right)$ for $t \geq 0$. In this case, we can let each agent i observe ε_{it} in each period t, yet $\left(x_i^{t-1}, \theta_i^t \right)$ will continue be a sufficient statistics for the agent's private information in period t with respect to what is payoff-relevant to the agent (even if the mapping $z_{it} \left(\theta_i^{t-1}, x_i^{t-1}, \cdot \right)$ is not one-to-one). This property will play an important role in establishing the envelope theorem below. Later we will show that *any* collection of kernels F admits a State representation. However, because this representation is not necessarily the most convenient one, our result will be established for an arbitrary given representation, provided that it satisfies the following condition.

Condition 4 (F-BIR$_0$) *Process Bounded Impulse Responses at* $t = 0$. *The kernels* F *admit a state representation with the following property: For each* $i = 1, .., N$, *define the functions*

$$\left\langle Z_{it} : \Theta_{i0} \times X_i^{t-1} \times \mathcal{E}_i^t \to \Theta_{it} \right\rangle_{t=0}^\infty$$

inductively as follows[28]

$$Z_{it}(\theta_{i0}, x_i^{t-1}, \varepsilon_i^t) = \begin{cases} \theta_{i0} & \text{for } t = 0, \\ z_{it}(Z_i^{t-1}(\theta_{i0}, x_i^{t-2}, \varepsilon_i^{t-1}), x_i^{t-1}, \varepsilon_{it}) & \text{for } t \geq 1 \end{cases} \tag{2}$$

with $Z_i^{t-1}(\theta_{i0}, x_i^{t-2}, \varepsilon_i^{t-1}) \equiv \left\langle Z_{i\tau}(\theta_{i0}, x_i^{\tau-1}, \varepsilon_i^\tau) \right\rangle_{\tau=0}^{t-1}$. *For each* $\theta_{i0} \in \Theta_{i0}$, $t \geq 0$, $x_i \in X_i$, $\varepsilon_i^t \in \mathcal{E}_i^t$, *the derivative* $\partial Z_{it}(\theta_{i0}, x_i^{t-1}, \varepsilon_i^t)/\partial \theta_{i0}$ *exists and is bounded in absolute value by* $C_{it}(\varepsilon_i)$, *where* $\mathbb{E}[||C_i(\tilde{\varepsilon}_i)||] < \infty$ *for all* θ_{i0}.

We are now ready to establish that the following condition is necessary for BNIC:

Definition 3 (ICFOC$_0$) *The choice rule* $\langle \chi, \Psi \rangle$ *satisfies ICFOC*$_0$ *(time-*0 *first-order condition for incentive compatibility) if, for each* $i = 1, ..., N$, *agent* i*' equilibrium expected payoff* $V_i^{\langle \chi, \Psi \rangle}(\theta_{i0})$ *is*

[28] Throughout, for all $i = 1, ..., N$, all $t \geq 0$, $\mathcal{E}_i^t \equiv \prod_{s=0}^t \mathcal{E}_{is}$.

Lipschitz continuous in θ_{i0}*, and its derivative is given a.e. by*

$$\frac{dV_i^{\langle\chi,\Psi\rangle}(\theta_{i0})}{d\theta_{i0}} = \mathbb{E}^{\lambda_i[\chi]|\theta_{i0}}\left[\sum_{t=0}^{\infty}\frac{\partial U_i(\tilde{\theta},\chi(\tilde{\theta}))}{\partial\theta_{it}} I_{it}(\tilde{\theta}_i^t,\chi_i^{t-1}(\tilde{\theta}))\right], \tag{3}$$

where[29]

$$I_{it}(\theta_i^t, x_i^{t-1}) \equiv \mathbb{E}\left[\left.\frac{\partial Z_{it}(\theta_{i0}, x_i^{t-1}, \tilde{\varepsilon}_i^t)}{\partial\theta_{i0}}\right| Z_i^t(\theta_{i0}, x_i^{t-1}, \tilde{\varepsilon}_i^t) = \theta_i^t\right]. \tag{4}$$

Theorem 1 (Necessity of ICFOC$_0$) *Under Conditions U-TD, U-ELC, F-BE*$_0$ *and F-BIR*$_0$*, any BNIC choice rule satisfies ICFOC*$_0$*.*

Intuitively, ICFOC$_0$ is derived by focusing on strategies in which agent i is allowed to misreport his initial type θ_{i0} but is then forced to report all the subsequent *shocks* truthfully (which corresponds to specific misreports of future types $\theta_i^{>0}$). This defines the subset of strategies alluded to above for which we will be able to establish that the expected payoff has nice properties, for all the elements of the sets, and over all possible mechanisms. In particular, we will show that, for any strategy in this set, the agent's expected payoff is equi-Lipschitz continuous and differentiable in θ_{i0}, in which case an envelope theorem of the type of those reported in Milgrom and Segal (2002) can be used to calculate the derivative of the agent's value function. The formal proof follows.

Proof of Theorem 1. Consider the fictitious environment in which the process is generated via independent shocks $\tilde{\varepsilon}$ and in each period $t \geq 1$ each agent i observes the shock ε_{it} and computes $\theta_{it} = z_{it}\left(\theta_i^{t-1}, x_i^{t-1}, \varepsilon_{it}\right)$ (recall that the function need not be invertible in ε_{it}). Consider the direct revelation mechanism in the fictitious environment in which each agent i reports θ_{i0} in period 0 and ε_{it} in each period $t \geq 1$ and which implements the decision rule $\hat{\chi}_t\left(\theta_0, \varepsilon^t\right) = \chi_t\left(Z^t\left(\theta_0, \hat{\chi}^{t-1}\left(\theta_0, \varepsilon^{t-1}\right), \varepsilon^t\right)\right)$ in each period t (defined recursively on t) with payment rule $\hat{\Psi}\left(\theta_0, \varepsilon\right) = \Psi\left(Z\left(\theta_0, \hat{\chi}\left(\theta_0, \varepsilon\right), \varepsilon\right)\right)$ (where $Z_t \equiv (Z_{it})_{i=1}^N$, $Z^t = (Z_s)_{s=0}^t$, and $Z \equiv (Z)_{t=0}^{\infty}$).

Suppose that agent i misreports his initial type to be $\hat{\theta}_{i0} \in \Theta_{i0}$ and then reports the shocks truthfully (and the others report truthfully). The agent's resulting payoff from this is

$$\begin{aligned}\hat{U}_i(\hat{\theta}_{i0}, \theta_{i0}, \theta_{-i,0}, \varepsilon) \;\; &\equiv \;\; U_i(Z(\theta_{i0}, \theta_{-i,0}, \hat{\chi}(\hat{\theta}_{i0}, \theta_{-i,0}, \varepsilon), \varepsilon), \hat{\chi}(\hat{\theta}_{i0}, \theta_{-i,0}, \varepsilon)) \\ &\quad + \hat{\Psi}_i(\hat{\theta}_{i0}, \theta_{-i,0}, \varepsilon).\end{aligned} \tag{5}$$

That the choice rule $\langle\chi, \Psi\rangle$ is BNIC implies that truthful reporting by each agent in each period must constitute a Bayesian-Nash Equilibrium of this new mechanism. In turn, this implies that each

[29] Note that the impulse responses are conditional expectations. As such, they are uniquely defined only with $\lambda[\chi]$-probability one.

agent i cannot improve his expected utility by misreporting his period-0 type and then reporting the subsequent shocks truthfully. That is, for any $\theta_{i0} \in \Theta_{i0}$,

$$V_i^{\langle \chi, \Psi \rangle}(\theta_{i0}) = \sup_{\hat{\theta}_{i0} \in \Theta_{i0}} W(\hat{\theta}_{i0}, \theta_{i0}) = W(\theta_{i0}, \theta_{i0}), \text{ where } W(\hat{\theta}_{i0}, \theta_{i0}) = \mathbb{E}\left[\hat{U}_i(\hat{\theta}_{i0}, \theta_{i0}, \tilde{\theta}_{-i,0}, \tilde{\varepsilon})\right].$$

Now we establish the following crucial lemma (proof in the Appendix).

Lemma 1 *Under the assumptions in the Theorem, for any* $i = 1, .., N$, $\hat{\theta}_{i0} \in \Theta_{i0}$, $W_i(\hat{\theta}_{i0}, \cdot)$ *is equi-Lipschitz continuous and differentiable in* θ_{i0}*, with the derivative at* $\theta_{i0} = \hat{\theta}_{i0}$ *given by*

$$\frac{\partial W_i(\hat{\theta}_{i0}, \hat{\theta}_{i0})}{\partial \theta_{i0}} = \mathbb{E}\left[\sum_{t=0}^{\infty} \frac{\partial U_i\left(Z\left(\hat{\theta}_{i0}, \tilde{\theta}_{-i,0}, \hat{\chi}\left(\hat{\theta}_{i0}, \tilde{\theta}_{-i,0}, \tilde{\varepsilon}\right), \tilde{\varepsilon}\right), \hat{\chi}\left(\hat{\theta}_{i0}, \tilde{\theta}_{-i,0}, \tilde{\varepsilon}\right)\right)}{\partial \theta_{it}} \cdot \frac{\partial Z_{it}\left(\hat{\theta}_{i0}, \hat{\chi}_i^{t-1}\left(\hat{\theta}_{i0}, \tilde{\theta}_{-i,0}, \tilde{\varepsilon}^{t-1}\right), \tilde{\varepsilon}_i^t\right)}{\partial \theta_{i0}} \right].$$

The equi-Lipschitz continuity established in the Lemma implies that the value function, which by BNIC must coincide with the equilibrium payoff $V_i^{\langle \chi, \Psi \rangle}(\theta_{i0})$ is Lipschitz continuous.[30] Furthermore, by Theorem 1 in Milgrom and Segal (2002), at any differentiability point of $V_i^{\langle \chi, \Psi \rangle}(\theta_{i0})$ we must have

$$\frac{dV_i^{\langle \chi, \Psi \rangle}(\theta_{i0})}{d\theta_{i0}} = \frac{\partial W_i(\theta_{i0}, \theta_{i0})}{\partial \theta_{i0}}.$$

Using Lemma 1, the Law of Iterated Expectations, and the definition of impulse responses (4) yields the result. ■

3.1 Interpretation: Impulse Responses

Throughout this section, for simplicity, we momentarily drop the index referring to agent i. The functions $I_t(\theta^t, x^{t-1})$ emerging from Theorem 1 are calculated as derivatives of the period-t signals with respect to the period-0 signal holding the intermediate independent shocks fixed. They can be interpreted as nonlinear impulse responses of the stochastic process. To see this, apply Theorem 1 to a situation with fixed decisions and no payments (i.e., $X_t = \{\hat{x}_t\}$ for each $t = 0, ..., \infty$, and $\Psi(\theta) = 0$ for all $\theta \in \Theta$), in which case the optimization is irrelevant, and we simply have $V^{\langle \chi, \Psi \rangle}(\theta_0) \equiv \mathbb{E}^{\lambda|\theta_0}[U(\tilde{\theta}, \hat{x})]$. Then (3) takes the form

$$\frac{d\mathbb{E}^{\lambda|\theta_0}[U(\tilde{\theta}, \hat{x})]}{d\theta_0} = \mathbb{E}^{\lambda|\theta_0}\left[\sum_{t=0}^{\infty} \frac{\partial U(\tilde{\theta}, \hat{x})}{\partial \theta_t} I_t(\tilde{\theta}^t, \hat{x})\right]. \tag{6}$$

[30] Since for each $\theta_{i0}, \theta'_{i0}$, $\left|V_i^{\langle \chi, \Psi \rangle}(\theta'_{i0}) - V_i^{\langle \chi, \Psi \rangle}(\theta_{i0})\right| \leq \sup_{\hat{\theta}_{i0} \in \Theta_{i0}} \left|W_i\left(\hat{\theta}_{i0}, \theta'_{i0}\right) - W_i\left(\hat{\theta}_{i0}, \theta'_{i0}\right)\right| \leq M\left|\theta'_{i0} - \theta_{i0}\right|$, where $M > 0$ is the constant of equi-Lipschitz continuity of W. This argument is similar to the first part of Milgrom and Segal's (2002) Theorem 2.

Note that the impulse response functions I_t are determined entirely by the stochastic process and should satisfy formula (6) for any utility function U satisfying Conditions U-TD and U-ELC.[31]

Assuming that for each $m \geq 1$, the function $z_m\left(\theta^{m-1}, x^{m-1}, \varepsilon_m\right)$ defined in Definition 2 is totally differentiable in θ^{m-1} for any $\left(x^{m-1}, \varepsilon_m\right)$, we can use the chain rule to calculate the impulse responses inductively on $t \geq 1$, which yields

$$\frac{\partial Z_t(\theta_0, x^{t-1}, \varepsilon^t)}{\partial \theta_0} = \sum_{\substack{K \in \mathbb{N},\ l \in \mathbb{N}^{K+1}: \\ 0 = l_0 < \ldots < l_K = t}} \prod_{k=1}^{K} \frac{\partial z_{l_k}(Z^{l_k - 1}(\theta_0, x^{l_k - 2}, \varepsilon^{l_k - 1}), x^{l_k - 1}, \varepsilon_{l_k})}{\partial \theta_{l_{k-1}}} \text{ for all } t \geq 1. \tag{7}$$

The derivative $\partial z_m / \partial \theta_l$ can be interpreted as the "direct impulse response" of the signal in period m to the signal in period $l < m$. The "total" impulse response $\partial Z_t / \partial \theta_0$ is then obtained by adding up the products of the direct impulse responses over all possible causation chains from period 0 to period t.

The simplest illustration is given by processes in which impulse responses are (time-dependent) constants:

Example 1 (AR(k)) Let θ_t evolve according to an autoregressive (AR) process that is independent of allocations:

$$\tilde{\theta}_t = z_t(\tilde{\theta}^{t-1}) \equiv \sum_{j=1}^{\infty} \phi_j \tilde{\theta}_{t-j} + \tilde{\varepsilon}_t,$$

where $\tilde{\theta}_t = 0$ for any $t < 0$, $\phi_j \in \mathbb{R}$ for any $j \in \mathbb{N}$ and $\tilde{\varepsilon}_t$ is a random variable distributed according to some c.d.f. G_t with support $\mathcal{E}_t \subset \mathbb{R}$ with all the $\tilde{\varepsilon}_t$ drawn independently of each other and of $\tilde{\theta}_0$. Thus, $\langle \mathcal{E}, G, z \rangle$ is a state representation for the kernels. Then the impulse responses (7) take the form

$$I_t = \sum_{\substack{K \in \mathbb{N},\ l \in \mathbb{N}^{K+1}: \\ 0 = l_0 < \ldots < l_K = t}} \prod_{k=1}^{K} \phi_{l_k - l_{k-1}} \text{ for } t \geq 1, \text{ while } I_0 = 1. \tag{8}$$

Note that the impulse responses are constant in this case since the functions z are linear. Condition F-BIR$_0$ then requires that $\|I\| \equiv \sum_{t=0}^{\infty} \delta^t |I_t| < \infty$. E.g., in the special case of an AR(1) process, $\phi_j = 0$ for all $j > 1$, hence $I_t = (\phi_1)^t$, and Condition F-BIR$_0$ is satisfied if and only if $\delta |\phi_1| < 1$. To verify Condition F-BE$_0$, write

$$\begin{aligned} \tilde{\theta}_t &= Z_t(\theta_0, \tilde{\varepsilon}^t) = I_t \theta_0 + \textstyle\sum_{\tau=1}^{t} I_{t-\tau} \tilde{\varepsilon}_\tau \text{ for all } t \geq 0, \text{ and so} \\ \mathbb{E}^{\lambda|\theta_0}\left[||\tilde{\theta}||\right] &\leq \|I\| \cdot |\theta_0| + \textstyle\sum_{t=1}^{\infty} \delta^t \sum_{\tau=1}^{t} |I_{t-\tau}| \cdot \mathbb{E}\left[|\tilde{\varepsilon}_\tau|\right] = \|I\| \cdot \left(|\theta_0| + \sum_{\tau=1}^{\infty} \delta^\tau \cdot \mathbb{E}\left[|\tilde{\varepsilon}_\tau|\right]\right). \end{aligned} \tag{9}$$

Hence, Condition F-BE$_0$ is ensured by assuming, in addition to $\|I\| < \infty$, that $\|\mathbb{E}[|\tilde{\varepsilon}|]\| < \infty$. \\

[31] We conjecture that this property uniquely defines the impulse response functions with probability 1. Thus, the calculation (4) of impulse responses from a state-space representation should yield the same answer with probability 1 regardless of which state-space representation is used.

Inspired by Example 1, we can ensure Conditions F-BE$_0$ and F-BIR$_0$ for more general processes by bounding them with an AR process (proof is straight forward and hence omitted).

Lemma 2 *Suppose that* $(\phi_t)_{t=1}^{\infty} \in \mathbb{R}_+^{\infty}$ *are the coefficients of an AR process whose impulse responses* I_t*, given by (8), satisfy* $\|I\| < \infty$*. Then*

(a) If $\langle \mathcal{E}, G, z \rangle$ *is a state representation of* F *where for each* $m \geq 1$*, each* $(\theta^{m-1}, x^{m-1}, \varepsilon_m) \in \Theta^{m-1} \times X^{m-1} \times \mathcal{E}_m$*, the function* $z_m\left(\theta^{m-1}, x^{m-1}, \varepsilon_m\right)$ *is totally differentiable in* θ^{m-1}*, with*

$$\left|\partial z_m\left(\theta^{m-1}, x^{m-1}, \varepsilon_m\right) / \partial \theta_l\right| \leq \phi_{m-l}$$

for all $l < m$*, then Condition F-BIR*$_0$ *holds.*

(b) If, for all $t \geq 1$*, all* $\left(\theta^{t-1}, x^{t-1}\right) \in \Theta^{t-1} \times X^{t-1}$*,* $\mathbb{E}^{F_t\left(\theta^{t-1}, x^{t-1}\right)}\left[\left|\tilde{\theta}_t\right|\right] \leq \sum_{j=1}^{t} \phi_j \left|\theta_{t-j}\right| + E_t$*, with* $\|E\| < \infty$*, then Condition F-BE*$_0$ *holds .*

3.2 Canonical State Representation

As promised earlier, we now show that *any* process admits at least one state representation, and use it to derive a simple formula for the impulse responses via derivatives of the kernels' c.d.f.'s:[32]

Proposition 1 (Canonical representation) *(a) The following is a State Representation (henceforth the* canonical representation*) of* F*: For each* $i = 1, ..., N$*, each* $t \geq 0$*,* $\mathcal{E}_{it} = (0,1)$*,* G_{it} *is the uniform distribution on* $(0,1)$ *and, for any* $(\theta_i^{t-1}, x_i^{t-1}) \in \Theta_i^{t-1} \times X_i^{t-1}$*,* $\varepsilon_{it} \in (0,1)$*,*

$$z_{it}(\theta_i^{t-1}, x_i^{t-1}, \varepsilon_{it}) = F_{it}^{-1}(\varepsilon_{it}|\theta_i^{t-1}, x_i^{t-1}) \equiv \inf\{\theta_{it} : F_{it}(\theta_{it}|\theta_i^{t-1}, x_i^{t-1}) \geq \varepsilon_{it}\}. \tag{10}$$

(b) Furthermore, if for all $i = 1, ..., N$*, all* $t \geq 1$*, all* $x_i^{t-1} \in X_i^{t-1}$*, the c.d.f.* $F_{it}(\theta_{it}|\theta_i^{t-1}, x_i^{t-1})$ *is continuously differentiable in* $(\theta_{it}, \theta_i^{t-1})$*, then for all* $i = 1, ..., N$*,* $t \geq 1$*,* $\chi \in \mathcal{X}$*,* $\theta_{i0} \in \Theta_{i0}$*,* $\lambda_i[\chi]\,|\theta_{i0}$*-almost all* (θ_i^t, x_i^{t-1})*, the impulse responses of the canonical representation take the form*

$$I_{it}(\theta_i^t, x_i^{t-1}) = \sum_{\substack{K \in \mathbb{N},\ l \in \mathbb{N}^{K+1}: \\ 0 = l_0 < ... < l_K = t}} \prod_{k=1}^{K} \left[-\frac{\partial F_{il_k}(\theta_{il_k}|\theta_i^{l_k-1}, x_i^{l_k-1})/\partial \theta_{il_{k-1}}}{f_{il_k}(\theta_{il_k}|\theta_i^{l_k-1}, x_i^{l_k-1})} \right]$$

where $f_{it}(\theta_{it}|\theta_i^{t-1}, x_i^{t-1}) \equiv \partial F_{it}(\theta_{it}|\theta_i^{t-1}, x_i^{t-1})/\partial \theta_{it}$ *is the density function of* $F_{it}(\theta_{it}|\theta_i^{t-1}, x_i^{t-1})$*.*

Proof. (a) By the probability integral transform theorem (see, e.g., Angus (1984)), for any $i = 1, ..., N$, any $t \geq 0$, any $(\theta_i^{t-1}, x_i^{t-1}) \in \Theta_i^{t-1} \times X_i^{t-1}$, the random variable $F_{it}^{-1}(\tilde{\varepsilon}_{it}|\theta_i^{t-1}, x_i^{t-1})$ is distributed according to the c.d.f. $F_{it}(\cdot|\theta_i^{t-1}, x_i^{t-1})$ thus proving that $\langle \mathcal{E}, G, z \rangle$ is a state representation of F.

[32] The construction in part (a) of Proposition 1 is similar to a standard proof of the Kolmogorov existence theorem—see the Second proof for countable T in Billingsley (1995, p.490).

(b) By continuity, (10) implies the identity

$$F_{im}(z_{im}(\theta_i^{m-1}, x_i^{m-1}, \varepsilon_{im})|\theta_i^{m-1}, x_i^{m-1}) = \varepsilon_{im}.$$

Restricting attention to $(\theta_i^{m-1}, x_i^{m-1}, \varepsilon_{im})$, for which $f_{im}(z_{im}(\theta_i^{m-1}, x_i^{m-1}, \varepsilon_{im})|\theta_i^{m-1}, x_i^{m-1}) > 0$, which happens with $\lambda_i[\chi]|\theta_{i0}$-probability 1 for any $\chi \in \mathcal{X}$, any $\theta_{i0} \in \Theta_{i0}$, the Implicit Function Theorem yields

$$\frac{\partial z_{im}(\theta_i^{m-1}, x_i^{m-1}, \varepsilon_{im})}{\partial \theta_{is}} = -\left.\frac{\partial F_{im}(\theta_{im}|\theta_i^{m-1}, x_i^{m-1})/\partial \theta_{is}}{f_{im}(\theta_{im}|\theta_i^{m-1}, x_i^{m-1})}\right|_{\theta_{im}=z_{im}(\theta_i^{m-1}, x_i^{m-1}, \varepsilon_{im})}.$$

Now use (7) and (4) to calculate the impulse responses. ■

4 The Relaxed Program and Distortions

In a static setting, the envelope formula for the agents' equilibrium payoffs permits one to calculate the agents' information rents, providing a useful tool for designing optimal mechanisms. We now show how this approach extends to a dynamic setting. We start by using ICFOC$_0$ to express the agents' information rents entirely in terms of the allocation rule χ. We then use this representation to set up a "relaxed program" incorporating only a subset of all the relevant constraints whose solution provides a candidate for the optimal allocation rule. In the next section, we will provide sufficient conditions for a solution to this relaxed program to satisfy all of the constraints and thus constitute an optimal mechanism. Before doing so, at the end of this section, we discuss how the allocation rule that solves the relaxed program differs from the efficient allocation rule and how the discrepancy between the two can be understood in terms of the impulse responses of the kernels.

Suppose that, in addition to the N agents, there is a "principal" (labeled "agent 0") who designs the mechanism and whose payoff takes the quasilinear form $U_0(x, \theta) - \sum_{i=1}^{N} P_i$ for some function $U_0 : \Theta \times X \to \mathbb{R}$. As standard in the literature, we assume that the principal makes a take-it-or-leave-it offer of a mechanism to the agents in period zero after each agent i has observed his initial type θ_{i0}. At that point, each agent can either accept the mechanism or reject it. By rejecting the mechanism, an agent obtains his reservation payoff, which for simplicity we normalize to zero for all agents and types.[33] For the mechanism to be accepted by all types of all agents, the equilibrium

[33] If an agent can accept the mechanism but then quit at a later stage, this would give rise to participation constraints in subsequent periods. However, in a quasilinear setting with unlimited transfers, the principal could ask the agent to post a sufficiently large bond upon acceptance, to be repaid later, so as to make it unprofitable to quit at any time during the mechanism. For this reason, we ignore participation constraints in all periods other than the initial period. Note, however, that in non-quasilinear settings where agents have a consumption-smoothing motive, bonding is costly, and so participation constraints may bind in many periods – see, e.g., Hendel and Lizzerri (2003) and Garrett and Pavan (2011b).

payoffs must satisfy the following participation constraints:[34]

$$V_i^{\langle\chi,\Psi\rangle}(\theta_{i0}) \geq 0 \quad \text{for all } i = 1, ..., N, \text{ all } \theta_{i0} \in \Theta_{i0}. \tag{11}$$

The principal's problem can then be described as choosing a BNIC choice rule $\langle\chi, \Psi\rangle$ that also satisfies participation constraints (11) to maximize her expected payoff.

Since the principal's problem is quite complex, we formulate a simpler problem in which the constraints on the choice rule are relaxed: Instead of requiring that the choice rule be BNIC, we only require it satisfy ICFOC$_0$. Also, instead of imposing all the participation constraints (11), we only imposes the constraints for each agent i's lowest type $\underline{\theta}_{i0}$, which we assume to be finite:

$$V_i^{\langle\chi,\Psi\rangle}(\underline{\theta}_{i0}) \geq 0 \text{ for all } i = 1, ..., N. \tag{12}$$

Formally, the Relaxed Program is stated as follows:

$$\max_{\chi\in\mathcal{X},\Psi:\Theta\to\mathbb{R}^N} \mathbb{E}^{\lambda[\chi]}[U_0(\tilde{\theta}, \chi(\tilde{\theta})) - \sum_{i=1}^{N} \Psi_i(\tilde{\theta})] \text{ s.t. ICFOC}_0 \text{ and (12)}. \tag{13}$$

Assuming that each distribution F_{i0} is absolutely continuous with density $f_{i0}(\theta_{i0}) > 0$ for all $\theta_{i0} \in \Theta_{i0}$ and integrating by parts, we can then use ICFOC$_0$ to express the ex-ante expectation of the transfers to the agents in terms of the allocation rule χ and of the expected equilibrium payoffs $V_i^{\langle\chi,\Psi\rangle}(\underline{\theta}_{i0})$ of the lowest period-0 types. Denoting by $\eta_{i0}(\theta_{i0}) \equiv f_{i0}(\theta_{i1})/(1 - F_{i0}(\theta_{i1}))$ the hazard rate of the distribution F_{i0} of each agent $i = 1, ..., N$, we then have the following result (proof is straight forward and hence omitted):

Proposition 2 (Principal's payoff) *Assume that Conditions U-TD, U-ELC, F-BE$_0$, F-BIR$_0$ hold, and $\underline{\theta}_{i0} > -\infty$ for all $i = 1, ..., N$. Then the principal's expected payoff under any BNIC choice rule $\langle\chi, \Psi\rangle$ is given by*

$$\begin{aligned}\mathbb{E}^{\lambda[\chi]}[U_0(\tilde{\theta}, \chi(\tilde{\theta})) - \sum_{i=1}^{N} \Psi_i(\tilde{\theta})] &= \mathbb{E}^{\lambda[\chi]}\left[\sum_{i=0}^{N} U_i(\tilde{\theta}, \chi(\tilde{\theta})) - \sum_{i=1}^{N} \frac{1}{\eta_{i0}(\tilde{\theta}_{i0})} \sum_{t=0}^{\infty} \frac{\partial U_i(\tilde{\theta}, \chi(\tilde{\theta}))}{\partial \theta_{it}} I_{it}(\tilde{\theta}_i^t, \chi_i^{t-1}(\tilde{\theta}))\right] \\ &\quad - \sum_{i=1}^{N} V_i^{\langle\chi,\Psi\rangle}(\underline{\theta}_{i0}).\end{aligned}$$

Finally, note that for any allocation rule χ, it is possible to construct a transfer rule to satisfy ICFOC$_0$ and to make the lowest types' participation constraints (12) bind (e.g., by taking $\Psi_i(\theta) = \int_{\underline{\theta}_{i0}}^{\theta_{i0}} V_i(q)\, dq - U_i(\theta, \chi(\theta))$, where $V_i(\theta_{i0})$ is given by the right hand side of (3)). It follows that a

[34] As usual, it is without loss of generality to restrict the principal to offer mechanisms that are always accepted in equilibrium, as long as the non-participation option can be made available within the mechanism.

feasible allocation rule χ is part of a solution to the Relaxed Program if and only if it maximizes the "*expected dynamic virtual surplus*"

$$\mathbb{E}^{\lambda[\chi]}\left[\sum_{i=0}^{N}U_i(\tilde{\theta},\chi(\tilde{\theta}))-\sum_{i=1}^{N}\frac{1}{\eta_{i0}(\tilde{\theta}_{i0})}\sum_{t=0}^{\infty}\frac{\partial U_i(\tilde{\theta},\chi(\tilde{\theta}))}{\partial\theta_{it}}I_{it}(\tilde{\theta}_i^t,\chi_i^{t-1}(\tilde{\theta}))\right] \tag{14}$$

among all feasible allocation rules.

4.1 Distortions

We now proceed to study allocation rules solving the relaxed program. Specifically, we compare these rules to the efficient allocation rules, i.e., to those maximizing ex-ante expected surplus

$$\mathbb{E}^{\lambda[\chi]}\left[\sum_{i=0}^{N}U_i(\tilde{\theta},\chi(\tilde{\theta}))\right]. \tag{15}$$

Of course, this analysis is useful only insofar as the solution to the relaxed program can be shown to satisfy all the constraints of the full program. We will give some sufficient conditions for this in Section 5 below. In particular, these conditions cover the examples discussed below.

Similarly to the static setting, the principal introduces distortions to reduce the agents' expected information rents. In contrast to the static setting, however, the expected information rent of agent i is now given by the expected present value of

$$\frac{1}{\eta_{i0}(\theta_{i0})}\frac{\partial U_i(\theta,x)}{\partial\theta_{it}}I_{it}(\theta_i^t,x_i^{t-1}),$$

and so it is determined not just by the properties of his utility function but also by those of the impulse responses. We illustrate the dependence on impulse responses with a few examples that help understand the findings of the earlier literature for the specific processes being considered, and the extent to which those findings generalize. For simplicity we focus on the case of $N=1$ and drop the subscript $i=1$ from each X_{it} and Θ_{it}.

Example 2 (Markov process) For all $t\geq 0$, $F_t\left(\theta^{t-1},x^{t-1}\right)$ depends only on θ_{t-1} (and hence denoted by $F_t(\theta_{t-1})$), $-\infty<\underline{\theta}_t<\bar{\theta}_t<+\infty$, $Supp[F_t(\theta_{t-1})]=\Theta_t$ all $\theta_{t-1}\in\Theta_{t-1}$, $X_t=\left[0,a+\bar{\theta}_t\right]$, with $a\in\mathbb{R}_{++}$, $U_1(\theta,x)=\sum_{t=0}^{\infty}\delta^t(a+\theta_t)x_t$ and $U_0(\theta,x)=-\sum_{t=0}^{\infty}\delta^t x_t^2/2$. The efficient allocation rule is then given by $\chi_t^*\left(\theta^t\right)=a+\theta_t$ for all $t\geq 0$. The period-t impulse response can be calculated from the canonical state representation of Proposition 1, assuming that each c.d.f. $F_t(\theta_t|\theta_{t-1})$ is continuously differentiable over $\Theta_t\times\Theta_{t-1}$:

$$I_t(\theta^t)=\prod_{\tau=1}^{t}\left[-\frac{\partial F_\tau(\theta_\tau|\theta_{\tau-1})/\partial\theta_{\tau-1}}{f_\tau(\theta_\tau|\theta_{\tau-1})}\right]\text{ for all }t\geq 1.$$

To understand the formula, note that when the kernels are Markov, the only causation chain from the initial type to the period-t type is the one that passes through all the adjacent intermediate periods, and so the impulse response is the product of the direct adjacent impulse responses. The solution to the relaxed program then takes the form

$$\chi_t\left(\theta^t\right)=\max\left\{0,a+\theta_t-\frac{1}{\eta_0\left(\theta_0\right)}\cdot\prod_{\tau=1}^{t}\left[-\frac{\partial F_\tau(\theta_\tau|\theta_{\tau-1})/\partial\theta_{\tau-1}}{f_\tau(\theta_\tau|\theta_{\tau-1})}\right]\right\}\text{ for all }t\geq 0.$$

Suppose that the kernels satisfy First-Order Stochastic Dominance, i.e., $F_t(\theta_t|\cdot)$ is nonincreasing in θ_{t-1} for all $t\geq 1$, all $\theta_t\in\Theta_t$. In that case, the impulse responses are all nonnegative, and therefore consumption is always weakly below the efficient level. Also, under the standard assumption that the period-0 hazard rate $\eta_0(\theta_0)$ is nondecreasing, the distortions in all periods vanish at the highest initial type $\theta_0=\bar{\theta}_0$, paralleling the classical finding of the static screening model. However, the dependence of distortions on subsequent types hinges on the more delicate properties of the impulse response functions. For example, note that $F_\tau(\bar{\theta}_\tau|\theta_{\tau-1})=1$ and $F_\tau(\underline{\theta}_\tau|\theta_{\tau-1})=0$ for any $\theta_{\tau-1}\in\Theta_{\tau-1}$, and therefore $\partial F_\tau(\theta_\tau|\theta_{\tau-1})/\partial\theta_{\tau-1}=0$ for $\theta_\tau\in\left\{\underline{\theta}_\tau,\bar{\theta}_\tau\right\}$, all $\tau\geq 1$. Therefore, if for any $\tau\geq 1$, any $\theta_{\tau-1}\in\Theta_{\tau-1}$, $f_\tau(\underline{\theta}_\tau|\theta_{\tau-1}), f_\tau(\bar{\theta}_\tau|\theta_{\tau-1})>0$, we see that the impulse response in period t turns to zero at both the highest and the lowest type. Due to the Markov nature of the process, this severs the causal effect of the initial type and eliminates distortions from that period onward. On the other hand, for intermediate types, distortions could be strict, in which case consumption in period t is nonmonotone in the intermediate types $\theta_2,...,\theta_{t-1}$. \\

Example 3 (Nonlinear AR process) Consider the following family of Markov processes, which specializes the one in the previous example: $\Theta_t=[0,1]$ for all t, and for any $t\geq 1$, $\theta_t=\phi\left(\theta_{t-1}\right)+\varepsilon_t$, where $\phi\left(\cdot\right)$ is an increasing differentiable function with $\phi\left(0\right)=0$ and $\phi\left(1\right)<1$, and the shocks ε_t are independent over time, with support $[0,1-\phi\left(1\right)]$. The impulse responses then take the form $I_t(\theta^t)=\prod_{\tau=0}^{t-1}\phi'\left(\theta_\tau\right)$ (this can be derived from the given state representation, or from the canonical representation assuming that the shock distributions G_t are differentiable), and the solution to the relaxed problem is given by

$$\chi_t\left(\theta^t\right)=\max\left\{0,a+\theta_t-\frac{1}{\eta_0\left(\theta_0\right)}\cdot\prod_{\tau=0}^{t-1}\phi'\left(\theta_\tau\right)\right\}\text{ for all }t.$$

Note, in particular, that χ exhibits downward distortions since $\phi'>0$. In this setting, increasing type θ_τ in a period $\tau\geq 1$ reduces distortions in subsequent periods if the function $\phi\left(\cdot\right)$ is concave, but increases these distortions if the function is convex. When $\phi\left(\cdot\right)$ is neither concave nor convex, distortions are nonmonotone in previous types. \\

Besanko (1985) considers the special case of this setting in which $\phi\left(\theta\right)$ is linear in θ. In this case, the allocation in period t depends only on θ_0 and θ_t and not on the intermediate types $\theta_2,...,\theta_{t-1}$.

As the example shows, for more general functions ϕ, the allocation in period t may be increasing, decreasing, or nonmonotone in the intermediate types.

Example 4 (Discrete Types) Next consider a setting where the type spaces Θ_t are discrete, to which the first-order approach is not directly applicable, but can be adapted by focusing on local downward incentive constraints instead of ICFOC and using discrete versions of the impulse responses. For simplicity, we begin with the setting of Battaglini (2005), who considers a Markov process over the binary type space $\Theta_t = \{L, H\}$ for each t. For this setting, we can use the following state representation $\langle \mathcal{E}, G, z \rangle$: in each period $t \geq 1$, G_t is the uniform distribution on $\mathcal{E}_t = (0, 1)$, and $z_t\left(\theta^{t-1}, \varepsilon_t\right) = H$ if $\varepsilon_t > 1 - q_{\theta_{t-1}}$ and $z_t\left(\theta^{t-1}, \varepsilon_t\right) = L$ otherwise. This induces a Markov process on the types with transition probabilities $\Pr\left\{\tilde{\theta}_t = H | \theta_{t-1}\right\} = q_{\theta_{t-1}}$, and the assumption $q_H > q_L$ ensures positive serial correlation. In this setting, the discrete one-period-ahead impulse response can be defined as $I\left(\theta_{t-1}, \theta_t\right) = \frac{1}{H-L}\mathbb{E}\left[z\left(H, \tilde{\varepsilon}_t\right) - z\left(L, \tilde{\varepsilon}_t\right) | z\left(\theta_{t-1}, \tilde{\varepsilon}_t\right) = \theta_t\right]$, i.e., the expected effect of the previous type type being H rather than L on the current type, given the observed previous and current types (θ_{t-1}, θ_t). It is then easy to see that $I\left(\theta_{t-1}, \theta_t\right) = 0$ whenever the type switches, i.e., $\theta_t \neq \theta_{t-1}$. For example, when type switches from $\theta_{t-1} = L$ to $\theta_t = H$, this means that $\varepsilon_t > 1 - q_L$, and therefore $\varepsilon_t > 1 - q_H$, hence the new type would also have been H had the previous type been H. Similarly, when the type switches from H to L, the new type would also have been L had the previous type been L. Due to the Markov nature of the process, the impulse response of the period-t type to the period-0 type takes the form $I_t\left(\theta^t\right) = \prod_{\tau=0}^{t-1} I\left(\theta_\tau, \theta_{\tau+1}\right)$, and therefore as soon as the type switches, the causal effect of the initial type is severed, ensuring efficiency from that point onward.

Observe also that the solution to the relaxed program is efficient when the initial type is H, since only type H's incentive constraint is considered in the relaxed program. These arguments yield Battaglini's "Generalized No Distortion at the Top Principle" (GNDTP): any switch to H yields efficiency from that period onward. Battaglini's other conclusion is the "Vanishing Distortions at the Bottom Principle" (VDBP): the distortions for histories $(L, ..., L)$ shrink with time. This conclusion can be understood by noting that the impulse response at $(\theta_{t-1}, \theta_t) = (L, L)$ is less than 1: indeed, for some of the shocks that leave type L unchanged, had the previous type been H instead of L, the current type would have been H.

This logic also demonstrates that GNDTP extends to any discrete type process satisfying FOSD: indeed, when type θ_{t-1} switches to the highest possible type $\theta_t = \bar{\theta}_t$, this implies that any previous type $\theta'_{t-1} > \theta_{t-1}$ would have also switched to $\bar{\theta}_t$. Since the relaxed program considers only downward adjacent incentive constraints, this means a period-$(t-1)$ type θ'_{t-1} who pretended to be θ_{t-1} is no longer distinguishable from θ_{t-1}, in which case distortions should be eliminated forever. This logic is similar to the finding of Example 2 for the case of positive densities at endpoints. However, in the discrete version of Example 2, the distortions need not be eliminated by switching from θ_{t-1} to the *lowest* possible type $\underline{\theta}_t$. Indeed, since it is the *downward* local IC constraints that bind in the relaxed program, and since types *above* θ_{t-1} may not have switched to $\underline{\theta}_t$ by experiencing the same shocks thus remaining distinguishable, this creates a reason to distort after reporting $\underline{\theta}_t$. We conclude that VDBP also generalizes to multiple discrete types.

Nonetheless, as the number of discrete types increases, the results become qualitatively closer to the continuous-type case than to Battaglini's two-type case, since the binding downward adjacent IC constraints start approximating the ICFOC constraints. For example, while GNDTP and VDPB still hold, their relevance decreases as the number of types grows, for these properties apply only to histories that occur with a small probability. Furthermore, as the distance between the two lowest types vanishes, distortions "at the bottom" become arbitrarily small immediately from $t = 1$. For intermediate types, distortions can in general be nonmonotonic in type, as in the previous two examples. Lastly note that distortions need not be shrinking over time: E.g., in Example 3, whenever $\phi'(\theta_t) > 1$, distortions in period $t + 1$ exceed those in period t. \\

The one property that is robust in all of the above examples is that of downward distortions. Now we establish sufficient conditions for this property to hold in a general setting. Intuitively, the principal introduces downward distortions to reduce the agents' information rents when these rents are increasing in the allocations x. To ensure this, we make the usual assumption that higher types have higher marginal utilities, extending it to allow for multidimensional allocations and types:

Condition 5 (U-SCP) *Utility Single-Crossing Property: X is a partially ordered set, and for each $i = 1, ..., N$, $U_i(\theta, x)$ has increasing differences in (θ_i, x).*

In addition, we make the following assumptions on the kernels:

Condition 6 (F-FOSD) *Process First-Order Stochastic Dominance: For all $i = 1, ..., N$, $t \geq 1$, $\theta_{it} \in \Theta_{it}$, $x_i^{t-1} \in X_i^{t-1}$, $F_{it}\left(\theta_{it}|\theta_i^{t-1}, x_i^{t-1}\right)$ is nonincreasing in θ_i^{t-1}.*

Condition 7 (F-AUT) *Process Decision-Autonomous: For all $i = 1, ..., N$, $t \geq 1$, $\theta_i^{t-1} \in \Theta_i^{t-1}$, the distribution $F_{it}\left(\theta_i^{t-1}, x_i^{t-1}\right)$ does not depend on x_i^{t-1}.*

Condition F-AUT implies that the impulse responses do not depend on the allocations and hence can be written as $I_{it}(\theta_i^t)$, while F-FOSD implies that the impulse responses are nonnegative (in Example 3, F-FOSD was implied by the assumption that $\phi'(\theta) \geq 0$). Intuitively, F-FOSD means that the type in each period $t > 1$ is positively linked to the period-0 type, and so making the agent's utility less sensitive to his future types reduces his ex-ante information rent. Finally, we also need to ensure that the different dimensions of the allocations are (weakly) complementary to each other in the objective function.

Condition 8 (U-COMP) *Utility Complementarity: X is a lattice,*[35] *and for each $i = 0, ..., N$, $\theta \in \Theta$, $U_i(\theta, x)$ is supermodular in x, and for all $i = 1, ..., N$, all $t \geq 0$, $-\partial U_i(\theta, x)/\partial\theta_{it}$ is supermodular in x.*

In particular, Condition U-COMP holds weakly in the special case where $X_t \subset \mathbb{R}$ in each period t and the payoffs $U_i(\theta, x)$ are additively separable in x_t, as in the examples above. More generally, U-COMP also allows for strict complementarity across time, e.g. as in "learning-by-doing" models where a higher production today reduces the marginal cost tomorrow. We then have the following result:

Proposition 3 (Downward distortions) *Suppose that Conditions F-AUT, F-FOSD, U-SCP, and U-COMP hold. If the allocation rule χ solves the Relaxed Program and the allocation rule χ^* is efficient, then the allocation rule given by $\chi_t(\theta^t) \wedge \chi_t^*(\theta^t)$ for all $t \geq 0$, all $\theta^t \in \Theta^t$ solves the Relaxed Program and the allocation rule given by $\chi_t(\theta^t) \vee \chi_t^*(\theta^t)$ for all $t \geq 0$, all $\theta^t \in \Theta^t$ is efficient.*

Proof. Since X is a lattice, the set $\mathcal{X}$ of feasible allocation rules is also a lattice with the meet and join operations defined pointwise (i.e. for each θ). Define $g : \mathcal{X} \times \{-1, 0\} \to \mathbb{R}$ as

$$g(\chi, q) \equiv \mathbb{E}^\lambda \left[\sum_{i=0}^{N} U_i(\tilde{\theta}, \chi(\tilde{\theta})) + q \sum_{i=1}^{N} \frac{1}{\eta_{i0}(\tilde{\theta}_{i0})} \sum_{t=1}^{\infty} I_{it}(\tilde{\theta}_i^t) \frac{\partial U_i(\tilde{\theta}, \chi(\tilde{\theta}))}{\partial \theta_{it}} \right].$$

Then $g(\chi, 0)$ is the expected total surplus and $g(\chi, -1)$ is the expected virtual surplus. Condition F-AUT ensures that the stochastic process $\lambda[\chi]$ doesn't depend on χ and that each $I_{it}(\theta_i^t, x_i^{t-1})$ does not depend on x_i^{t-1}, which is reflected in the formula. Condition F-FOSD ensures that each $I_{it}(\theta_i^t) \geq 0$. Together with Condition U-SCP, this ensures that g has increasing differences in (χ, q). Together with U-COMP, this ensures that g is supermodular in χ. The result then follows from Topkis's Theorem (see, e.g., Topkis (1998)). ■

In particular, the Proposition implies that if either the solution χ to the relaxed program or the efficient solution χ^* is uniquely defined with probability one, then $\chi(\theta) \leq \chi^*(\theta)$ with probability one. (More generally, it means that the set of allocation rules solving the relaxed program is below the set of efficient allocation rules in the strong set order.)

Downward distortions appear in most applications featuring a single-agent, where Conditions F-AUT, F-FOSD, U-SCP, and U-COMP are often implicitly assumed (for an example that violates F-FOSD and where distortions can indeed be upward, see Courty and Li (2000)). On the contrary, Condition U-COMP is less likely to be satisfied in settings featuring multiple-agents, because of possible capacity constraints that prevent the choice set X to be a lattice. As a result, distortions

[35] The assumption that X is a lattice is not innocuous when $N > 1$: 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).

need not be downward in these applications, even if the other conditions in Proposition 3 are met (an example of a multi-agent setting violating Conditions F-AUT and U-COMP and where distortions can be upward is the bandit-auction in Section 6).

5 PBE-Implementability

We now provide sufficient conditions for the solution to the relaxed program (more generally, for a given choice rule $\langle \chi, \Psi \rangle$) to be implementable in a Perfect Bayesian Equilibrium of a direct revelation mechanism.[36] In such a mechanism, in each period t, each agent i, after observing his type $\theta_{it} \in \Theta_{it}$, submits a report $\hat{\theta}_{it} \in \Theta_{it}$. Given the reports $\hat{\theta}^t$, the mechanism then implements the allocation $\chi_{it}(\hat{\theta}^t)$ for each agent i, which is observed by the agent (and no other information is disclosed to the agent). The mechanism also makes payments with NPV $\Psi_i(\hat{\theta})$ to each agent i. Below, we will show how these payments can be spread over time in a way that they can be disclosed to the agent without affecting his incentives to report truthfully.[37]

Recall that each agent i's reporting strategy in the mechanism specifies a report $\hat{\theta}_{it}$ in each period t as a function $\sigma_{it}(\theta_i^t, \hat{\theta}_i^{t-1}, x_i^{t-1})$ of the true type history θ_i^t, the reported type history $\hat{\theta}_i^{t-1}$, and the allocation history x_i^{t-1}. For the direct mechanism to implement the choice rule, we restrict attention to strategies that are *on-path truthful*, i.e., specify truthtelling whenever the agent has been truthful in the past: $\sigma_{it}(\theta_i^t, \theta_i^{t-1}, x_i^{t-1}) = \theta_{it}$ for all $i = 1, ..., N$, all $\theta_i^{t-1} \in \Theta_i^{t-1}$, all $x_i^{t-1} \in X_i^{t-1}$.

In addition, the specification of a PBE requires describing each agent i's beliefs at each of his information set $(\theta_i^t, \hat{\theta}_i^{t-1}, x_i^{t-1})$ about the unobserved past moves of Nature (θ_{-i}^{t-1}) and of the other agents $(\hat{\theta}_{-i}^{t-1})$. (The agent's beliefs about the *contemporaneous* types of agents $j \neq i$ then follow by applying the kernels $F_{jt}(\theta_j^{t-1}, \chi_j^{t-1}(\hat{\theta}_i^{t-1}, \hat{\theta}_{-i}^{t-1}))$.) We restrict these beliefs to satisfy two natural conditions:

B(i) For any $i = 1, ..., N$, any $t \geq 0$, agent i's beliefs at each time-t history $(\theta_i^t, \hat{\theta}_i^{t-1}, x_i^{t-1}) \in \Theta_i^t \times \Theta_i^{t-1} \times X_i^{t-1}$ are independent of his true type history θ_i^t.

B(ii) For any $i = 1, ..., N$, any $t \geq 0$, agent i's beliefs at each time-t history $(\theta_i^t, \hat{\theta}_i^{t-1}, x_i^{t-1}) \in \Theta_i^t \times \Theta_i^{t-1} \times X_i^{t-1}$ assign probability one to the other agents having reported truthfully, i.e., to the event that $\hat{\theta}_{-i}^{t-1} = \theta_{-i}^{t-1}$.

[36] Note that while we do provide sufficient conditions for PBE implementability in direct mechanisms, we do not claim that restricting attention to direct mechanisms is without loss of generality for the PBE solution concept.

[37] Our focus is on mechanisms with minimal information disclosure as this maximizes the set of implementable choice rules. However, all of our results apply verbatim to more general information disclosure policies by simply reinterpreting x_i to contain all information disclosed by the mechanism to agent i. (For example, one can let $X_{it} = \hat{X}_{it} \times \Theta_{-i}^t$, where $\hat{X}_{it}$ is the set of physical allocations, and Θ_{-i}^t is a "vocabulary." Then by an appropriate choice of χ the mechanism can reveal to agent i any information about the other agents' messages in periods $s = 0, \ldots, t$. Even random disclosure can be allowed by conditioning on reports of a fictitious agent.)

Condition B(i) is in the same spirit as condition B(i) in Fudenberg and Tirole (1991, p.333). It is motivated by the fact that, given agent i's reports $\hat{\theta}_i^{t-1}$ and observed allocations x_i^{t-1}, the distribution of his true types θ_i^t is independent of the other agents' types or reports (since in each period $\tau \leq t$, $\theta_{i\tau}$ is drawn from distribution $F_{i\tau}\left(\theta_i^{\tau-1}, x_i^{\tau-1}\right)$, independently of the other agents' types). This condition will be important for deriving the envelope formula below (if agent i's beliefs over θ_{-i}^{t-1} could vary with θ_i^t, his information rent could be extracted as in Cremer and McLean (1988)). Condition B(ii) in turn says that agent i always believes that his opponents have been following their equilibrium strategies. This condition is also natural — e.g., agents could be constrained to send reports in the support of their kernels, in which case no report history would be inconsistent with truthtelling. Furthermore, with continuous distributions, any particular type profile has zero probability, which explains why condition B(ii) cannot be derived from Bayes' rule and has to be imposed.

Under the two conditions, we can describe agent i's beliefs as contingent probability distributions $\Gamma_{it} : \Theta_i^{t-1} \times X_i^{t-1} \rightarrow \Delta\left(\Theta_{-i}^{t-1}\right)$, where $\Gamma_{it}(\hat{\theta}_i^{t-1}, x_i^{t-1})$ represents agent i's beliefs over the other agents' past types (=reports) θ_{-i}^{t-1} given that he reported $\hat{\theta}_i^{t-1}$ and observed the allocations $x_i^{t-1} = \chi_i^{t-1}(\hat{\theta}_i^{t-1}, \theta_i^{t-1})$. Hereafter, we denote by $\mathbf{\Gamma}$ the set of beliefs satisfying the properties above.

Definition 4 (On-path truthful PBE) *We say that an on-path truthful strategy profile σ and a belief system Γ form an on-path truthful PBE of the direct mechanism $\langle \chi, \Psi \rangle$ if (a) beliefs Γ are consistent with Bayes' rule on all positive-probability events, and (b) each agent's strategy maximizes his expected payoff at each information set given the beliefs Γ, assuming that the other agents follow their equilibrium strategies (i.e., in this case, continue to report truthfully).*

Part (a) of the definition simply means that the belief system is a system of regular conditional probability distributions. The existence of a system of regular conditional probability distributions is well known (see, e.g., Dudley (2002)). Given an allocation rule χ and a belief system Γ, we let $\lambda_i[\chi, \Gamma]|\theta^{t-1}, \theta_{it}$ denote the stochastic process over Θ from the eyes of agent i in period t when the agent has reported truthfully in the past, the complete report history is θ^{t-1}, and agent i's current type is θ_{it}. Formally, $\lambda_i[\chi, \Gamma]|\theta^{t-1}, \theta_{it}$ is the unique probability measure over Θ obtained by first drawing $\tilde{\theta}_{-i}^{t-1}$ from the beliefs $\Gamma_{it}\left(\theta_i^{t-1}, \chi_i^t\left(\theta^{t-1}\right)\right)$, then drawing each $\tilde{\theta}_{jt}$, $j \neq i$, from the distribution $F_{jt}(\theta_j^{t-1}, \chi_j^{t-1}(\theta_i^{t-1}, \theta_{-i}^{t-1}))$, and then applying the kernels F and the allocation rule χ starting with the period-t state $(\theta_i^t, \left(\theta_{-i}^{t-1}, \theta_{-i,t}\right))$.

From now on, we take the belief system Γ as given and focus on part (b) of the PBE definition. In particular, this part requires that truthtelling remains optimal for each agent at all *on-path* information sets, i.e., at those information sets reached via truthful reporting:

Definition 5 (On-path BIC) *Fix $i \in \{1, ..., N\}$ and $s \geq 0$. The choice rule $\langle \chi, \Psi \rangle$ is on-path Bayes Incentive Compatible for agent i in period s, given the belief system Γ, if for each $(\theta^{s-1}, \theta_{is}) \in$*

$\Theta^{s-1} \times \Theta_{is}$, *agent i's equilibrium time-s expected payoff*

$$V_{is}^{\langle \chi, \Psi \rangle, \Gamma}(\theta^{s-1}, \theta_{is}) \equiv \mathbb{E}^{\lambda_i[\chi,\Gamma]|\theta^{s-1},\theta_{is}} \left[U_i(\chi(\tilde{\theta}), \tilde{\theta}) + \Psi_i(\tilde{\theta}) \right] \tag{16}$$

cannot be raised by the deviation to any reporting strategy σ_i *such that* $\sigma_{it}(\theta_i^t, \theta_i^{t-1}, \chi_i^{t-1}(\theta^{t-1})) \equiv \theta_{it}$ *for all* $t < s$. *That is, for any such strategy, we have* $V_{is}^{\langle \chi \circ \sigma_i, \Psi \circ \sigma_i \rangle, \Gamma}(\theta^{s-1}, \theta_{is}) \leq V_{is}^{\langle \chi, \Psi \rangle, \Gamma}(\theta^{s-1}, \theta_{is})$. *The choice rule* $\langle \chi, \Psi \rangle$ *is on-path Bayes Incentive Compatible (on-path BIC) if the above property holds for all agents* $i = 1, ..., N$, *all periods* $s \geq 0$.

The concept of on-path BIC extends the concept of BNIC defined in Section 3 from period 0 to all periods. Thus, Theorem 1 can be extended to an arbitrary period to state necessary first-order conditions for on-path BIC. For this purpose, we extend Conditions F-BE$_0$ and F-BIR$_0$ to apply to all periods as follows:[38]

Condition 9 (F-BE) *Process bounded expectations: For any* $i = 1, ..., N$, $t \geq 0$, $(\theta^{t-1}, \theta_{it}) \in \Theta^{t-1} \times \Theta_{it}$, $\chi \in \mathcal{X}$, $\Gamma \in \mathbf{\Gamma}$,

$$\mathbb{E}^{\lambda_i[\chi,\Gamma]|\theta^{t-1},\theta_{it}} \left[||\tilde{\theta}_i|| \right] < \infty.$$

Condition 10 (F-BIR) *The kernels F admit a state representation* $\langle \mathcal{E}, G, z \rangle$ *(see Definition 2) with the following property: For any* $i = 1, ..., N$, $s \geq 0$, *define the functions* $\left\langle Z_{i,(s),t} : \Theta_i^s \times X_i^{t-1} \times \mathcal{E}_i^t \to \Theta_{it} \right\rangle_{t=0}^{\infty}$ *inductively as follows*

$$Z_{i,(s),t}(\theta_i^s, x_i^{t-1}, \varepsilon_i^t) = \begin{cases} \theta_{it} & \text{for } t \leq s, \\ z_{it}(Z_{i,(s)}^{t-1}(\theta_i^s, x_i^{t-2}, \varepsilon_i^{t-1}), x_i^{t-1}, \varepsilon_{it})) & \text{for } t > s \end{cases} \tag{17}$$

with $Z_{i,(s)}^{t-1}(\theta_i^s, x_i^{t-2}, \varepsilon_i^{t-1}) \equiv \left\langle Z_{i,(s),\tau}(\theta_i^s, x_i^{\tau-1}, \varepsilon_i^\tau) \right\rangle_{\tau=0}^{t-1}$. *For each* $i = 1, ..., N$, $s \geq 0$, $t \geq s$, $\theta_i^s \in \Theta_i^s$, $x_i \in X_i$, $\varepsilon_i \in \mathcal{E}_i$ *the derivative* $\partial Z_{i,(s),t}(\theta_i^s, x_i^{t-1}, \varepsilon_i^t)/\partial \theta_{is}$ *exists and is bounded in absolute value by* $C_{i,(s),t-s}(\varepsilon_i)$, *where* $\mathbb{E}[||C_{i,(s)}(\tilde{\varepsilon}_i)||] \leq B_i < \infty$.

For any $i = 1, ..., N$, $s \geq 0$, and $t \geq s$, then let

$$I_{i,(s),t}(\theta_i^t, x_i^{t-1}) \equiv \mathbb{E} \left[\frac{\partial Z_{i,(s),t}(\theta_i^s, x_i^{t-1}, \varepsilon_i^t)}{\partial \theta_{is}} \middle| Z_i^t(\theta_i^s, x_i^{t-1}, \tilde{\varepsilon}_i^t) = \theta_i^t \right] \tag{18}$$

denote the impulse response of agent i's period-t type to the period-s type. Then the definition of ICFOC and its necessity extend as follows (proof follows from the same arguments as for $s = 0$).

[38] These conditions could be in turn ensured by verifying the simple sufficient conditions on the kernels F given in Lemma 2 above.

Definition 6 (ICFOC$_s$) *Fix $i \in \{1, ..., N\}$ and $s \geq 0$. The choice rule $\langle \chi, \Psi \rangle$ with accompanying belief system Γ satisfies ICFOC$_{i,s}$ if, for any $\theta^{s-1} \in \Theta^{s-1}$, agent i's equilibrium time-s expected payoffs $V_{is}^{\langle \chi, \Psi \rangle, \Gamma}(\theta^{s-1}, \cdot)$ is Lipschitz continuous in θ_{is}, with the derivative given a.e. by*

$$\frac{\partial V_{is}^{\langle \chi, \Psi \rangle, \Gamma}(\theta^{s-1}, \theta_{is})}{\partial \theta_{is}} = \mathbb{E}^{\lambda_i[\chi, \Gamma]|\theta^{s-1}, \theta_{is}} \left[\sum_{t=s}^{\infty} \frac{\partial U_i(\tilde{\theta}, \chi(\tilde{\theta}))}{\partial \theta_{it}} I_{i,(s),t}(\tilde{\theta}_i^t, \chi_i^{t-1}(\tilde{\theta})) \right] \tag{19}$$

For any $s \geq 0$, the choice rule $\langle \chi, \Psi \rangle$ with accompanying belief system Γ satisfies ICFOC$_s$ if it satisfies ICFOC$_{i,s}$ for all agents $i = 1, ..., N$.

Proposition 4 (Necessity of ICFOC$_s$) *Assume Conditions U-TD, U-ELC, F-BE, F-BIR hold. Then any on-path BIC choice rule $\langle \chi, \Psi \rangle$ satisfies ICFOC$_s$ in each period $s \geq 0$.*

5.1 Payment Construction

We now show how, for any given allocation rule χ and accompanying beliefs Γ, one can construct transfers that satisfy ICFOC$_s$ in all periods $s \geq 0$. In addition, with infinite horizon, the transfers cannot be postponed until the mechanism is over, and we want to be able to spread the transfers over time, in the following sense:

Definition 7 (Spreadable payments) *Given the choice rule $\langle \chi, \Psi \rangle$ with accompanying beliefs Γ, the payment rule $\Psi : \Theta \rightarrow \mathbb{R}^N$ is spread over time with flow payments $\left\langle \psi_t : \Theta^t \rightarrow \mathbb{R}^N \right\rangle_{t=0}^{\infty}$ if the following are true for any $i = 1, ..., N$, $s \geq 0$, $\theta^s \in \Theta^s$: (i) for $\lambda_i[\chi, \Gamma]|\theta^{s-1}, \theta_{is}$-almost all θ, $\Psi_i(\theta) = \sum_{t=0}^{\infty} \delta^t \psi_{it}(\theta^t)$; and (ii) $\mathbb{E}^{\lambda_i[\chi, \Gamma]|\theta^{s-1}, \theta_{is}} \left[||\psi_i(\tilde{\theta})|| \right] < \infty$*

Part (ii) of the definition means that the flow payments to agent i have a finite Expected Present Value (EPV), in the sense of double Lebesgue-Stieltjes integration, at any of his information sets. This will allow us to appeal to Fubini's theorem to interchange expectation and infinite summation in calculating expected present values, and also ensures that the series in (i) converges absolutely with probability one.

To be able to construct transfers that can be spread over time, we make the following assumption on the utility function:[39]

[39] For the result in Theorem 2 below the following weaker condition actually suffices:

Condition 11 (U-SPRST) *Utility Spreadable for some type: For any $i = 1, ..., N$, there exists $\theta_i^{spr} \in \Theta_{i\delta}$ and a sequence of functions*

$$\left\langle v_{it} : \Theta_{-i}^t \times X^t \rightarrow \mathbb{R} \right\rangle_{t=0}^{\infty}$$

with $\left| v_{it}(\theta_{-i}^t, x^t) \right| \leq K_{it}$ with $||K_i|| < \infty$ such that, for any $(\theta_{-i}, x) \in \Theta_{-i} \times X$, $U_i(\theta_i^{spr}, \theta_{-i}, x) = \sum_{t=0}^{\infty} \delta^t v_{it}(\theta_{-i}^t, x^t)$.

Condition 12 (U-SPR) *Utility Spreadable: For any $i = 1, ..., N$, there exists a sequence of functions*

$$\left\langle u_{it} : \Theta^t \times X^t \to \mathbb{R} \right\rangle_{t=0}^{\infty}$$

such that $U_i(\theta, x) = \sum_{t=0}^{\infty} \delta^t u_{it}\left(\theta^t, x^t\right)$ with $\left|u_{it}\left(\theta^t, x^t\right)\right| \leq L_i|\theta_{it}| + M_{it}$ for all $(\theta, x) \in \Theta \times X$, with $L_i, ||M_i|| < \infty$.

Now, to construct transfers satisfying ICFOC$_s$, for each $i = 1, ..., N$, $s \geq 0$, let

$$\begin{aligned} D_{is}^{\chi,\Gamma}\left(\theta^{s-1}, \theta_{is}\right) &= \mathbb{E}^{\lambda_i[\chi,\Gamma]|\theta^{s-1},\theta_{is}}\left[\sum_{t=s}^{\infty} I_{i,(s),t}(\tilde{\theta}_i^t, \chi_i^{t-1}(\tilde{\theta}))\frac{\partial U_i(\tilde{\theta}, \chi(\tilde{\theta}))}{\partial \theta_{it}}\right] \qquad (20) \\ \text{and } Q_{is}^{\chi,\Gamma}\left(\theta^{s-1}, \theta_{is}\right) &= \int_{\hat{\theta}_{is}}^{\theta_{is}} D_{is}^{\chi,\Gamma}\left(\theta^{s-1}, q\right) dq \end{aligned}$$

for some arbitrarily fixed type sequence $\hat{\theta}_i \in \Theta_{i\delta}$. Then for all $i = 1, ..., N$, $\theta \in \Theta$, let

$$\begin{aligned} \Psi_i\left(\theta\right) &= \sum_{t=0}^{\infty} \delta^t \psi_{it}\left(\theta^t\right), \text{ where for all } t \geq 0, \qquad (21) \\ \psi_{it}\left(\theta^t\right) &= \delta^{-t} Q_{it}^{\chi,\Gamma}(\theta^{t-1}, \theta_{it}) - \delta^{-t}\mathbb{E}^{\lambda_i[\chi,\Gamma]|\theta^{t-1},\theta_{it}}\left[Q_{i,t+1}^{\chi,\Gamma}(\tilde{\theta}^t, \tilde{\theta}_{i,t+1})\right] - \mathbb{E}^{\lambda_i[\chi,\Gamma]|\theta^{t-1},\theta_{it}}\left[u_{it}(\tilde{\theta}^t, \chi^t(\tilde{\theta}^t))\right]. \end{aligned}$$

We then have the following result (proof in the Appendix):

Theorem 2 (Payment construction) *Assume Conditions U-TD, U-ELC, F-BE, F-BIR , and U-SPR hold. Take any allocation rule χ and accompanying belief system Γ. Then (i) the transfers Ψ in (21) are spread over time with the flow payments ψ, and (ii) the choice rule $\langle \chi, \Psi \rangle$ satisfies ICFOC$_s$ in all periods $s \geq 0$.*

Remark 1 Note that the flow payments $\psi_{it}\left(\theta^t\right)$ in (21) are measurable with respect to $\left(\theta_i^t, \chi_i^t\left(\theta^t\right)\right)$. This means that they do not reveal to agent i any information in addition to that contained in the allocations x_{it}. Hence they can be disclosed to the agent without affecting his beliefs (and hence his incentives). \\

5.2 Payment Equivalence

Having established existence of a payment rule Ψ that ensures ICFOC$_s$ in all periods $s \geq 0$, we now examine to what extent ICFOC pins down these payments for a given allocation rule χ.

First note that ICFOC$_0$ nails down, up to a constant, the expected payments $\mathbb{E}^{\lambda_i[\chi]|\theta_{i0}}\left[\Psi_i(\tilde{\theta})\right]$ for each starting type θ_{i0}, in any mechanism implementing a given allocation rule χ as a Bayes-Nash

Note that, with finite horizon, this condition boils down to requiring that utility be bounded at one particular type. When the horizon is infinite, the stronger condition U-SPR is needed to validate the one-step-deviation principle of Proposition 5.

equilibrium. This extends the celebrated Revenue Equivalence result for static mechanism design to the dynamic setting.

Below, we derive an even stronger version of Revenue Equivalence for many dynamic environments, by using ICFOC in all periods and not just in period 0. We start by considering a single-agent environment and show that the payments Ψ that implement a given allocation rule χ are pinned down by ICFOC *with probability one*, and not just in expectation. We then explain in what sense this finding extends to a setting with multiple agents, and also remark that it extends to BNE implementation as long as the kernels have connected supports.

First consider the single-agent case, where beliefs are vacuous, and omit the agent index for the time being. Using ICFOC and the Law of Iterated Expectations, for any two on-path IC choice rules $\langle\chi,\Psi\rangle$ and $\langle\chi,\bar{\Psi}\rangle$ implementing the same allocation rule χ, any $s\geq 0$, $\theta^s\in\Theta^s$, we can write

$$\begin{aligned} V_s^{\langle\chi,\Psi\rangle}(\theta^s)-V_{s-1}^{\langle\chi,\Psi\rangle}\left(\theta^{s-1}\right) &= V_s^{\langle\chi,\Psi\rangle}(\theta^s)-\mathbb{E}^{F_s\left(\theta^{s-1},\chi^{s-1}\left(\theta^{s-1}\right)\right)}\left[V_s^{\langle\chi,\Psi\rangle}(\theta^{s-1},\tilde{\theta}_s)\right] \\ &= \mathbb{E}^{F_s\left(\theta^{s-1},\chi^{s-1}\left(\theta^{s-1}\right)\right)}\left[\int_{\tilde{\theta}_s}^{\theta_s} D_s^{\chi}(\theta^{s-1},q)dq\right] \\ &= V_s^{\langle\chi,\bar{\Psi}\rangle}(\theta^s)-V_{s-1}^{\langle\chi,\bar{\Psi}\rangle}(\theta^{s-1}). \end{aligned} \tag{22}$$

Substituting the definitions (16) of expected payoffs and rearranging terms yields

$$\mathbb{E}^{\lambda[x]|\theta^s}\left[\Psi(\tilde{\theta})\right]-\mathbb{E}^{\lambda[x]|\theta^s}\left[\bar{\Psi}(\tilde{\theta})\right]=\mathbb{E}^{\lambda[x]|\theta^{s-1}}\left[\Psi(\tilde{\theta})\right]-\mathbb{E}^{\lambda[x]|\theta^{s-1}}\left[\bar{\Psi}(\tilde{\theta})\right].$$

Thus, we see by induction that for each $T\geq 1$,

$$\mathbb{E}^{\lambda[x]|\theta^T}\left[\Psi(\tilde{\theta})\right]-\mathbb{E}^{\lambda[x]|\theta^T}\left[\bar{\Psi}(\tilde{\theta})\right]=\mathbb{E}^{\lambda[x]}\left[\Psi(\tilde{\theta})\right]-\mathbb{E}^{\lambda[x]}\left[\bar{\Psi}(\tilde{\theta})\right]\equiv K. \tag{23}$$

Now, when the payments can be spread over time, one can also verify that as $T\to\infty$, for $\lambda[\chi]$–almost all θ, $\mathbb{E}^{\lambda[x]|\theta^T}\left[\Psi(\tilde{\theta})\right]\to\Psi(\theta)$.[40] Thus, when both payment rules $\bar{\Psi}$ and Ψ can be spread over time, we have that

$$\Psi(\theta)-\bar{\Psi}(\theta)=K \text{ for } \lambda[\chi]\text{–almost all } \theta$$

This conclusion proves important for solving problems in which the principal cares not just about the expectation of the payments (as in our Relaxed Program above), but also about how the payments vary with the state θ. For example, this includes models where $\psi_t(\theta)$ is interpreted as the "utility payment" to the agent in period t whose monetary cost to the principal is some convex function $\gamma(\psi_t(\theta))$, as in models where the agent is risk-averse with respect to the payments and has preferences for consumption smoothing. In such models, knowing the NPV $\Psi(\theta)$ of "utility payments" required to implement a given allocation rule helps computing the cost-minimizing distribution of the payments

[40] The proof for this result is in Lemma 5 in the Appendix.

over time and hence the principal's expected cost of sustaining a given allocation rule. This is in turn instrumental to the selection of the profit-maximizing rule (see Garrett and Pavan (2011b) for an application of this approach to managerial compensation).

The arguments above extend to settings with multiple agents, under the following condition.

Definition 8 (No leakage) *The allocation rule χ leaks no information to agent i if for each $t \geq 0$, $\theta_i^{t-1}, \hat{\theta}_i^{t-1} \in \Theta_i^{t-1}$, the distribution $F_{it}(\theta_i^{t-1}, \chi_i^{t-1}(\hat{\theta}_i^{t-1}, \theta_{-i}^{t-1}))$ does not depend on θ_{-i}^{t-1} (and so can be written as $\hat{F}_{it}(\theta_i^{t-1}, \hat{\theta}_i^{t-1})$).*

This condition means that the observation of θ_{it} never gives agent i any information about the other agents' types. Clearly, any allocation rule satisfies this property when the agent's process evolves autonomously from the allocations, i.e., when Condition F-AUT holds, or (trivially) in a single-agent setting.

Now observe that when on-path IC holds for agent i, then it must also hold in the "blind" setting in which agent i is not shown x_i. (Indeed, the set of agent i's contingent reporting plans in the blind settings is a subset of those in the original setting, but still includes truthtelling.) Furthermore, if the allocation rule χ leaks no information to agent i, then we can interpret the "blind" setting as a single-agent setting in which agent i's allocation in period t is simply his report $\hat{\theta}_{it}$, and his utility is $\hat{U}_i(\theta_i, \hat{\theta}_i) = \mathbb{E}^{\lambda_i[\chi]|\hat{\theta}_i}\left[U_i(\theta_i, \tilde{\theta}_{-i}, \chi(\hat{\theta}_i, \tilde{\theta}_{-i}))\right]$, where $\lambda_i[\chi]|\hat{\theta}_i$ denotes the process over the other agents' types given that agent i's reports are fixed at $\hat{\theta}_i$. (Intuitively, the other agents' types can be viewed as being realized only after agent i has finished reporting, and $\hat{U}_i$ is the expectation taken over such realizations.) Applying to this setting the result established above for the single-agent case, we see that agent i's expected payment $\mathbb{E}^{\lambda_i[\chi]|\theta_i}\left[\Psi_i\left(\theta_i, \tilde{\theta}_{-i}\right)\right]$ is pinned down, up to a constant, by the allocation rule χ with probability 1.

Finally, we observe that the same result can be derived from BNIC rather than on-path BIC, as long as the kernels have a connected support. The complication arising in this case is that BNIC does not ensure that the equilibrium continuation payoff after period zero satisfy the envelope formula (19). Indeed, the continuation payoff need not even be continuous in θ_{it}, since the choice rule may be modified arbitrarily at zero-probability events while preserving BNIC. Nonetheless, the argument behind Proposition 4 still applies to the *value function* (i.e., the supremum expected payoff achieved by possible reporting plans starting at a given history). Namely, it shows that the value function must be Lipschitz continuous and its derivative must satisfy ICFOC at each on-path history at which truthful reporting is optimal. Furthermore, the principle of dynamic programming ensures that in a BNIC choice rule, truthtelling is optimal for probability-1 histories, i.e., almost everywhere on the support of the process. Hence, when the kernels' supports are connected, the value function satisfies condition (22). Since the value function coincides with the equilibrium expected payoff at probability-1 histories, all the preceding arguments extend as well, and we obtain the following result (proof follows from the arguments above):

Theorem 3 (Payment equivalence) *Assume that conditions U-TD, U-ELC, F-BE, and F-BIR hold. Take any two BNIC choice rules* $\langle\chi,\Psi\rangle$ *and* $\langle\chi,\bar{\Psi}\rangle$ *and suppose that the allocation rule* χ *leaks no information to agent* i*, and that the payments* Ψ *and* $\bar{\Psi}$ *can be spread over time. Then there exists a constant* K_i *such that for* $\lambda[\chi]$*-almost all* θ_i*,*

$$\mathbb{E}^{\lambda_i[\chi]|\theta_i}\left[\Psi_i\left(\theta_i,\tilde{\theta}_{-i}\right)\right]=\mathbb{E}^{\lambda_i[\chi]|\theta_i}\left[\bar{\Psi}_i\left(\theta_i,\tilde{\theta}_{-i}\right)\right]+K_i.$$

5.3 Sufficient conditions

We now provide sufficient conditions for a given allocation rule χ to be implementable in a PBE of a direct revelation mechanism. We will then apply these conditions to the solution to the relaxed program, and also verify that the solution satisfies all the participation constraints.

From now on, we focus on sustaining strong truthtelling strategies, where an agent reports truthfully at any private history. This focus is clearly restrictive, since in general an agent who lied in the past may find it optimal to continue lying. Characterizing such sequentially optimal contingent strategies is quite difficult, which explains why we focus on settings in which strong truthtelling can indeed be sustained.[41] In particular, these settings include Markov environments, defined as follows:

Definition 9 (Markov) *The environment is Markov in period* t *if for each* $i=1,...,N$*: (a) for any* $\tau>t$*, any* $x_i^{\tau-1}\in X_i^{\tau-1}$*, the kernel* $F_{i\tau}\left(\theta_i^{\tau-1},x_i^{\tau-1}\right)$ *does not depend on* θ_i^{t-1} *and (b) there exist functions* $\bar{U}_{it}:\Theta^{t-1}\times X^{t-1}\to\mathbb{R}$ *and* $\hat{U}_{it}:\Theta_i^{\geq t}\times\Theta_{-i}\times X\to\mathbb{R}$ *such that, for any* $(\theta,x)\times\Theta\times X$*,* $U_i(\theta,x)=\bar{U}_{it}(\theta^{t-1},x^{t-1})+\hat{U}_{it}(\theta_i^{\geq t},x,\theta_{-i})$*.*

This definition ensures that, after observing θ_i^t, the kernels governing agent i's future types as well as his vNM preferences over future lotteries depend on his type history θ_i^t only through θ_{it} (but they can depend on past decisions x^{t-1}, as well as, in the case of interdependent values, on other agents' past types, θ_{-i}^{t-1}).[42] This guarantees that the agent's reporting incentives in period t depend only on his previous reports, and not on whether these reports have been truthful.

[41] Townsend (1982) is able to restrict attention to strong truthtelling without loss of generality by having each agent report his complete "state," rather than just the latest signal. That is, with each report the agent is given a chance to "correct" his past lies. Since this approach would entail multidimensional reports in each period, it is not clear that it would help in verifying incentive compatibility.

[42] All subsequent results that do not hinge on payment construction such as Theorem 4 below extend to the following more general definition of Markov payoffs

$$U_i(\theta,x)=\bar{U}_{it}(\theta^{t-1},x^{t-1})+B_{it}(\theta^{t-1},x^{t-1})\hat{U}_{it}(\theta_i^{\geq t},\theta_{-i},x)$$

with $B_{it}(\cdot)>0$, which includes multiplicatively separable and additively separable payoffs as special cases. The stronger condition in the main text is however needed in settings with payments to guarantee that the net payoffs satisfy all the ICFOC conditions.

To state our main sufficiency result, for any $\hat{\theta}_{it} \in \Theta_{it}$, let $\chi \circ \hat{\theta}_{it}$ denote the allocation rule obtained from χ by replacing θ_{it} with $\hat{\theta}_{it}$, i.e., such that, for any θ, $(\chi \circ \hat{\theta}_{it})(\theta) = \chi\left(\hat{\theta}_{it}, \theta_{i,-t}, \theta_{-i}\right)$. We then have the following result.

Theorem 4 (Sufficiency) *Suppose that Conditions U-TD, U-ELC, F-BE, and F-BIR hold, and consider a choice rule* $\langle \chi, \Psi \rangle$ *with accompanying system of beliefs* Γ. *Suppose that (i) the environment is Markov in period* t, *(ii) the choice rule satisfies ICFOC*$_t$, *(iii) for any* $i = 1, ..., N$, $\theta^{t-1} \in \Theta^{t-1}$, $\hat{\theta}_{it} \in \Theta_{it}$, *a.e.* $\theta_{it} \in \Theta_{it}$,

$$\left[D_{it}^{\chi,\Gamma}(\theta^{t-1}, \theta_{it}) - D_{it}^{\chi\circ\hat{\theta}_{it},\Gamma}(\theta^{t-1}, \theta_{it})\right] \cdot \left(\theta_{it} - \hat{\theta}_{it}\right) \geq 0, \tag{24}$$

where the functions D *are as defined in (20), and (iv) either strong truthtelling is optimal at all period-*$(t+1)$ *histories (including non-truthful ones), or the choice rule satisfies ICFOC*$_{t+1}$ *and the environment is Markov in period* $t+1$. *Then, for all agents, a one-step deviation from strong truthtelling is not profitable at any period-t history.*

Proof. We start by establishing the following lemma.[43]

Lemma 3 *Consider a function* $\Phi : (\underline{\theta}, \overline{\theta})^2 \to \mathbb{R}$. *Suppose that (a) for all* $\hat{\theta} \in (\underline{\theta}, \overline{\theta})$, *the function* $\Phi(\cdot, \hat{\theta})$ *is Lipschitz continuous in* θ, *(b) the function* $\bar{\Phi}(\theta) \equiv \Phi(\theta, \theta)$ *is Lipschitz continuous in* θ, *and (c) for any* $\hat{\theta} \in (\underline{\theta}, \overline{\theta})$, *a.e.* $\theta \in (\underline{\theta}, \overline{\theta})$, $\left(\bar{\Phi}'(\theta) - \partial\Phi(\theta, \hat{\theta})/\partial\theta\right) \cdot \left(\theta - \hat{\theta}\right) \geq 0$. *Then* $\bar{\Phi}(\theta) \geq \Phi(\theta, \hat{\theta})$ *for all* $(\theta, \hat{\theta}) \in (\underline{\theta}, \overline{\theta})^2$.

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

$$g(\theta, \hat{\theta}) = \int_{\hat{\theta}}^{\theta} \frac{\partial g(q, \hat{\theta})}{\partial \theta} dq = \int_{\hat{\theta}}^{\theta} \left[\bar{\Phi}'(q) - \frac{\partial \Phi(q, \hat{\theta})}{\partial \theta}\right] dq.$$

By (c), the integrand is nonnegative for a.e. $q \geq \hat{\theta}$ and nonpositive for a.e. $q \leq \hat{\theta}$. Therefore, $g(\theta, \hat{\theta}) \geq 0$ for both $\theta \geq \hat{\theta}$ and $\theta < \hat{\theta}$. For any $\theta \in (\underline{\theta}, \overline{\theta})$, the function $\Phi(\theta, \cdot)$ is thus maximized at $\hat{\theta} = \theta$. ■

Now, to apply the lemma, fix $\theta^{t-1} \in \Theta^{t-1}$ and let $\Phi(\theta_{it}, \hat{\theta}_{it})$ denote agent i's expected payoff in period t when, after having reported θ_i^{t-1} truthfully in the past and having observed $x_i^{t-1} = \chi_i^{t-1}(\theta_i^{t-1}, \theta_{-i}^{t-1})$, he observes the period-t signal θ_{it}, reports $\hat{\theta}_{it}$, and then reverts to strong truthtelling from period-$(t+1)$ onwards (assuming that all other agents follow a truthful strategy in each period). Observe that, since the environment is Markov in period t, it suffices to show the suboptimality of deviations from strong truthtelling at truthful histories, i.e., for $\hat{\theta}_i^{t-1} = \theta_i^{t-1}$. The remaining part of the proof establishes that conditions (i)-(iv) in the theorem imply assumptions (a)-(c) in Lemma 3.

[43] A similar approach has been applied by Garcia (2005) to establish incentive-compatibility in a static problem with one-dimensional information and multidimensional decisions, but under stronger assumptions.

- Assumption (ii) of the theorem (ICFOC$_t$) ensures that assumption (b) of the Lemma holds, and that, for a.e. $\theta_{it} \in \Theta_{it}$, $\bar{\Phi}'(\theta_{it}) = D_{it}^{\chi,\Gamma}(\theta^{t-1}, \theta_{it})$ where $\bar{\Phi}(\theta_{it}) \equiv \Phi(\theta_{it}, \theta_{it})$ and $\bar{\Phi}'(\theta_{it}) = d\bar{\Phi}(\theta_{it})/d\theta_{it}$.

- Next, we show that assumption (a) of the Lemma holds and that, for all $\hat{\theta}_{it} \in \Theta_{it}$, a.e. $\theta_{it} \in \Theta_{it}$, $\partial\Phi(\theta_{it}, \hat{\theta}_{it})/\partial\theta_{it} = D_{it}^{\chi\circ\hat{\theta}_{it},\Gamma}(\theta^{t-1}, \theta_{it})$. Note that this means that, for agent i, the choice rule $\left\langle \chi \circ \hat{\theta}_{it}, \Psi \circ \hat{\theta}_{it} \right\rangle$ that is obtained from $\langle \chi, \Psi \rangle$ by ignoring agent i's report θ_{it} and replacing it with $\hat{\theta}_{it}$ satisfies ICFOC$_{i,t}$. The first case of Condition (iv) of the Theorem implies that the choice rule $\left\langle \chi \circ \hat{\theta}_{it}, \Psi \circ \hat{\theta}_{it} \right\rangle$ is on-path BIC for agent i in period t and therefore, by Proposition 4, it must satisfy ICFOC$_{i,t}$. The second case of Condition (iv) ensures that the new choice rule satisfies ICFOC$_{i,t+1}$ (indeed, since the environment is Markov in period $t+1$, the satisfaction of ICFOC$_{t+1}$ does not depend on whether or not $\theta_{it} = \hat{\theta}_{it}$). That the new choice rule satisfies ICFOC$_{i,t}$ then follows from the following lemma (proof in the Appendix):

Lemma 4 *Suppose that the environment is Markov in period $t+1$ and that the choice rule $\langle \chi, \Psi \rangle$ ignores θ_{it}, and satisfies ICFOC$_{i,t+1}$ with accompanying beliefs Γ. Then the choice rule satisfies ICFOC$_{i,t}$.*

- Given the previous observations, assumption (c) of the Lemma then follows from assumption (iii) of the theorem.

■

Recall that ICFOC is necessary for PBE implementability (Proposition 4), and that in quasilinear environments we can always construct transfers to satisfy ICFOC (Theorem 2). Thus, the key contribution of Theorem 4 is in identifying the single-crossing condition (24) that ensures implementability in quasilinear environments.[44] This condition can be viewed as a generalization of monotonicity of the allocation rule in the static setting (as we discuss in the next subsection); however, in contrast to the static setting, it is not necessary for implementability but only sufficient. The value of the theorem is in providing a route to sufficiency both for Markov environments as well as for environments where truthtelling in future periods can be established either by backward induction (as in finite-horizon models) or by assuming certain separability conditions, as in the applications of subsection 6.3 below.

Given the result in the previous theorem, the next step is to establish the suboptimality of arbitrary, and not just one-step, deviations from strong truthtelling. For this purpose, we establish the following version of the one-stage deviation principle (proof in the Appendix):[45]

[44] Theorem 4 also applies to nonquasilinear environments simply by letting $\Psi(\theta) \equiv 0$, but in this case ICFOC cannot be guaranteed through payments. To verify that a given allocation rule χ is PBE-implementable, one must thus verify that it satisfies both ICFOC and the single-crossing condition (24).

[45] The usual version of the Principle, found, e.g., in Fudenberg and Tirole (1991, p.110), is not applicable in our

Proposition 5 (OSDP) *Assume Condition U-SPR and the assumption of part (b) of Lemma 2 hold for each agent* $i = 1, ..., N$*. Let* $\langle \chi, \Psi \rangle$ *be a choice rule where the transfers* Ψ *can be spread over time with flows* ψ *satisfying for all* $i = 1, ..., N$*, (i)* $\psi_{it}\left(\theta^t\right) \leq K_{it}$ *for all* $t \geq 0$*, with* $\|K_i\| < \infty$*, and (ii)* $\left|\psi_{it}\left(\theta^t\right)\right| \leq L_i^{\psi} \left|\theta_{it}\right| + M_{it}^{\psi}$*, for all* $t \geq 0$*, with* $L_i^{\psi}, ||M_i^{\psi}|| < \infty$*. Then if one-stage deviations from strong truthtelling are not profitable at any information set, arbitrary deviations from strong truthtelling are not profitable at any information set.*

Intuitively, bound (i) on the payments ensures that for any possible deviation strategy, the net present value of the payments received by the agent from period T onward is either small or negative when T is large. Bound (ii), together with the assumption of part (b) of Lemma 2, ensure that if the agent reverts to strong truthtelling in period T, the expected payments received after the reversion are small in absolute value. Hence, the reversion will not cost the agent much in terms of expected payments. Under Condition U-SPR, the same is true for the agent's expected non-monetary utility. Hence, in considering potential profitable deviations from strong truthtelling, it suffices to check only those deviations that revert to strong truthtelling at some finite period T. The usual backward-induction argument then ensures that these deviations are unprofitable when no one-step profitable deviation exists.

Theorem 4 and Proposition 5 can be applied to some non-Markov environments. For example, in Section 6.3, we use Theorem 4 to establish the optimality of truthtelling in period 0 (in which Markovness is vacuously satisfied), while the optimality of strong truthtelling in subsequent periods is established in a different way. Also, in that setting we can verify that the transfers satisfy the assumptions of Proposition 5. Nevertheless, these results can be applied more systematically to environments that are Markov in all periods, as we now demonstrate.

5.4 Sufficiency for Markov Environments

Note that in Markov environments (that is, in environments that are Markov in all periods), each agent i's utility function must be additively separable in θ_{it}. Condition U-SPR is then strengthened as follows:

Condition 13 (U-MSPR) *Utility Markov spreadable: For each* $i = 1, ..., N$*, there exists a sequence of functions* $\left\langle u_{it} : \Theta_{it} \times \Theta_{-i}^t \times X^t \to \mathbb{R} \right\rangle_{t=0}^{\infty}$ *such that* $U_i(\theta, x) = \sum_{t=0}^{\infty} \delta^t u_{it}\left(\theta_{it}, \theta_{-i}^t, x^t\right)$ *with* $\left|u_{it}\left(\theta_{it}, \theta_{-i}^t, x^t\right)\right| \leq L_i \left|\theta_{it}\right| + M_{it}$ *for all* $(\theta, x) \in \Theta \times X$*, where* $L_i, \|M_i\| < \infty$*.*

As for the kernels, we assume the following condition, which is the Markov analog of the assumption of part (b) of Lemma 2, and in particular implies Condition F-BE.

setting since the payoffs are not continuous at infinity when flow payments are not bounded – and the flow payments (21) need not be bounded when Θ is unbounded.

Condition 14 (F-MB) *Process Markov bounded: For all $i = 1, ..., N$, $t \geq 0$, $\left(\theta_{it}, x_i^t\right) \in \Theta_{it} \times X_i^t$, $\mathbb{E}^{F_{i,t+1}\left(\theta_{it}, x_i^t\right)}\left[\left|\tilde{\theta}_{t+1}\right|\right] \leq \phi_i \left|\theta_{it}\right| + E_{it}$, with $\delta\phi_i < 1$ and $\|E_i\| < \infty$.*

Under these two conditions the one-stage deviation principle can be ensured (proof in the Appendix):

Proposition 6 (OSDP-Markov) *Assume that the environment is Markov in all periods and satisfies Conditions U-TD, U-ELC, U-MSPR, F-BIR, and F-MB. Take an allocation rule χ with accompanying belief system Γ and let the payments Ψ be given by (21). Then if one-stage deviations from strong truthtelling are not profitable at any information set, arbitrary deviations from strong truthtelling are not profitable at any information set.*

Combining Theorem 4 with Proposition 6 immediately yields the following important result.

Corollary 1 (Markov-implementability) *Assume that the environment is Markov in all periods and satisfies Conditions U-TD, U-ELC, F-BIR, U-MSPR, and F-MB. Then any allocation rule χ that, together with the accompanying beliefs Γ satisfies the single-crossing conditions (24) in each period t can be implemented in a strong truthtelling PBE with payments given by (21).*

Using this Corollary, we can ensure that the allocation rule solving the relaxed program also solves the full program. All that is left for this is to ensure that all participation constraints are satisfied. This can be done as follows:

Theorem 5 (Optimal mechanisms) *Assume that the environment is Markov in all periods and satisfies Conditions U-TD, U-ELC, U-MSPR, F-BIR, and F-MB. In addition, assume that $U_i(\theta, x)$ is nondecreasing in θ_i for each $i = 1, ..., N$ and that condition F-FOSD holds. Let χ be an allocation rule that maximizes the expected dynamic virtual surplus (14), and suppose that, with an accompanying belief system Γ, it satisfies the single-crossing condition (24) in all periods. Then*

(a) There exists a payment rule Ψ that can be spread over time (with flow payments that can be disclosed to the agents) such that (i) strong truthtelling arises in a PBE of the direct mechanism for $\langle\chi, \Psi\rangle$, (ii) the lowest types' participation constraints bind, and (iii) all types' participation constraints are satisfied;

(b) The above choice rule maximizes the principal's expected payoff among all BNIC-IR choice rules;

(c) In any BNIC-IR choice rule that is optimal for the principal, the allocation rule must maximize the expected dynamic virtual surplus (14).

(d) The principal's expected payoff cannot be increased using randomized mechanisms.

Proof. For part (a), consider the payments $\hat{\Psi}_i(\theta) = \Psi_i(\theta) - V_i^{\langle\chi,\Psi\rangle}(\underline{\theta}_{i0})$ for each $i = 1, ..., N$, where Ψ is the payment rule given by (21). Since $\hat{\Psi}$ only differs from Ψ by a constant, property (i) follows from Corollary 1. Property (ii) holds by construction. For (iii), recall that the equilibrium expected payoffs are $V_{i0}^{\langle\chi,\hat{\Psi}\rangle}(\theta_{i0}) = \int_{\underline{\theta}_{i0}}^{\theta_{i0}} D_{i0}^{\chi,\Gamma}(q)\,dq$. Under F-FOSD and $U_i(\theta, x)$ nondecreasing in θ_i, $D_{i0}^{\chi,\Gamma}(q) \geq 0$ all q, which implies that $V_{i0}^{\langle\chi,\hat{\Psi}\rangle}(\theta_{i0}) \geq 0$ all θ_{i0}.

Next, consider parts (b) and (c). These results are implied by Proposition 2. Finally, for part (d) note that a randomized mechanism is equivalent to a mechanism that conditions on the random types of some fictitious agent. Since the expected virtual surplus in this expanded setting is independent of the signals of the fictitious agent and still takes the form (14), it is still maximized by the non-randomized allocation rule χ. Thus, applying part (a) of the Theorem to the expanded setting implies that the deterministic choice rule $\langle\chi, \Psi\rangle$ still maximizes the principal's expected payoff. (A similar point was made in a static mechanism design setting by Strausz (2006)). ■

5.5 Verification of the Single-Crossing Condition

We found that the crucial condition ensuring the PBE implementability of an allocation rule is the single-crossing condition (24). While this condition need not always be easy to check, in many cases of interest it can be checked by verifying that the allocation rule satisfies one of the following simpler (but less general) conditions, which can be interpreted as different forms of monotonicity of the allocation rule (proof is immediate and hence omitted):

Proposition 7 (Monotonicities) *The single-crossing condition (24) of Theorem 4 is ensured by any of the following conditions (listed in decreasing order of generality):*

- *For each $i = 1, ..., N$, $t \geq 0$, $(\theta^{t-1}, \theta_{it}) \in \Theta^{t-1} \times \Theta_{it}$, $D_{it}^{\chi\circ\hat{\theta}_{it},\Gamma}(\theta^{t-1}, \theta_{it})$ is nondecreasing in $\hat{\theta}_{it}$* ***(weak monotonicity)****;*

- *Condition F-AUT holds and, for any $i = 1, ..., N$, $t \geq 0$, $\theta \in \Theta$,*

$$\sum_{\tau=t}^{\infty} I_{i,(t),\tau}\left(\theta_i^{\tau}\right) \frac{\partial U_i\left(\theta, \chi(\hat{\theta}_{it}, \theta_i^{-t}, \theta_{-i})\right)}{\partial \theta_{i\tau}} \tag{25}$$

  is nondecreasing in $\hat{\theta}_{it}$ ***(ex-post monotonicity);***

- *In addition to conditions F-AUT, F-FOSD, for any $i = 1, ..., N$, $U_i(\theta, x)$ has increasing differences in (θ_i, x_i) and is independent of x_{-i}, and $\chi_i(\theta)$ is nondecreasing in θ_i, for all $\theta_{-i} \in \Theta_{-i}$* ***(strong monotonicity)****.*

To the best of our knowledge, all results about optimal dynamic mechanisms in the literature, with the exception of the result in Courty and Li (2000) for second-order-stochastic-dominance, establish sufficiency via strong monotonicity.[46] However, there are natural problems where the solution to the

[46] The SOSD case of Courty and Li (2000) is covered by "weak monotonicity."

relaxed program fails to be strongly monotone (see, e.g., the discussion at the end of this section, or the bandit auction of Section 6.1 that violates F-AUT and where we employ weak monotonicity instead). In this case, ex-post monotonicity, which is still quite easy to check but weaker, offers a convenient way for establishing sufficiency. Another advantage of ex-post monotonicity is that it allows us to construct payments that guarantee that each agent i has incentives to report truthfully in each period s, even if he knows the entire history θ_{-i}^{s} of the other agents' types. These payments can be constructed as in (21) but with the measures $\lambda_i [\Gamma] | \theta^{t-1}, \theta_{it}$ replaced by the measure $\lambda | \theta^t$, which is the (exogenous) process over Θ according to the kernels F starting at state θ^t. We then have the following result.

Corollary 2 (Periodic Ex-post IC) *Assume that the environment is Markov in all periods and satisfies Conditions U-TD, U-ELC, U-MSPR, F-BIR, F-MB, and F-AUT. For any allocation rule χ satisfying ex-post monotonicity, one can construct payments Ψ spreadable over time such that, in each period $s \geq 0$, each agent $i = 1, ..., N$ finds it optimal to report truthfully when he knows the other agents report truthfully, regardless of his beliefs about the other agents' type history θ_{-i}^{s}.*

In other words, the Corollary ensures that the allocation rule χ can be implemented in a periodic ex-post equilibrium, in the sense of Athey and Segal (2007). In particular, it implies that truthtelling remains a PBE of the direct mechanism even if all the reports and allocations are made public.

Remark 2 With a finite horizon T, the result in Corollary 2 can be strengthened further, ensuring incentive compatibility even if each agent could find out about the other agents' *future* types – a property that might be called "other-ex-post incentive compatibility". This is achieved by making all the payments after period T, with payments (21) using measures $\lambda_i | \theta_i^t, \theta_{-i}^T$, which put probability 1 on the other agents' realized types θ_{-i}^T and consider the evolution of agent i's type according to his kernels. Furthermore, the same arguments as in Subsection 5.2 imply that the payments that ensure other-ex-post incentive compatibility are unique, up to a constant. (With infinite horizon, where transfers must be spread over time, ensuring other-ex-post incentive compatibility seems more difficult.) \\

Remark 3 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. For example, in the special case with a single agent with linear utility (as in Example 3), a necessary condition for a differentiable allocation rule to be implementable is the symmetry of the substitution matrix, which means that for all $\tau < t$, $\partial \chi_t \left(\theta^t \right) / \partial \theta_\tau = \partial \chi_\tau \left(\theta^\tau \right) / \partial \theta_t = 0$. Hence, if we want robustness to the agent observing his own future types, the allocation rule can depend only on the current type, that is, $\chi_t \left(\theta_t \right)$. Thus, while some authors have drawn analogies between dynamic mechanism design and static multidimensional mechanism design problems (see, e.g., Courty and Li (2000) and Rochet and Stole (2003)), here we highlight an important difference: requiring the agents to learn (and report) the dimensions of their types sequentially over time allows to implement more allocation rules than in the static setting where all dimensions are learned at the same time. \\

We now turn to the even stronger property of strong monotonicity, which requires that the allocation components that agent i cares about are those that he observes and that the latter are nondecreasing in all of the agent's signals. Using standard monotone comparative statics techniques, we can provide some sufficient conditions for the allocation rule χ that solves the relaxed program (i.e., maximizes the expected dynamic virtual surplus (14)) to have this property. We can do this for two kinds of settings: (a) when all allocations are (weakly) complementary to each other, in the sense of Condition U-COMP, and (b) when each agent's allocation in each period is one-dimensional, and in addition the following condition holds:

Condition 15 (U-DSEP) *Utility decision-separable:* $X = \prod_{t=1}^{\infty} X_t$ *and, for all* $i = 0, ..., N$, *all* $(\theta, x) \in \Theta \times X$, $U_i(\theta, x) = \sum_{t=1}^{\infty} \delta^t u_{it}(\theta^t, x_t)$.

Observe that when F-AUT and U-DSEP hold, the expected virtual surplus (14) can be maximized simultaneously for all periods and states, obviating the need to solve a dynamic programming problem. That is, the solution requires that for all periods t, λ–almost all θ^t,

$$\chi_t(\theta^t) \in \arg\max_{x_t \in X_t} \left[\sum_{i=0}^{N} u_{it}(\theta^t, x_t) - \sum_{i=1}^{N} \frac{1}{\eta_{i0}(\theta_{i0})} \sum_{\tau=0}^{t} I_{i\tau}(\theta_i^\tau) \frac{\partial u_{it}(\theta^\tau, x_\tau)}{\partial \theta_{i\tau}} \right] \tag{26}$$

We then have the following result (proof in the Appendix):

Proposition 8 (Strong Monotonicity-Primitives) *Assume that conditions F-AUT and F-FOSD hold, that, for any* $i = 0, ..., N$, $t \geq 0$, X_{it} *is a subset of an Euclidean space, that an allocation rule maximizing expected virtual surplus (14) exists, and that either of the following holds:*

(a) Condition U-COMP holds, and the virtual utility of each agent $i = 1, ..., N$,

$$U_i(\theta, x) - \frac{1}{\eta_{i0}(\theta_{i0})} \sum_{t=0}^{\infty} I_{it}(\theta_i^t) \partial U_i(\theta, x) / \partial \theta_{it},$$

has increasing differences (ID) in (θ, x), *and the same is true of the principal's utility* $U_0(x, \theta)$, *or*

(b) U-DSEP holds, and for each agent $i = 1, ..., N$, *each* $t \geq 0$, $X_{it} \subset \mathbb{R}$, *and there exists a nondecreasing function* $\phi_{it} : \Theta_i^t \to \mathbb{R}^t$ *such that the agent's virtual flow utility,*

$$u_{it}(\theta^t, x_t) - \frac{1}{\eta_{i0}(\theta_{i0})} \sum_{\tau=0}^{t} I_{i\tau}(\theta_i^\tau) \partial u_{it}(\theta^t, x_t) / \partial \theta_{i\tau},$$

depends only on $\varphi_{it}\left(\theta_i^t\right)$ *and* x_{it} *and has Strict Increasing Differences (SID) in* $\left(\varphi_{it}\left(\theta_i^t\right), x_{it}\right)$*, while the principal's flow utility depends only on* x_t*.*

Then the Relaxed Program (13) has a solution in which for all $i = 1, ..., N$*,* $\chi_i\left(\theta_i, \theta_{-i}\right)$ *is nondecreasing in* θ_i *for any* $\theta_{-i} \in \Theta_{-i}$*.*

Let us discuss the meaning of the proposition, for simplicity starting with the case where U-DSEP holds and X_t is one-dimensional (so that either case (a) or case (b) in the proposition can be applied). As in the static setting, the key assumption that ensures the monotonicity of the solution to the relaxed program is that each agent's virtual utility have increasing differences (ID) in consumption and types. The difference from the static setting lies in how this increasing difference condition could be ensured. Namely, for the static setting, as well as for the initial period $t = 0$ of the dynamic setting, it suffices to assume that each period-0 utility $u_{i0}(\theta_0, x_0)$ has ID in consumption and types, for any $i = 0, ..., N$, and that the partials $-\partial u_{i0}\left(\theta_0, x_0\right) / \partial \theta_{i0}$ for each $i = 1, ..., N$ have ID (which is a third-derivative condition on the utility function), and that the hazard rates $\eta_{i0}(\theta_{i0})$ are nondecreasing. However, to ensure ID of virtual values in periods $t \geq 1$ in the dynamic setting, in addition to extending the previous assumptions to future utility flows, we also need to add the assumption that the impulse responses $I_{i\tau}(\theta_i^\tau)$ are nondecreasing in types. Indeed, those assumptions together ensure that the term capturing the agent's information rent, $-\frac{1}{\eta_{i0}(\theta_{i0})} I_{i\tau}(\theta_i^\tau) \partial u_{it}\left(\theta^\tau, x_\tau\right) / \partial \theta_{i\tau}$ has ID in consumption and types. Intuitively, nondecreasing impulse responses serve to ensure that higher types receive lower distortions, which helps ensure strong monotonicity of the solution to the relaxed program.

With multidimensional allocations, it no longer suffices to ensure that virtual utilities have increasing differences in (x, θ): In addition, one needs to ensure that this direct effect is not outweighed by the interactions across different allocation dimensions, both across periods for a given agent (as in a dynamic programming problem) and across agents. In case (a) of Proposition 8, all the interactions are positive by Condition U-COMP. Examples where U-COMP is satisfied include (i) provision of a public good, and (ii) models with "learning by doing" where $u_i(\theta_i, x_i)$ is strictly supermodular in x_i, e.g. through a cost complementarity. On the other hand, U-COMP is not satisfied when the problem is to allocate private goods in limited supply (as in auctions), which creates negative interactions across agents. To deal with such situations, case (b) of Proposition 8 assumes separability across periods (U-DSEP) and uses the "aggregation method" to ensure that each agent's consumption is nondecreasing in his own type. This aggregation method requires that payoffs depend only on the components of the allocation that are observable to the agents, x_{it}, and that values are private. Given these assumptions, the functions ϕ_{it} then capture the aspects of θ_i^t that are relevant for agent i's payoff. The aggregation assumption then ensures that a given increase in θ_i^t either increases all differences from x_{it} to x_{it}' – by increasing the aggregator ϕ_{it} – or keeps them all constant. (Without this extra assumption, we could have a problem selecting one maximizer out of many — e.g. we can always select the maximizer so that χ_i is nondecreasing in θ_i for any given i, but it might be impossible to find a selection for which this monotonicity holds for all i.)

While the assumptions in Proposition 8 are rather strong, we want to emphasize that they are merely sufficient but not necessary for our first-order approach to be valid. Namely, even when the solution to the relaxed program fails to be strongly monotone, it could still be implementable, which could be verified by checking directly the single-crossing condition (24), or one of the conditions in Proposition 7. For an illustration, consider Example 3 (see also section 6.1). In this example, Proposition 8 can be applied to ensure strong monotonicity when $\phi''(\theta) \leq 0$, which means nonincreasing impulse responses. Conversely, when this condition does not hold, the solution to the relaxed program violates strong monotonicity, since the agent's consumption is actually *decreasing* in his past types. Nevertheless, in such cases one may still be able to establish that the solution to the relaxed program is implementable, for example by verifying its ex-post monotonicity. For the case of a linear utility function of Example 3, ex-post monotonicity means that

$$\sum_{\tau=t}^{\infty} \delta^{\tau-t} I_{(t),\tau}\left(\theta_i^\tau\right) \chi_t(\hat{\theta}_{it}, \theta_i^{-t})$$

– the present discount value of future consumption weighted by the impulse responses – is nondecreasing in $\hat{\theta}_{it}$. To ensure that the solution to the relaxed program in Example 3 satisfies ex-post monotonicity, it suffices to assume that $\eta_0(\theta) \geq 1/a$ and $\phi'(\theta) \leq 1$ for all θ (to rule out corner solutions), that $\eta_0(\theta_0)$ is nondecreasing in θ_0, and also that for all θ, $\phi''(\theta) \leq (1-\delta)/(\delta a)$. In particular, if $\phi''(\theta) \in [0, (1-\delta)/(\delta a)]$, strong monotonicity is violated since raising θ_t reduces the agent's future consumption, but because of discounting, this is outweighed by the "efficiency effect" of raising current consumption, and ex-post monotonicity still obtains.

6 Applications

We now show how the results in the previous sections can be put to work in a few applications. First, we consider the design of selling mechanisms for new experience goods, where bidders refine their valuations endogenously upon consumption. Next, we consider the optimal stopping problem faced by a durable-good monopolist selling an indivisible item to a population of agents whose valuations change exogenously with time. Lastly, we consider a family of time-separable environments where, by design, incentives separate over time and where optimal mechanisms can be easily constructed also for non-Markov processes.

6.1 Bandit Auctions

Consider the profit-maximization problem of a seller who repeatedly auctions off an indivisible, nonstorable, good to a set of $N \geq 1$ bidders who refine their valuations upon winning the auction. This setup captures novel applications such as repeated sponsored search auctions where the advertisers privately learn about the expected profitability of clicks on their ads, or repeated procurement with learning-by-doing. It provides a natural example of an environment where the type distributions depend on past allocations.

Formally, consider the following Markov environment. Let $X_{it} \equiv \{0,1\}$ for each $i = 0, ..., N$, $t \geq 0$, and let $X = \left\{x \in \prod_{t=0}^{\infty} \prod_{i=0}^{N} X_{it} : \sum_{i=0}^{N} x_{it} = 1 \; \forall t \geq 0\right\}$. The seller's payoff function takes the form $U_0(\theta, x) = -\sum_{t=0}^{\infty} \delta^t \sum_{i=1}^{N} x_{it} c_i$, where $c_i \in \mathbb{R}$ is the cost of allocating the object to agent i (notice that we have normalized $c_0 = 0$). The bidders' payoff functions take the form $U_i(\theta, x) = \sum_{t=0}^{\infty} \delta^t \theta_{it} x_{it}$ for $i = 1, ..., N$. These functions trivially satisfy Conditions U-TD, U-ELC, and U-MSPR.

The type process for any bidder $i = 1, ..., N$ is determined as follows. Let $R_i = (R_i(\cdot|k))_{k \in \mathbb{N}}$ be a sequence of absolutely continuous and strictly increasing c.d.f.'s with mean bounded in absolute value uniformly in k. The first period valuation θ_{i0} is drawn from $\Theta_{i0} = (\underline{\theta}_{i0}, \bar{\theta}_{i0})$, with $\underline{\theta}_{i0} > -\infty$, according to an absolutely continuous, strictly increasing c.d.f. F_{i0}. For any $t > 0$, $\theta_i^t \in \Theta_i^t$, and $x_i^{t-1} \in X_i^{t-1}$, if $x_{i,t-1} = 1$, then

$$F_{it}(\theta_{it}|\theta_{i,t-1}, x_i^{t-1}) = R_i(\theta_{it} - \theta_{i,t-1}| \sum_{\tau=0}^{t-1} x_{i\tau});$$

If, instead, $x_{i,t-1} = 0$, then

$$F_{it}(\theta_{it}|\theta_{i,t-1}, x_i^{t-1}) = \begin{cases} 0 & \text{if } \theta_{it} < \theta_{i,t-1}, \\ 1 & \text{if } \theta_{it} \geq \theta_{i,t-1}. \end{cases}$$

This formulation embodies the following key assumptions: (1) Bidders' valuations change only upon winning the auction (i.e., if $x_{i,t-1} = 0$, then $\theta_{ti} = \theta_{i,t-1}$ almost surely); (2) The valuation processes are time-homogenous (i.e., if bidder i wins the object in period $t-1$, then the distribution of his period t valuation depends only on his valuation in period $t-1$ and the total number of times he won in the past).[47]

In order to derive the relaxed problem, we use the canonical state representation from Proposition 1 to find the impulse responses. By (10) we have that, for all $i = 1, \ldots, N$, $t > 0$, $\varepsilon_{it} \in (0,1)$, and

[47] This kind of structure arises, for example, in a Bayesian learning model with Gaussian signals. That is, suppose that each bidder i has a constant but unknown true valuation v_i for the object and starts with a prior belief $v_i \sim N(\theta_{i0}, \tau_i)$ where the precision τ_i is common knowledge. Each time upon winning the auction, bidder i receives a private signal $s_i \sim N(v_i, \sigma_i)$ and updates his expectation of v_i using standard projection formulae. Take θ_{it} to be bidder i's posterior expectation in period t. Then $R_i(\cdot|k)$ is the c.d.f. for the change in the posterior expectation due to the k-th signal, which is indeed independent of the current value of θ_{it}. (Standard calculations show that $R_i(\cdot|k)$ is in fact a Normal distribution with mean zero and variance decreasing in k.) Alternative specifications for R_i can be used to model learning-by-doing, habit formation, etc.

all $(\theta_{i,t-1}, x_i^{t-1}) \in \Theta_{i,t-1} \times X_i^{t-1}$,

$$z_{it}(\theta_{i,t-1}, x_i^{t-1}, \varepsilon_{it}) = F_{it}^{-1}(\varepsilon_{it}|\theta_{i,t-1}, x_i^{t-1}) = \theta_{i,t-1} + \mathbf{1}_{\{x_{i,t-1}=1\}} R_i^{-1}(\varepsilon_{it}| \sum_{\tau=0}^{t-1} x_{i\tau}).$$

Thus using (7) and (4) we let $I_{it}(\theta_i^t, x_i^{t-1}) = 1$ for all $i = 1, ..., N$, $t \geq 0$, $\theta_i^t \in \Theta_i^t$, and $x_i^{t-1} \in X_i^{t-1}$. It is straightforward to verify that these processes satisfy Conditions F-BIR and F-MB (and hence Condition F-BE) so that Proposition 2 applies with the dynamic virtual surplus taking the form[48]

$$\mathbb{E}^{\lambda[\chi]} \left[\sum_{t=0}^{\infty} \delta^t \sum_{i=1}^{N} \left(\tilde{\theta}_{it} - c_i - \frac{1}{\eta_{i0}(\tilde{\theta}_{i0})} \right) \chi_{it}(\tilde{\theta}^t) \right].$$

The relaxed problem is then a standard multiarmed-bandit problem with a safe arm corresponding to not selling and yielding a sure payoff equal to zero, and N risky arms with arm $i = 1, ..., N$ corresponding to bidder i and yielding in each period $t \geq 0$ a current payoff $\theta_{it} - c_i - [\eta_{i0}(\theta_{i0})]^{-1}$. The solution takes the form of a virtual index policy. That is, define the virtual (Gittins) index of bidder $i = 1, ..., N$ in period $t \geq 0$ given history $(\theta_i^t, x_i^{t-1}) \in \Theta_i^t \times X_i^{t-1}$ as

$$\gamma_{it}(\theta_i^t, x_i^{t-1}) \equiv \max_T \mathbb{E}^{\lambda[\bar{\chi}_i]|\theta_{it}, x_i^{t-1}} \left[\frac{\sum_{\tau=t}^{T} \delta^\tau \left(\tilde{\theta}_{i\tau} - c_i - \frac{1}{\eta_{i0}(\theta_{i0})} \right)}{\sum_{\tau=t}^{T} \delta^\tau} \right],$$

where T is a stopping time, and $\bar{\chi}_i$ is the allocation rule that assigns the object to bidder i in all periods. (Note that the index depends on θ_i^t and x_i^{t-1} only through θ_{i0}, θ_{it}, and $\sum_{\tau=0}^{t-1} x_{i\tau}$.) The index of the seller is identically equal to zero in all periods and for convenience we write it as $\gamma_{0t}(\theta_0^t, x_0^{t-1}) \equiv 0$. We then have the following result (proof in the Appendix):

Proposition 9 (Bandit auctions) *The allocation rule χ solves the relaxed program for the bandit auctions problem described above if for all $i = 1, ..., N$, $t \geq 0$, $\theta^t \in \Theta^t$,*

$$\chi_{it}(\theta^t) = 1 \quad \Rightarrow \quad i \in \arg \max_{j \in \{0,...,N\}} \gamma_{jt}(\theta_j^t, \chi_j^{t-1}(\theta^{t-1})),$$

i.e., if χ is a virtual index policy. Furthermore, if the hazard rates $\eta_{i0}(\theta_{i0})$, $i = 1, \ldots, N$, of the initial type distributions are nondecreasing in θ_{i0}, then any virtual index policy can be implemented as a strong truthtelling PBE and is part of a profit-maximizing mechanism.[49]

The optimality of virtual index policies in the relaxed program follows by standard arguments (see, e.g., Whittle, 1982 and Bergemann and Välimäki, 2008). Before discussing implementability, we

[48] For F-MB, note that $\mathbb{E}^{F_{it}(\theta_{i,t-1}, x_i^{t-1})} \left[\left| \tilde{\theta}_t \right| \right] \leq |\theta_{i,t-1}| + B$ for some constant B independent of t by the uniform bound on the means of $R_i(\cdot|k)$, $k \in \mathbb{N}$.

[49] In fact, it is possible to implement any virtual index policy in a periodic ex-post equilibrium by simply changing the probability distributions in the definition of the transfers in the proof of Theorem 5 to condition on the contemporaneous types of the other agents.

first compare the optimal allocation rule to its efficient counterpart, which maximizes the expected social surplus $\mathbb{E}^{\lambda[\chi]}\left[\sum_{t=0}^{\infty}\delta^t\sum_{i=1}^{N}(\tilde{\theta}_{it}-c_i)\chi_{it}(\tilde{\theta}^t)\right]$. The latter is also an index policy where the index of bidder $i=1,...,N$ with current type $\theta_{it}\in\Theta_{it}$ and allocation history $x_i^{t-1}\in X_i^{t-1}$ is

$$\gamma_{it}^*(\theta_{it},x_i^{t-1})=\max_T \mathbb{E}^{\lambda[\bar{\chi}_i]|\theta_{it},x_i^{t-1}}\left[\frac{\sum_{\tau=t}^{T}\delta^\tau(\tilde{\theta}_{i\tau}-c_i)}{\sum_{\tau=t}^{T}\delta^\tau}\right].$$

The efficient index policy can be implemented with the Team mechanism of Athey and Segal (2007), or with the Dynamic Pivot mechanism of Bergemann and Välimäki (2010). By inspection of the indices, we see that $\gamma_{it}(\theta_i^t,x_i^{t-1})=\gamma_{it}^*(\theta_{it},x_i^{t-1})-[\eta_{i0}(\theta_{i0})]^{-1}$ for all $i=1,...,N$, $t\geq 0$, $\theta_i^t\in\Theta_i^t$, $x_i^{t-1}\in X_i^{t-1}$. That is, the profit-maximizing virtual index for bidder i is obtained from the efficient one by simply subtracting the "handicap" $[\eta_{i0}(\theta_{i0})]^{-1}$ which is responsible for i's information rent, and which is entirely determined by his initial type θ_{i0} and is constant over time. As the seller's index is zero in both problems, the constant handicaps imply *permanent* distortions under the profit-maximizing policy. Moreover, when initial hazard rates are nondecreasing, bidders with higher initial types have lower handicaps and in effect receive preferential treatment in all periods over those with lower initial types.

Also note that, with $N\geq 2$, distortions are in general in both directions. The presence of the handicaps leads the seller to keep the object too often, which implies downward distortions at some histories. However, since the efficient indices are independent of θ_{i0}, there can exist histories $(\theta^t,x^{t-1})\in\Theta^t\times X^{t-1}$ such that for some bidders i and j, $\gamma_{it}(\theta_i^t,x_i^{t-1})>\gamma_{jt}(\theta_j^t,x_j^{t-1})$, $\gamma_{it}(\theta_i^t,x_i^{t-1})>0$, and $\gamma_{it}^*(\theta_{it},x_i^{t-1})<\gamma_{jt}^*(\theta_{jt},x_j^{t-1})$. At any such history the current allocation of bidder i is distorted upwards. For example, in the special case of pure Bayesian learning about the true valuation of each bidder i, this implies that the profit-maximizing policy settles on an inefficient bidder more often than the efficient policy. Furthermore, despite the handicaps being constant, distortions may increase over time. To see this, consider the case where all bidders are ex-ante identical and consider a state at which $\gamma_{i0}(\theta_{i0})>0$ all i. Symmetry then implies that the profit-maximizing rule allocates the good efficiently in the first period. This, however, need not be the case in subsequent periods, because of the distortions in the relative ranking of the virtual indexes introduced by the handicaps. Furthermore, because the indexes also depend on the number of past wins, such inefficiencies can cumulate over time, making distortions increase with the number of preceding periods.

As for the implementability of virtual index policies, observe that the processes satisfy Condition F-FOSD. Thus, in order to apply Theorem 5, it suffices to show that when initial hazard rates are monotone, any virtual index policy χ, together with the appropriate beliefs Γ (recall that these are essentially determined by χ) satisfies the single-crossing condition (24) in each period s. We do this by verifying weak monotonicity of Proposition 7, which here requires that for any $i=1,...,N$, $s\geq 0$,

$\theta^{s-1} \in \Theta^{s-1}$, $\theta_{is} \in \Theta_{is}$, bidder i's expected discounted consumption

$$\mathbb{E}^{\lambda_i[\chi \circ \hat{\theta}_{is},\Gamma]|\theta^{s-1},\theta_{is}} \left[\sum_{t=s}^{\infty} \delta^t (\chi \circ \hat{\theta}_{is})_{it}(\tilde{\theta}) \right]$$

is non-decreasing in his current bid $\hat{\theta}_{is}$. (Above, $(\chi \circ \hat{\theta}_{is})_{it}(\tilde{\theta})$ denotes bidder i's period-t allocation under the rule $\chi \circ \hat{\theta}_{is}$.) Heuristically, this follows since a higher bid today will result in bidder i consuming sooner in the sense that, for any $k \in \mathbb{N}$, the expected waiting time until he wins the auction for the k-th time after period s is decreasing in the current bid $\hat{\theta}_{is}$. (For details, see the Appendix.) Theorem 5 then implies that there exists spreadable payments Ψ such that the direct mechanism for the choice rule $\langle \chi, \Psi \rangle$ has a PBE in strong-truthtelling strategies and it maximizes the principal's profits among all BNIC-IR mechanisms.

Lastly, note that the optimal mechanism for bandit auctions is essentially a long-term contract that charges initial fees and grants permanent preferential treatment based on the bidders' initial types (which may, e.g., be their initial private estimates about their true valuations). These features are markedly different from running a sequence of second price auctions with a reserve price, and suggest potential advantages of building long-term contractual relationships in repeated procurement and sponsored search.

6.2 Durable-Good Monopolist

In this setting, the principal is a seller of a single indivisible durable good facing N potential buyers whose valuations change exogenously with time (for example, as the result of the collection of information from private sources). Once the good is sold to a buyer, it cannot be reallocated. This scenario could correspond, for example, to the sale of a licence by the government that cannot be traded in a secondary market. The problem that the seller faces then consists in designing a profit-maximizing sequence of auctions, whereby the allocation and price rules in each round may depend on the bidding in the preceding rounds.

We model this problem by letting $X_{it} = \{0, 1\}$ for each $i = 1, ..., N$, $t \geq 0$, so that $x_{it} \in \{0, 1\}$ describes buyer i's consumption of the good in period t. The utility function for each player i (including the seller) is $U_i(x, \theta) = \sum_{t=0}^{\infty} \delta^t \theta_{it} x_{it}$, where θ_{0t} are publicly observed and allowed to follow an arbitrary Markov process. The feasible set is given by

$$X = \left\{ x \in \prod_{t=0}^{\infty} \prod_{i=1}^{N} X_{it} : \sum_{i=1}^{N} x_{it} \leq 1 \text{ for all } t \geq 0 \text{ and } x_{it} \text{ is nondecreasing in } t \text{ for all } i = 1, ..., N \right\}.$$

Together the two constraints say that there is only one good to be allocated and that, once the good is sold to a buyer, it cannot be reallocated (i.e., rentals are not allowed). Note that the latter constraint makes the problem a dynamic programming one—both for the first-best and for the profit-maximizing allocation.

In the first-best program, there are two possible motives for the seller to postpone the sale: (1) his flow values can exceed the buyers' values, and (2) waiting permits her to receive more accurate information pertaining to the efficient allocation of the object. Under profit-maximization, a third motive for delay comes from the fact that delay may reduce the agents' information rents. This motive is similar to the familiar durable-good monopolist incentive to postpone sales to discriminate among consumers (see, e.g., Board (2008)). However, a key difference here is that the optimal mechanism asks the buyers to participate, send messages (e.g., bids), and make payments even prior to allocating the object to them. For example, a buyer can be asked to make an initial payment to reduce the reservation prices he will face during the subsequent bidding and thus anticipate the expected time he will start consuming the good.

Now note that the payoffs described above clearly satisfy Conditions U-TD, U-ELC, and U-MSPR. As for the processes, we assume that they are all Markov and satisfy Conditions F-BIR, F-MB, F-AUT, and F-FOSD, so that all the conditions of Theorem 5 apply. We now establish sufficient conditions for a solution to the relaxed program to satisfy strong monotonicity, and therefore to be PBE-implementable and part of an optimal mechanism. Note that, while this framework does not formally fit either case (a) or case (b) of Proposition 8 (the feasible set is not a lattice and the problem is not time-separable), the arguments in the proof (in the Appendix) are similar to those in that proposition.

Proposition 10 (Durable-good monopolist) *In addition to the conditions described above, assume that, for all $i = 1, ..., N$, $s \geq 0$, $t \geq s$, the impulse responses $I_{i,(s),t}\left(\theta_i^t\right)$ are strictly positive and nonincreasing in θ_i^t, and that the hazard rates $\eta_{i0}(\theta_{i0})$ of the initial type distributions are nondecreasing. Then, if the relaxed problem admits a solution, it also admits a solution χ that is strongly monotone (i.e., in which $\chi_{it}\left(\theta^t\right)$ is nondecreasing in θ_i^t for each $i = 1, ..., N$, $t \geq 0$, $\theta_{-i}^t \in \Theta_{-i}^t$), in which case the conclusions of Theorem 5 and Corollary 2 apply (In particular, (i) the above allocation rule is part of an optimal mechanism and (ii) it can be implemented in a periodic ex-post equilibrium, implying that the information disclosed to the agents is irrelevant).*

As for the optimal stopping times for the profit-maximizing and the efficient rule[50] they are obtained as follows. First, normalize the seller's types to zero and let the bidders' types define their valuations, net of the seller's values. For each $i = 1, ..., N$, $s \geq 0$, $\theta_i^s \in \Theta_i^s$, then let

$$J_{is}(\theta_i^s) \equiv \mathbb{E}^{\lambda|\theta_i^s}\left[\sum_{t=s}^{\infty} \delta^t \left(\tilde{\theta}_{it} - \frac{1}{\eta_{i0}(\theta_{i0})} I_{it}(\tilde{\theta}_i^t)\right)\right]$$

denote the expected virtual NPV of selling the good to buyer i in period s when the latter's type history is θ_i^s. Then let $W(\theta^t) = \max_i J_{it}(\theta_i^t)$. Finally, let $SN(\theta^t)$ denote the Snell Envelope of the

[50] To the best of our knowledge, the efficient rule for this problem is not known. The result in Proposition 10, however, implies that if an efficient rule exists, there also exists one that is strongly monotone.

process corresponding to $W(\theta^t)$.[51] The optimal stopping time $\tau^{stop}(\tilde{\theta})$ corresponds to the first time that $W(\theta^t) = SE(\theta^t)$. The good is then allocated to the buyer whose expected virtual NPV $J_{i\tau^{stop}}(\theta_i^t)$ in period τ^{stop} is the highest.

The optimal stopping time for the efficient rule is obtained in a similar way by replacing the virtual NPVs J_{it} with the true ones J_{it}^*. Both the profit-maximizing and the efficient rule can be implemented by a sequence of "score" auctions. In the case of efficiency, because the processes are Markov and satisfy F-FOSD, these scores are simply monotone transformations of current values. In the case of profit-maximization, the scores are more complicated increasing functions of current and past bids. The good is then sold the first time a bidder's score exceeds all the opponents' score by a pre-specified amount that is increasing in the opponents' bids and decreasing in his own bids.

6.3 Time-Separable Environments

Consider now the following family of models that can fit many applications including sequential auctions, procurement, and regulation. The two key assumptions are that payoffs and decisions separate over time and that impulse responses depend only on current and initial types and are nonincreasing in $(\theta_{i0}, \theta_{it})$. Formally, let $X = \prod_{t=1}^{\infty} X_t$ with each $X_t \subset \mathbb{R}^{N+1}$, and for any $i = 1, ..., N$, assume that

$$U_i\left(\theta, x\right) = \sum_{t=0}^{\infty} \delta^t u_{it}(\theta_{it}, x_{it}),$$

with the flow payoff u_{it} satisfying the following assumptions: (i) $|u_{it}\left(\theta_{it}, x_{it}\right)| \leq L_i \left|\theta_{it}\right| + M_{it}$ for all $(\theta_{it}, x_{it}) \in \Theta_{it} \times X_{it}$, with $L_i, \|M_i\| < \infty$, and (ii) $u_{it}\left(\theta_{it}, x_{it}\right)$ is differentiable in θ_{it} with partial derivative $\partial u_{it}(\theta_{it}, x_{it})/\partial\theta_{it}$ positive, bounded uniformly by a constant $K_i > 0$, strictly increasing in x_{it}, and strictly submodular in (θ_{it}, x_{it}). For $i = 0$ (i.e., for the principal), simply assume that $U_0\left(\theta, x\right) = \sum_{t=0}^{\infty} \delta^t u_{0t}(x_t)$.

As for the stochastic processes, assume that they satisfy F-BE, F-BIR, F-AUT and F-FOSD. In addition, assume that first-period hazard rates $\eta_{i0}\left(\theta_{i0}\right)$ are nondecreasing and that, for any $i = 1, ..., N, t \geq 0$, the impulse responses $I_{it}(\theta_i^t)$ depend only on $(\theta_{i0}, \theta_{it})$ and are nonincreasing in $(\theta_{i0}, \theta_{it})$. An example of such processes is when the agents' types follow an AR(k) process as in Example 1 with impulse responses and shock terms satisfying $\|I_i\|, \|\mathbb{E}\left[|\tilde{\varepsilon}_i|\right]\| < \infty$, all $i = 1, ..., N$. Notice that the processes are not restricted to be Markov and that the environment satisfies all the conditions in part (b) of Proposition 8, with the functions φ given by $\varphi_{it}(\theta_i^t) = (\theta_{i0,}\theta_{it})$ all $i = 1, ..., N$, $t \geq 0$. We then have the following result:

[51] This is the smallest supermartingale that dominates W.

Proposition 11 (Separable environments) *Consider the separable environment described above and assume that the relaxed program admits a solution.*[52] *Then there exists an optimal mechanism* (χ, Ψ) *such that*

(a) the allocation rule χ *is strongly monotone, with* $\chi_t(\theta^t)$ *depending only on* (θ_0, θ_t) *all* $t \geq 0$, $\theta^t \in \Theta^t$ *and satisfying, for all* $t \geq 0$, λ*-almost all* θ^t,

$$\chi_t(\theta^t) \in \arg\max_{x_t \in X_t} \left\{ u_{0t}(\theta_{0t}, x_{0t}) + \sum_{i=1}^{N} \left(u_{it}(\theta_{it}, x_{it}) - \frac{1}{\eta_{i0}(\theta_{i0})} I_{it}(\theta_{i0}, \theta_{it}) \frac{\partial u_{it}(\theta_{it}, x_{it})}{\partial \theta_{it}} \right) \right\}, \tag{27}$$

(b) the payments Ψ *are spread over time as follows:*

$$\psi_{it}(\theta_0, \theta_t) = -u_{it}(\theta_{it}, \chi_{it}(\theta_0, \theta_t)) + \int_{\underline{\theta}_{it}}^{\theta_{it}} \frac{\partial u_{it}(r, \chi_{it}(\theta_0, (r, \theta_{-i,t})))}{\partial \theta_{it}} dr \quad \text{all } t \geq 1 \tag{28}$$

and

$$\psi_{i0}(\theta_0) = -\mathbb{E}^{\lambda|\theta_0} \left[U_i(\tilde{\theta}_i, \chi(\tilde{\theta})) + \sum_{t=1}^{\infty} \delta^t \psi_{it}(\theta_0, \tilde{\theta}_t) \right] + \int_{\underline{\theta}_{i0}}^{\theta_{i0}} \mathbb{E}^{\lambda_i|r} \left[\sum_{t=0}^{\infty} \delta^t I_{it}(\tilde{\theta}_{i0}, \tilde{\theta}_{it}) \frac{\partial u_{it}(\tilde{\theta}_{it}, \chi_{it}(\tilde{\theta}_t))}{\partial \theta_{it}} \right] dr; \tag{29}$$

(c) reporting truthfully at all histories is a periodic ex-post equilibrium.

Proof. The result follows from combining Theorem 4 with Proposition 8. First note that, because the environment is time-separable, an allocation rule χ maximizes the dynamic virtual surplus–and hence solves the relaxed program–if and only if for all $t \geq 0$, λ-almost all $\theta^t \in \Theta^t$, $\chi_t(\theta^t)$ satisfies (27). Furthermore, because the environment satisfies all the conditions of Proposition 8, it is easy to see that there exists a solution to the relaxed program that is strongly monotone and such that for all $t \geq 0$, all $\theta^t \in \Theta^t$, $\chi_t(\theta^t)$ depends only on (θ_0, θ_t). To see that the proposed payments implement this rule and satisfy properties (b) and (c) in the proposition, then note that, under the proposed mechanism, *incentives separate over time*, starting from $t = 1$ onwards. That is, in each period $t \geq 1$, and for any history, each agent maximizes his payoff by simply choosing the current message so as to maximize his flow payoff $u_{it} + \psi_{it}$. This follows from the fact, at any $t \geq 1$, each agent i's message m_{it} has no direct effect on the allocations in periods $\tau > t$ and, because the environment is time-separable, the allocation in period t has no direct effect on the future flow payoffs. That, given the proposed payments, the agent finds it optimal to report truthfully, irrespective of his beliefs about the other agents' types and messages, and irrespective of whether or not he has been truthful in the past, then follows from standard results from static mechanism design by observing that (i) the allocation $\chi_{it}(\theta_0, \theta_t)$ is monotone in θ_{it}, for all $(\theta_{-i,t}, \theta_0)$, and (ii) that values are private., i.e., each u_{it} depends only on own θ_{it}. In other words, it is as if each agent i were facing a single-agent static decision problem indexed by $(\theta_{-i,t}, \theta_0)$. In particular, note that reporting truthfully remains optimal

[52] This can be ensured, for example, by assuming that, for all $t \geq 0$, X_t is compact and that, for all θ, all $t \geq 0$, u_{it}, $i = 0, .., N$, are continuous over X_t.

even if each agent were able to observe both his own future types, the other agents' past, current, and future types, and the messages sent by the other agents, thus making his beliefs completely irrelevant.

As for period zero, because strong truthtelling is optimal at all period-1 histories, irrespective of whether or not the processes are Markov, the optimality of truthtelling at $t = 0$ follows from Theorem 4 along with the result in Proposition 7.[53] Also note that, because each $u_{it}(\theta_{it}, x_{it})$ is nondecreasing in θ_{it}, under the proposed mechanism, participating is optimal for all period-0 types. That the proposed mechanism maximizes the principal's payoff then follows from essentially the same arguments as in the proof of Theorem 5. ■

We conclude by discussing possible implementations of the profit-maximizing rule. First, consider the special case where flow payoffs are linear (i.e., $u_{it}(\theta_{it}, x_{it}) = \theta_{it}x_{it}$) as in auctions, and where types evolve according to an AR(k) process (or, more generally, any process for which the impulse responses I_{it} depend only on the initial types). The implementation is then particularly simple. Suppose that there is no allocation in the first period and assume that the agents do not observe the other agents' types (both assumptions simplify the discussion but are not essential for the argument). In period zero, each agent i is then asked to choose from a menu of "handicaps" $(I_{it}\eta_{i0}^{-1}(\theta_{i0}))_{t=1}^{\infty}$, indexed by θ_{i0}, with each handicap costing $\psi_{i0}(\theta_{i0})$, where $\psi_{i0}(\theta_{i0})$ is as in (29) but with the measure $\lambda|\theta_0$ replaced by the measure $\lambda_i|\theta_{i0}$. Then in each period $t \geq 1$, a "handicapped" VCG mechanism is played with transfers as in (28). (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 and more general processes, in the sense that in the first period the agents still choose from a menu of future plans (indexed by the first-period type). However, in general, in the subsequent periods the distortions will depend also on the current reports through the partial derivatives $\partial u_{it}(\theta_{it}, x_{it})/\partial\theta_{it}$ and through the impulse responses $I_{it}(\theta_{i0}, \theta_{it})$. However, as long as payoff separate over time and the impulse responses depend only on initial and current types, the intermediate reports (i.e., reports in periods $1, \ldots, t-1$) remain irrelevant both for the period-t allocation and for the period-t payments.

Finally, note that the logic we used to establish sufficiency in this family of problems (VCG payments in each period $t > 0$ along with monotonicity of the allocation rule and payments at $t = 0$ given by (29)) extends to a few environments where payoffs are not time-separable, but where virtual

[53] That in period 0 the agent finds it optimal to report truthfully irrespective of his beliefs about the other agents' types follows from the fact that truthful reporting is optimal for each $\theta_{-i,0}$. Together with the fact that truthful reporting is also optimal in each subsequent period irrespective of the beliefs, this establish that the proposed payments implement the allocation rule as a periodic ex-post equilibrium. Because valuations are private, the reader may wonder whether truthful reporting is actually dominant in the proposed mechanism. The answer is negative; under arbitrary strategies, the opponents' future actions may depend on their observed allocations in a way that may induce the agent to depart from truthful reporting.

payoffs continue to be an affine transformation of the true payoffs with constants that depend only on the initial reports. Consider, for example, an economy where all processes are AR(1) and where payoffs are given by $U_i = \sum_{t=0}^{\infty} \delta^t \theta_{it} x_{it} - c_{it}(x^t)$, all $i = 0, ..., N$. The dynamic virtual surplus then coincides with the true surplus of a fictitious economy where all agents' payoffs are as in the original economy and where the principal's payoff is given by $U_0 - \sum_{i=1}^{N} \eta_{i0}^{-1}(\theta_{i0}) \sum_{t=0}^{\infty} \delta^t x_{it}$. Then note that, irrespective of whether or not the agents reported truthfully in period zero, the allocation rule χ that solves the relaxed program maximizes the surplus of this fictitious economy from period $t = 1$ onwards. This property, together with the fact that values are private from $t = 1$ onwards in this fictitious economy, then implies that incentives for truthful reporting at any period $t \geq 1$ can be provided by using either the team payments of Athey and Segal (2007) or the pivot payments of Bergemann and Välimäki (2010). Furthermore, as long as the rule χ is weakly monotone in the initial reports θ_0, then incentives can also be provided in period 0 by adding to the aforementioned payments an initial payment given by (29).[54]

7 Conclusions

We considered the design of optimal mechanisms for dynamic settings where information arrives gradually and where decisions are made over time. The model allows for both finite and infinite horizon, fairly general payoffs, many agents, a continuum of types, and decision-controlled processes that are not necessarily Markov and possibly have mass points. The generality of the model served two purposes. First, it permitted us to unify the existing literature, identifying general principles and highlighting what drives similarities and differences in the special cases considered. Second, it offered a flexibility that facilitates novel applications.

The core of the analysis is a formula that describes how equilibrium payoffs respond to the arrival of new information. This formula is the dynamic analog of the familiar envelope formula for static settings. Part of the contribution is in having identified primitive conditions that validate this formula as a necessary condition for incentive compatibility in any Bayes-Nash incentive compatible mechanism. This step is key to the design of *optimal* mechanisms.

We then showed how to construct payment schemes that guarantee that this formula is satisfied not only at the initial contracting stage but throughout the entire execution of the mechanism, which is necessary for implementability under stronger solution concepts such as PBE. We qualified in what sense such payments are essentially pinned down by the allocation rule, thus extending the celebrated "revenue equivalence" result to dynamic settings.

Next, we showed how optimal mechanisms can be obtained by first solving a relaxed program

[54] For an application of similar ideas to a family of problems where the dynamic virtual surplus takes the form of a multiplicative transformation of the true surplus, with the scale depending only on the initial reports, see the recent paper by Kakade et al. (2011).

that consists in searching for an allocation rule that maximizes the dynamic virtual surplus (the latter uses only the first-order conditions for incentive compatibility at the initial contracting stage to express the principal's payoff entirely in terms of the allocation rule and the lowest period-0 types) and then verifying that the solution to the relaxed program is truly incentive compatible. For this last step, we identified sufficient conditions that establish PBE-implementability either by backward induction (e.g., in finite-horizon settings), or by establishing the suboptimality of one-step deviations from truthtelling and then identifying conditions that guarantee that this implies the suboptimality of all other deviations. We showed that, under certain technical conditions, sufficiency is facilitated by two properties: (i) Markoveness of the environment (this permits one to restrict attention to truthful histories when verifying the suboptimality of one-step deviations) and (ii) a certain single-crossing condition for the derivative of the equilibrium payoffs with respect to the true types in the reported message. This single crossing condition uses the same envelope formula derived for the necessary conditions. We then showed how, under supermodular payoffs, this single crossing condition is implied by the allocation rule being sufficiently monotone, in a sense made precise in the paper, and identified primitive conditions that guarantee that the solution to the relaxed program is indeed monotone.

Equipped with the aforementioned results, we then illustrated how the dynamics of the distortions is driven by the dynamics of the impulse responses of the subsequent types to the initial ones and used this insight to interpret the results obtained in the literature for specific processes.

Lastly, we showed how the flexibility of the model can facilitate novel applications by considering the design of bandit auctions (i.e., auctions for new experience goods) and of sale mechanisms for durable goods, such as government licenses or real estate. The results appear useful in various other applications (see, e.g., Garrett and Pavan (2011a,b), Farhi and Werning (2010), Golosov et al. (2010)).

Throughout, we maintained the assumption that the agents are time-consistent. Relaxing this assumption is challenging and represents an interesting line for future research (For some preliminary work in this direction, see Galperti (2011)).

Appendix

Proof of Lemma 1. Let us focus on those ε for which $\partial Z_{it}(\theta_{i0}, \hat{\chi}_i^{t-1}(\hat{\theta}_{i0}, \theta_{-i,0}, \varepsilon_i^{t-1}), \varepsilon_i^t)/\partial\theta_{i0} < C_{it}(\varepsilon_i)$ for all i, t, θ_0, with $\|C_i(\varepsilon_i)\| < \infty$, and $\|Z_i(\theta_{i0}, \hat{\chi}_i(\theta_{i0}, \theta_{-i,0}, \varepsilon), \varepsilon_i)\| < \infty$, which under Conditions F-BE$_0$, F-BIR$_0$ occurs with probability 1, and temporarily drop arguments ε, $\theta_{-i,0}$, $\hat{\theta}_{i0}$, $x = \hat{\chi}(\hat{\theta}_{i0}, \theta_{-i,0}, \varepsilon)$, and i to simplify notation.

The classical chain rule (using Conditions U-TD and F-BIR$_0$) yields that, for any given T,

$$\Delta_T(\theta_0, h) \equiv \frac{1}{h} U\left(Z^T(\theta_0 + h), Z^{>T}(\theta_0)\right) - \frac{1}{h} U(Z(\theta_0)) - \sum_{t=0}^{T} \frac{\partial U(Z(\theta_0))}{\partial \theta_t} Z_t'(\theta_0) \to 0 \text{ as } h \to 0. \quad (30)$$

Note that

$$\frac{1}{h} U\left(Z^T(\theta_0 + h), Z^{>T}(\theta_0)\right) \to \frac{1}{h} U(Z(\theta_0 + h)) \text{ as } T \to \infty$$

uniformly in h since, using U-ELC, the difference is uniformly bounded by

$$A\frac{1}{h} \sum_{t=T+1}^{\infty} \delta^t |Z_t(\theta_0 + h) - Z_t(\theta_0)| \leq A \sum_{t=T+1}^{\infty} \delta^t C_t,$$

and the right-hand side converges to zero as $T \to \infty$ since $\|C\| < \infty$.

Also, the series in (30) converges uniformly in h by the Weierstrass M-test, since, using Conditions U-ELC and F-BIR$_0$

$$\sum_{t=0}^{T} \left| \frac{\partial U(Z(\theta_0))}{\partial \theta_t} \right| \cdot \left| Z_t'(\theta_0) \right| \leq \sum_{t=0}^{T} A\delta^t \cdot C_t \to A \|C\| \text{ as } T \to \infty$$

Hence, we have

$$\Delta_T(\theta_0, h) \to \frac{1}{h} \left[\hat{U}(\theta_0 + h) - \hat{U}(\theta_0) \right] - \sum_{t=0}^{\infty} \frac{\partial U(Z(\theta_0))}{\partial \theta_t} Z_t'(\theta_0) \text{ as } T \to \infty$$

uniformly in h. The uniform convergence allows us to interchange the order of limits and use (30):

$$\begin{aligned} \lim_{h \to 0} \left[\frac{1}{h} \left[\hat{U}(\theta_0 + h) - \hat{U}(\theta_0) \right] - \sum_{t=0}^{\infty} \frac{\partial U(Z(\theta_0))}{\partial \theta_t} Z_t'(\theta_0) \right] &= \lim_{h \to 0} \lim_{T \to \infty} \Delta_T(\theta_0, h) \\ &= \lim_{T \to \infty} \lim_{h \to 0} \Delta_T(\theta_0, h) = 0. \end{aligned}$$

This yields (putting back all the missing arguments)

$$\frac{\partial \hat{U}_i(\hat{\theta}_{i0}, \theta_0, \varepsilon)}{\partial \theta_{i0}} = \sum_{t=0}^{\infty} \frac{\partial U_i\left(Z\left(\theta_0, \varepsilon, \hat{\chi}(\hat{\theta}_{i0}, \theta_{-i,0}, \varepsilon)\right), \hat{\chi}\left(\hat{\theta}_{i0}, \theta_{-i,0}, \varepsilon\right)\right)}{\partial \theta_{it}} \cdot \frac{\partial Z_{it}\left(\theta_{i0}, \varepsilon_i^t, \hat{\chi}_i^{t-1}\left(\hat{\theta}_{i0}, \theta_{-i,0}, \varepsilon^{t-1}\right)\right)}{\partial \theta_{i0}}.$$

Next, note that, being a composition of Lipschitz continuous functions, $\hat{U}_i(\hat{\theta}_{i0}, \cdot, \theta_{-i,0}, \varepsilon)$ is equi-Lipschitz continuous in θ_{i0} with constant $A \cdot \|C_i(\varepsilon)\|$. Since, by F-BIR$_0$, $\mathbb{E}[\|C_i(\tilde{\varepsilon})\|] < \infty$, by the Dominated Convergence Theorem we can write

$$\begin{aligned} \frac{\partial W_i\left(\hat{\theta}_{i0}, \theta_{i0}\right)}{\partial \theta_{i0}} &= \lim_{h \to 0} \mathbb{E} \left[\frac{\hat{U}_i\left(\hat{\theta}_{i0}, \theta_{i0} + h, \tilde{\theta}_{-i,0}, \tilde{\varepsilon}\right) - \hat{U}_i\left(\hat{\theta}_{i0}, \theta_{i0}, \tilde{\theta}_{-i,0}, \tilde{\varepsilon}\right)}{h} \right] \\ &= \mathbb{E} \lim_{h \to 0} \left[\frac{\hat{U}_i\left(\hat{\theta}_{i0}, \theta_{i0} + h, \tilde{\theta}_{-i,0}, \tilde{\varepsilon}\right) - \hat{U}_i\left(\hat{\theta}_{i0}, \theta_{i0}, \tilde{\theta}_{-i,0}, \tilde{\varepsilon}\right)}{h} \right] = \mathbb{E} \left[\frac{\partial \hat{U}_i\left(\hat{\theta}_{i0}, \theta_{i0}, \tilde{\theta}_{-i,0}, \tilde{\varepsilon}\right)}{\partial \theta_{i0}} \right]. \end{aligned}$$

Hence, for any $\hat{\theta}_{i0}$, $W_i\left(\hat{\theta}_{i0},\cdot\right)$ is differentiable and equi-Lipschitz continuous in θ_{i0}, with the Lipschitz constant $A\cdot\mathbb{E}\left[||C_i\left(\tilde{\varepsilon}\right)||\right]$, and with derivative at $\theta_{i0}=\hat{\theta}_{i0}$ given by the formula in the lemma. ■

Proof of Theorem 2. For part (i), we show that each of the three terms in ψ_{it} has a finite EPV in distribution $\lambda_i[\chi,\Gamma]|\theta^{s-1},\theta_{is}$ for any $i=1,...,N$, $s\geq 0$, and starting history $\left(\theta^{s-1},\theta_{is}\right)\in\Theta^{s-1}\times\Theta_{is}$ (which also implies that the series in (21) converges with probability 1 in that distribution). For the first term, using U-ELC and F-BIR, we have

$$\left|D_{it}^{\chi,\Gamma}\left(\theta^{t-1},\theta_{it}\right)\right|\leq\mathbb{E}^{\lambda_i[\chi,\Gamma]|\theta^{t-1},\theta_{it}}\left[\sum_{\tau=t}^{\infty}\left|I_{i,(t),\tau}(\tilde{\theta}_i,\chi_i(\tilde{\theta}))\right|A_i\delta^{\tau}\right]\leq\delta^t A_i B_i, \tag{32}$$

where $A_i>0$ is the constant of equi-Lipschitz continuity of the utility function, and where $B_i>0$ is the bound on the impulse responses, as implied by Condition F-BIR. This means that

$$\left|\delta^{-t}Q_{it}^{\chi,\Gamma}\left(\theta^{t-1},\theta_{it}\right)\right|\leq A_iB_i\left|\theta_{it}-\hat{\theta}_{it}\right|\leq A_iB_i\left(\left|\theta_{it}\right|+\left|\hat{\theta}_{it}\right|\right). \tag{33}$$

Hence, this term has a finite EPV under Condition F-BE and $||\hat{\theta}_i||<\infty$. For the second term, using (33) for $t+1$ and the Law of Iterated Expectations, the EPV of its absolute value is bounded by

$$A_iB_i\left(\sum_{\tau=0}^{t-1}\delta^{\tau+1}\mathbb{E}^{\lambda_i[\chi,\Gamma]|\theta^{\tau-1},\theta_{i\tau}}[\tilde{\theta}_{i\tau+1}]+\mathbb{E}^{\lambda_i[\chi,\Gamma]|\theta^{s-1},\theta_{is}}\left[||\tilde{\theta}_i||\right]+||\hat{\theta}_i||\right)$$

which is finite by Condition F-BE and $||\hat{\theta}_i||<\infty$. Finally, for the third term, note that this has a finite EPV because of Conditions U-SPR and F-BE.

For part (ii), write the time-s equilibrium expected payoff (16) given history θ^{s-1},θ_{is} using Fubini's Theorem and the Law of Iterated Expectations as

$$\begin{aligned}
V_{is}^{\langle\chi,\Psi\rangle,\Gamma}(\theta^{s-1},\theta_{is}) &= \lim_{T\to\infty}\sum_{t=0}^{T}\delta^t\mathbb{E}^{\lambda_i[\chi,\Gamma]|\theta^{s-1},\theta_{is}}\left[u_{it}(\tilde{\theta}^t,\chi^t(\tilde{\theta}^t))+\psi_{it}(\tilde{\theta}^t)\right]\\
&= \sum_{t=0}^{s-1}\delta^t\left(\mathbb{E}^{\Gamma_i\left(\theta_i^{s-1},\chi_i^{s-1}(\theta^{s-1})\right)}\left[u_{it}(\theta_i^t,\tilde{\theta}_{-i}^t,\chi^t(\theta_i^t,\tilde{\theta}_{-i}^t))\right]+\psi_{it}(\theta^t)\right)\\
&\quad +Q_{is}^{\chi,\Gamma}(\theta^{s-1},\theta_{is})-\lim_{T\to\infty}\mathbb{E}^{\lambda_i[\chi,\Gamma]|\theta^{s-1},\theta_{is}}\left[Q_{i,T+1}^{\chi,\Gamma}(\tilde{\theta}^T,\tilde{\theta}_{i,T+1})\right]
\end{aligned}$$

(the expectations of the other terms for $t\geq s$ cancel out by the Law of Iterated Expectations). Note that the second line is independent of θ_{is}, and that the last limit equals zero by (33), Condition F-BE, and $||\hat{\theta}_i||<\infty$. Finally, note that by (20) and (32), $Q_{is}^{\chi,\Gamma}\left(\theta^{s-1},\cdot\right)$ is Lipschitz continuous in θ_{is}, and its derivative equals a.e. $D_{is}^{\chi,\Gamma}\left(\theta^{s-1},\theta_{is}\right)$, which is the right-hand side of (19). ■

Lemma 5 *Suppose that the payment rule Ψ can be spread over time. Then, as $T\to\infty$, for $\lambda[\chi]$–almost all θ, $\mathbb{E}^{\lambda[\chi]|\theta^T}\left[\Psi(\tilde{\theta})\right]\to\Psi(\theta)$.*

Proof. Using the Law of Iterated Expectations,

$$\mathbb{E}_{\tilde{\theta}}^{\lambda[x]}\left[\left|\mathbb{E}_{\tilde{\tilde{\theta}}}^{\lambda[x]|\tilde{\theta}^T}\left[\Psi(\tilde{\tilde{\theta}})\right]-\Psi(\tilde{\theta})\right|\right] = \mathbb{E}_{\tilde{\theta}}^{\lambda[x]}\left[\left|\mathbb{E}_{\tilde{\tilde{\theta}}}^{\lambda[x]|\tilde{\theta}^T}\left[\sum_{t=T+1}^{\infty}\delta^t\psi_t(\tilde{\tilde{\theta}}^t)\right]-\sum_{t=T+1}^{\infty}\delta^t\psi_t(\tilde{\theta}^t)\right|\right]$$
$$\leq 2\mathbb{E}_{\tilde{\theta}}^{\lambda[x]}\left[\sum_{t=T+1}^{\infty}\delta^t\left|\psi_t(\tilde{\theta}^t)\right|\right].$$

Since $\mathbb{E}^{\lambda[x]|\theta_0}[||\psi(\tilde{\theta})||]<\infty$, the last expectation goes to zero as $T\to\infty$, which implies the result. ■

Proof of Lemma 4. Take an arbitrary choice rule $\langle\chi,\Psi\rangle$ that is independent of θ_{it}, and let Γ be an accompanying belief system. For a fixed θ^{t-1}, let $\tilde{\theta}_{-i}^{t-1}$ be drawn from $\Gamma\left(\theta_i^{t-1},\chi_i^{t-1}\left(\theta^{t-1}\right)\right)$, and then let each $\tilde{\theta}_{jt}$ be drawn from the distribution $F_{jt}(\theta_j^{t-1},\chi_j^{t-1}(\theta^{t-1}))$ for each $j\neq i$, to obtain θ_{-i}^t. Then we can calculate

$$V_{it}^{\langle\chi,\Psi\rangle,\Gamma}\left(\theta^{t-1},\theta_{it}\right)=\mathbb{E}_{\tilde{\theta}_{-i}^t}\mathbb{E}_{\tilde{\varepsilon}_{i,t+1}}\left[V_{i,t+1}^{\langle\chi,\Psi\rangle,\Gamma}\left(\tilde{\theta}_{-i}^t,\theta_i^t,Z_{i,t+1}\left(\theta_{it},\chi_i^t\left(\theta_i^{t-1},\tilde{\theta}_{-i}^t\right),\tilde{\varepsilon}_{i,t+1}\right)\right)\right]. \quad (34)$$

To differentiate this expression with respect to θ_{it}, use the chain rule and the assumption that χ does not depend on θ_{it} to write

$$\frac{dV_{i,t+1}^{\langle\chi,\Psi\rangle,\Gamma}\left(\theta^t,Z_{i,t+1}\left(\theta_{it},\chi_i^t\left(\theta^t\right),\varepsilon_{i,t+1}\right)\right)}{d\theta_{it}}=\frac{\partial V_{i,t+1}^{\langle\chi,\Psi\rangle,\Gamma}\left(\theta^t,Z_{i,t+1}\left(\theta_{it},\chi_i^t\left(\theta^t\right),\varepsilon_{i,t+1}\right)\right)}{\partial\theta_{it}} \quad (35)$$
$$+\frac{\partial V_{i,t+1}^{\langle\chi,\Psi\rangle,\Gamma}\left(\theta^t,Z_{i,t+1}\left(\theta_{it},\chi_i^t\left(\theta^t\right),\varepsilon_{i,t+1}\right)\right)}{\partial\theta_{i,t+1}}\cdot\frac{\partial Z_{i,t+1}\left(\theta_{it},\chi_i^t\left(\theta^t\right),\varepsilon_{i,t+1}\right)}{\partial\theta_{it}}.$$

For the first term in (35), note that since the environment is Markov in period $t+1$ and χ does not depend on θ_{it}, we have that

$$\frac{\partial V_{i,t+1}^{\langle\chi,\Psi\rangle,\Gamma}\left(\theta^t,\theta_{i,t+1}\right)}{\partial\theta_{it}}=\mathbb{E}^{\lambda_i[\chi,\Gamma]|\theta^t,\theta_{i,t+1}}\left[\frac{\partial U_i(\tilde{\theta},\chi(\tilde{\theta}))}{\partial\theta_{it}}\right].$$

As for the second term, by ICFOC$_{i,t+1}$, $\partial V_{i,t+1}^{\langle\chi,\Psi\rangle,\Gamma}\left(\theta^t,\theta_{i,t+1}\right)/\partial\theta_{i,t+1}=D_{i,t+1}^{\chi,\Gamma}(\theta^t,\theta_{i,t+1})$, whereas

$$\frac{\partial Z_{i,t+1}\left(\theta_{it},\chi_i^t\left(\theta^t\right),\varepsilon_{i,t+1}\right)}{\partial\theta_{it}}=I_{i,(t),t+1}(\theta_i^{t+1},\chi_i^t\left(\theta^t\right)).$$

Now, the Dominated Convergence Theorem allows us to differentiate (34) with respect to θ_{it} since all the functions being differentiated are all equi-Lipschitz continuous in $(\theta_{i,t+1},\theta_{-i}^t\ )$ by U-ELC, ICFOC$_{i,t+1}$ and F-BIR. Doing this differentiation and using (18) and the above chain formula yields

$$\frac{dV_{it}^{\langle\chi,\Psi\rangle,\Gamma}\left(\theta^{t-1},\theta_{it}\right)}{d\theta_{it}}$$
$$=\mathbb{E}^{\lambda_i[\chi,\Gamma]|\theta^{t-1},\theta_{it}}\left[\frac{\partial U_i(\tilde{\theta},\chi(\tilde{\theta}))}{\partial\theta_{it}}+I_{i,(t),t+1}\left(\tilde{\theta}_i^{t+1},\chi_i^t(\tilde{\theta})\right)\cdot\sum_{\tau=t+1}^{\infty}I_{i,(t+1),\tau}\left(\tilde{\theta}_i^{\tau},\chi_i^{\tau-1}(\tilde{\theta})\right)\frac{\partial U_i(\tilde{\theta},\chi(\tilde{\theta}))}{\partial\theta_{i\tau}}\right].$$

Now, Markovness in period $t+1$ and the chain rule expression for impulse responses (7) imply that $I_{i,(t),t+1}\left(\theta_i^{t+1},x_i^t\right)\cdot I_{i,(t+1),\tau}\left(\theta_i^{\tau},x_i^{\tau-1}\right)=I_{i,(t),\tau}\left(\theta_i^{\tau},x_i^{\tau-1}\right)$. Using also the fact that $I_{i,(t),t}=1$, the above formula yields ICFOC$_{i,t}$. ■

Proof of Proposition 5. The usual backward-induction argument establishes that, if no one-step deviation from strong truthtelling is profitable, then an agent's expected payoff at any history can never be increased by any *finite-stage* deviation from strong truthtelling. To establish that infinite-stage deviations are also unprofitable, suppose in negation that there exists an agent i and a period-s history $h_{is}=\left(\theta_i^s,\hat{\theta}_i^{s-1},x_i^{s-1}\right)$ (not necessarily truthful for agent i) such that agent i raises his expected payoff at h_{is} by some $\varepsilon>0$ by deviating from strong truthtelling to some strategy $\hat{\sigma}_i$, assuming that all other agents report truthfully. (All the expectations below will be given history h_{is}.)

We then show that there exists some finite $T>s$ such that reversion from $\hat{\sigma}_i$ to strong truthtelling starting in period T cannot reduce the agent's expected payoff by more than $\varepsilon/2$. For this purpose, note first that under the assumption of part (b) of Lemma 2, the agent's time-s expectation of his type in each period t under *any* strategy is bounded as follows (by analogy to (9), and letting $I_{i,t-\tau}=0$ for $\tau>t$, where I denote the impulse responses of the AR(k) process that bounds F, in the sense of Lemma 2):

$$
\begin{aligned}
\mathbb{E}[|\tilde{\theta}_{it}|] &\leq \textstyle\sum_{\tau=0}^{s} I_{i,t-\tau}\left|\theta_{i\tau}\right|+\sum_{\tau=s+1}^{t} I_{i,t-\tau}E_{i\tau}\equiv S_{it}, \text{ where}\\
\|S_i\| &\leq \|I_i\|\cdot\left(\textstyle\sum_{\tau=0}^{s}\delta^{\tau}\left|\theta_{i\tau}\right|+\sum_{\tau=s+1}^{\infty}\delta^{\tau}\left|E_{i\tau}\right|\right)\leq\|I_i\|\cdot\left(\textstyle\sum_{\tau=0}^{s}\delta^{\tau}\left|\theta_{i\tau}\right|+\|E_i\|\right)<\infty.
\end{aligned}
$$

Hence, by Condition U-SPR, the EPV of non-monetary utility flows starting in period T under any strategy is absolutely bounded by $\sum_{t=T}^{\infty}\delta^t\left[L_iS_{it}+M_{it}\right]$, and so reversion from strategy $\hat{\sigma}_i$ to strong truthtelling in period T can reduce this EPV by at most twice this amount. Likewise, condition (ii) on payments implies that the EPV of payment flows from strong truthtelling starting in period T is absolutely bounded by $\sum_{t=T}^{\infty}\delta^t\left[L_i^{\psi}S_{it}+M_{it}^{\psi}\right]$. As for the EPV of payment flows from period T onwards under strategy $\hat{\sigma}_i$, by condition (i) it is bounded above by $\sum_{t=T}^{\infty}\delta^tK_{it}$. Adding up, we see that the total expected loss due to the reversion in period T does not exceed

$$
\sum_{t=T}^{\infty}\delta^t\left[\left(2L_i+L_i^{\psi}\right)S_{it}+2M_{it}+M_{it}^{\psi}+K_{it}\right].
$$

Since $L_i,L_i^{\psi},\|S_i\|,\|M_i\|,||M_i^{\psi}||,\|K_i\|<\infty$, the sum goes to zero as $T\to\infty$, and so falls below $\varepsilon/2$ for some T large enough. But then the finite-stage deviation to $\hat{\sigma}_i$ between periods s and T is profitable, which contradicts the first statement of the proof. ■

Proof of Proposition 6. We prove the result by showing that the payment rule given by (21) satisfies conditions (i) and (ii) of Proposition 5. This is done by establishing the bounds for each of the three terms in (21).

We begin with condition (ii). The first term in (21) satisfies this bound by (33) and $||\hat{\theta}_i|| < \infty$. For the second term, using (33) and Condition F-MBE,

$$\begin{aligned}\left|\delta^{-t}\mathbb{E}^{\lambda_i[\chi,\Gamma]|\theta^{t-1},\theta_{it}}\left[Q^{\chi,\Gamma}_{i,t+1}(\tilde{\theta}^t,\tilde{\theta}_{i,t+1})\right]\right| &\leq A_iB_i\left(\mathbb{E}^{F_{i,t+1}(\theta_{it},\chi_i^{t-1}(\theta^{t-1}))}\left[\left|\tilde{\theta}_{i,t+1}\right|\right]+\left|\hat{\theta}_{i,t+1}\right|\right) \\ &\leq A_iB_i\left(\phi_i\left|\theta_{it}\right|+E_{it}+\left|\hat{\theta}_{i,t+1}\right|\right),\end{aligned} \tag{36}$$

which satisfies the bound since $||E_i||, ||\hat{\theta}_i|| < \infty$. The third term satisfies the bound by Condition U-MSPR.

Next, we turn to condition (i). Using (20) and the fact that the payments (21) satisfy ICFOC (by Proposition 4), we can write

$$\begin{aligned}Q^{\chi,\Gamma}_{it}\left(\theta^{t-1},\theta_{it}\right) &= V^{\langle\chi,\Psi\rangle,\Gamma}_{it}\left(\theta^{t-1},\theta_{it}\right)-V^{\langle\chi,\Psi\rangle,\Gamma}_{it}(\theta^{t-1},\hat{\theta}_{it}), \\ Q^{\chi,\Gamma}_{i,t+1}\left(\theta^{t},\theta_{i,t+1}\right) &= V^{\langle\chi,\Psi\rangle,\Gamma}_{i,t+1}\left(\theta^{t},\theta_{i,t+1}\right)-V^{\langle\chi,\Psi\rangle,\Gamma}_{i,t+1}(\theta^{t},\hat{\theta}_{i,t+1}).\end{aligned}$$

Substituting in (21), omitting superscripts on V and Q to save space, and noting that

$$\mathbb{E}^{\lambda_i[\chi,\Gamma]|\theta^{t-1},\theta_{it}}\left[V_{i,t+1}\left(\tilde{\theta}^t,\tilde{\theta}_{i,t+1}\right)\right]=V_{it}\left(\theta^{t-1},\theta_{it}\right),$$

we can write

$$\begin{aligned}\psi_{it}\left(\theta^t\right) = {} & \delta^{-t}\left(\mathbb{E}^{\lambda_i[\chi,\Gamma]|\theta^{t-1},\theta_{it}}\left[V_{i,t+1}(\theta^t_i,\tilde{\theta}^t_{-i},\hat{\theta}_{i,t+1})\right]-V_{it}(\theta^{t-1},\hat{\theta}_{it})\right) \\ & -\mathbb{E}^{\lambda_i[\chi,\Gamma]|\theta^{t-1},\theta_{it}}\left[u_{it}\left(\theta_{it},\tilde{\theta}^t_{-i},\chi^t(\tilde{\theta}^t)\right)\right].\end{aligned} \tag{37}$$

Now, the assumption that no agent has a profitable one-step deviation from strong truthtelling implies in particular that, given true history $(\theta^{t-1},\hat{\theta}_{it})$, agent i does not gain by reporting θ_{it}, followed by strong truthtelling. For the Markov environment, this can be written as follows:

$$V_{it}(\theta^{t-1},\hat{\theta}_{it}) \geq \mathbb{E}^{\lambda_i[\chi,\Gamma]|\theta^{t-1},\hat{\theta}_{it}}\left[\begin{array}{c}V_{i,t+1}\left(\theta^t_i,\tilde{\theta}^t_{-i},\tilde{\theta}_{i,t+1}\right)+\delta^t u_{it}(\hat{\theta}_{it},\tilde{\theta}^t_{-i},\chi^t(\theta^t_i,\tilde{\theta}^t_{-i})) \\ -\delta^t u_{it}(\theta_{it},\tilde{\theta}^t_{-i},\chi^t(\theta^t_i,\tilde{\theta}^t_{-i}))\end{array}\right]$$

(since in the Markov environment, once $\tilde{\theta}_{i,t+1}$ is realized following the misreport, the only payoff effect of the past true type being $\hat{\theta}_{it}$ rather than θ_{it} is through period-t utility). Using this inequality, and also noting that the distribution of $\tilde{\theta}^t_{-i}$ in probability measure $\lambda_i\left[\chi,\Gamma\right]|\theta^{t-1},\theta_{it}$ does not depend on θ_{it}, so we can replace the measure in (37) with $\lambda_i\left[\chi,\Gamma\right]|\theta^{t-1},\hat{\theta}_{it}$, the payments (37) are bounded above as follows:

$$\begin{aligned}\psi_{it}\left(\theta^t\right) \leq {} & \delta^{-t}\mathbb{E}^{\lambda_i[\chi,\Gamma]|\theta^{t-1},\hat{\theta}_{it}}\left[V_{i,t+1}(\theta^t_i,\tilde{\theta}^t_{-i},\hat{\theta}_{i,t+1})-V_{i,t+1}(\theta^t_i,\tilde{\theta}^t_{-i},\tilde{\theta}_{i,t+1})\right] \\ & -\mathbb{E}^{\lambda_i[\chi,\Gamma]|\theta^{t-1},\hat{\theta}_{it}}\left[u_{it}(\hat{\theta}_{it},\tilde{\theta}^t_{-i},\chi^t(\theta^t_i,\tilde{\theta}^t_{-i}))\right].\end{aligned}$$

Finally, use ICFOC$_{t+1}$ and (20) to rewrite the bound as

$$\psi_{it}\left(\theta^t\right) \leq -\delta^{-t}\mathbb{E}^{\lambda_i[\chi,\Gamma]|\theta^{t-1},\hat{\theta}_{it}}\left[Q_{i,t+1}(\theta^t_i,\tilde{\theta}^t_{-i},\tilde{\theta}_{i,t+1})\right]-\mathbb{E}^{\lambda_i[\chi,\Gamma]|\theta^{t-1},\hat{\theta}_{it}}\left[u_{it}(\hat{\theta}_{it},\tilde{\theta}^t_{-i},\chi^t(\theta^t_i,\tilde{\theta}^t_{-i}))\right],$$

and note that both terms are bounded above by a finite-norm sequence: the first term by (36) and $||E_i||, ||\hat{\theta}_i|| < \infty$, and the second term by Condition U-MSPR. ■

Proof of Proposition 8.

Case (a): We construct a nondecreasing solution $\chi_s(\theta)$ sequentially for $s = 0, 1....$ Suppose we have a solution χ in which $\chi^{s-1}(\theta)$ is nondecreasing. Consider the problem of choosing the optimal continuation allocation rule in period s given type history θ^s and allocation history $\chi^{s-1}(\theta^{s-1})$. Using the state representation (17) from period s onward, we can write the continuation rule for $t \geq s$ as a collection of functions $\hat{\chi}_t(\varepsilon)$ of the shocks ε.

First, note that, because X is a sublattice, $\prod_{t\geq s} X_t$ is a lattice. This means that the set of feasible shock-contingent plans $\hat{\chi}$ is also a lattice under pointwise meet and join operations (i.e. for each ε). Next, note that, under the assumptions in the proposition, each agent i's virtual utility

$$U_i\left(Z_s(\theta^s,\varepsilon), \chi^{s-1}\left(\theta^{s-1}\right), x^{\geq s}\right) - \frac{1}{\eta_{i0}(\theta_{i0})}\sum_{t=0}^{\infty}\frac{\partial U_i\left(Z_s(\theta^s,\varepsilon), \chi^{s-1}\left(\theta^{s-1}\right), x^{\geq s}\right)}{\partial \theta_{it}} I_{it}(Z_{i,(s)}^t(\theta^s,\varepsilon))$$

is supermodular in $x^{\geq s}$ and has ID in $\left(\theta^s, x^{\geq s}\right)$ (observe that $Z_s(\theta^s,\varepsilon)$ is nondecreasing in θ^s under F-FOSD, and $\chi^{s-1}\left(\theta^{s-1}\right)$ is nondecreasing in θ^{s-1} by construction). Therefore, adding up over agents and taking expectation over ε, we obtain that the expected virtual surplus starting with period-s history θ^s is supermodular in the continuation plan $\hat{\chi}$ and has ID in $(\theta^s, \hat{\chi})$. Topkis's Theorem then implies that the set of optimal continuation plans is nondecreasing in θ^s in the strong set order. In particular, focus on the first component $\chi_s \in X_s$ of such plans. By Theorem 2 of Kukushkin (2009), there exists a nondecreasing selection of optimal values, $\hat{\chi}_s(\theta^s)$. Therefore, the relaxed program admits a solution in which $\chi^s(\theta^s) = \left(\chi^{s-1}\left(\theta^{s-1}\right), \hat{\chi}_s(\theta^s)\right)$ is nondecreasing in θ^s.

Case (b): In this case, the solution to the relaxed problem is a collection of independent rules $\chi_t\left(\theta^t\right)$, one for each t, with each $\chi_t\left(\theta^t\right)$ satisfying the optimality condition (26) for all θ^t. For any t, we can then choose $\chi_t\left(\theta^t\right)$ to depend only on $\left(\varphi_{1t}\left(\theta_1^t\right), ..., \varphi_{Nt}\left(\theta_N^t\right)\right)$: $\chi_t\left(\theta^t\right) = \bar{\chi}_t\left(\varphi_{1t}\left(\theta_1^t\right), ..., \varphi_{Nt}\left(\theta_N^t\right)\right)$. Now fix any $i \geq 1$ and for any $x_{it} \in X_{it}$, let $X_t(x_{it}) \equiv \{x_t' \in X_t : \ x_{it}' = x_{it}\}$. The optimality condition (26) implies that

$$\bar{\chi}_{it}(\varphi_t) \in \arg\max_{x_{it}\in X_{it}} \left[\bar{u}_{it}(\varphi_{it}, x_{it}) + g_{it}\left(\varphi_{-i,t}, x_{it}\right)\right]$$

where $\bar{u}_{it}(\phi_{it}, x_{it})$ is the virtual utility of agent i, and

$$g_{it}\left(\varphi_{-i,t}, x_{it}\right) \equiv \max_{x_t'\in X_t(x_{it})}\left[U_{0t}\left(x_t'\right) + \sum_{j\neq i}\bar{u}_{jt}\left(\varphi_{jt}, x_{jt}'\right)\right]$$

Since $\bar{u}_{it}(\varphi_{it}, x_{it}) + g_{it}\left(\varphi_{-i,t}, x_{it}\right)$ has strict ID in (φ_{it}, x_{it}), by the Monotone Selection Theorem of Milgrom and Shannon (1994), $\bar{\chi}_{it}\left(\varphi_{it}, \varphi_{-i,t}\right)$ must be nondecreasing in φ_{it}, and so $\chi_{it}\left(\theta_i^t, \theta_{-i}^t\right)$ is nondecreasing in θ_i^t. ■

Proof of Proposition 9. Optimality of virtual index policies in the bandit problem is well known (see Whittle 1982). Assume then that the initial type distributions have monotone hazard rates. Let χ be a virtual index policy, and let Γ denote beliefs where at each history, each bidder i assigns probability 1 to bidders $-i$ having been truthful. As noted in the text, it suffices to verify weak monotonicity, i.e., that for all i, s, θ^{s-1}, θ_{is},

$$\mathbb{E}^{\lambda_i[\chi\circ\hat{\theta}_s,\Gamma]|\theta^{s-1},\theta_{is}}\left[\sum_{t=s}^{\infty}\delta^t(\chi\circ\hat{\theta}_{is})_{it}(\tilde{\theta})\right]$$

is nondecreasing in $\hat{\theta}_{is}$. We show this for $s=0$. The argument for $s>0$ is analogous but simpler since $\hat{\theta}_{is}$ does not affect bidder i's handicap when $s>0$.

We can think of the processes being generated as follows: First, draw a sequence of innovations $\omega_i=(\omega_{ik})_{k=1}^{\infty}$ according to $\prod_{k=1}^{\infty}R_i(\cdot|k)$ for each i, independently across i, and draw initial types θ_{i0} according to F_{i0} independently of the innovations ω_i and across i. Letting $K_t\equiv\sum_{\tau=1}^{t}x_\tau$, bidder i's type in period t is then

$$\theta_{it}=\theta_{i0}+\sum_{k=1}^{K_t}\omega_{ik}.$$

It is straightforward to verify that this generates the same conditional distributions (and hence the same process) as the kernels defined in the main text.

Fix bidder i, a realization $(\theta_0,\omega)\in\Theta_0\times(\mathbb{R}^N)^{\infty}$, and $\theta'_{i0},\theta''_{i0}\in\Theta_0$ with $\theta''_{i0}>\theta'_{i0}$. We show by induction on k that, for any $k\in\mathbb{N}$, the k-th time that i wins the object if he initially reports θ''_{i0} (and reports truthfully in periods $t>0$) comes weakly earlier than if he reports θ'_{i0}. As the realization $(\theta_0,\omega)\in\Theta_0\times(\mathbb{R}^N)^{\infty}$ is arbitrary, this implies that the expected time to the k-th win is decreasing in the initial report, which in turn implies weak monotonicity.

As a preliminary observation, note that the period-t virtual index of bidder i is increasing in the (reported) period-0 type θ_{i0} since the handicap is decreasing in θ_{i0}, and (in case $t=0$) $\mathbb{E}^{\lambda[\bar{\chi}_i]|\theta_{i0}}[\theta_{i\tau}]$ is increasing in θ_{i0} for all $\tau\geq 0$.

Base case: Suppose towards contradiction that the first win given initial report θ'_{i0} comes in period t' whereas it comes in period $t''>t'$ given report θ''_{i0}. As the realization (θ_0,ω) is fixed, the virtual indices of bidders $-i$ in period t' are the same in both cases. But $\gamma_{it'}((\theta''_{i0},\theta_{i0},\ldots,\theta_{i0}),0)>\gamma_{it'}((\theta'_{i0},\theta_{i0},\ldots,\theta_{i0}),0)$ so that then i must win in period t' also with initial report θ''_{i0}, which contradicts $t''>t'$.

Induction step: Suppose the claim is true for some $k\geq 1$. Suppose towards contradiction that the $k+1$-th win given report θ'_{i0} comes in period t' whereas it comes in period $t''>t'$ given θ''_{i0}. We have the following observations: (i) In both cases i wins the auction $k-1$ times in periods $t<t'$. Furthermore, since the realization (θ_0,ω) is fixed, this implies that (ii) bidder i's current type θ_{it} is the same in both cases, and (iii) each bidder $j\neq i$ wins the object in periods $t<t'$ equally many

times in both cases, and hence the virtual indices of bidders $-i$ in period t' are the same in both cases. By (i) and (ii) i's virtual index in period t' is identical in both cases save for the initial report. As the index is increasing in the initial report, (iii) implies that i must then win in period t' also with initial report θ''_{i0}, contradicting $t'' > t'$. Hence the claim is true for $k+1$. ■

Proof of Proposition 10. Without loss of generality, normalize the seller's types to zero by letting each agent i's type θ_{it} denote the flow utility that agent i obtains, net to the seller's value. Inductively for $s = 0, 1...$, we construct a solution to the relaxed program in which, for each agent $i \geq 1$, (a) $\chi_i^s(\theta^s)$ is nondecreasing in θ_i^s , and (b) $\chi_0^s(\theta^s)$ depends only on θ_{-i}^s and $\chi_i^s\left(\theta_i^s, \theta_{-i}^s\right)$ – i.e., whether the object is sold depends on agent i's types only through his own allocation. For this purpose, using the state representation from period s onward, write the future types of each agent i as a function $Z_{i,(s),t}(\theta_i^s, \varepsilon_i)$ of shocks ε_i, which is strictly increasing given the assumption of positive impulse responses. Combined, the various assumptions imply that, for any $t \geq s$, any $i \geq 1$, the virtual value

$$\nu_{i,(s),t}(\theta_i^s, \varepsilon_i) \equiv Z_{i,(s),t}(\theta_i^s, \varepsilon_i) - \frac{1}{\eta_{i0}(\theta_{i0})} \prod_{\tau=1}^{t} I_{i,(\tau-1),\tau}(Z_{i,(s)}^t(\theta^s, \varepsilon_i))$$

of having agent i consuming the good in period t are strictly increasing in θ_i^s.

When the good is not previously sold (i.e., $\chi_{0,s-1}\left(\theta^{s-1}\right) = 1$), we allocate the good optimally in period s using the principle of dynamic programming, i.e., assuming that, in case the good is not allocated in period s, then an optimal feasible shock-contingent plan $\hat{\chi}$ is followed from period $s+1$ onward. Also, when indifferent among optimal allocations, we select an allocation that does not sell the good (i.e., $x_{0t} = 1$) whenever possible. That is, we let

$$\begin{aligned} \chi_s(\theta^s) &\in \arg\max_{x_s \in X_s^*(\theta^s)} x_{0s}, \text{ where} \\ X_s^*(\theta^s) &\equiv \arg\max_{x_s \in \{0,1\}^{N+1}: \sum_{i=0}^N x_{is}=1} \left\{ \begin{array}{c} \sum_{i=1}^{N} \sum_{t=s}^{\infty} \delta^t \mathbb{E}_\varepsilon \left[\nu_{i,(s),t}(\theta_i^s, \tilde{\varepsilon}_i)\right] x_{is} \\ +x_{0s} \cdot \sup_{\hat{\chi}} \mathbb{E}_\varepsilon \left[\sum_{t=s+1}^{\infty} \sum_{i=1}^{N} \delta^t \nu_{i,(s),t}(\theta_i^s, \tilde{\varepsilon}_i) \hat{\chi}_t(\varepsilon) \right] \end{array} \right\} \end{aligned}$$

This selection ensures that property (b) extends from period $s-1$ to period s.

Now to verify property (a), fix $i \geq 1$ and consider an increase from θ_i^s to $\hat{\theta}_i^s \geq \theta_i^s$, holding θ_{-i}^s fixed. One of the following must hold:

(i) $\chi_{0,s-1}(\hat{\theta}_i^{s-1}, \theta_{-i}^{s-1}) = \chi_{0,s-1}\left(\theta_i^{s-1}, \theta_{-i}^{s-1}\right) = 0$. In this case, $\chi_{is}(\hat{\theta}_i^s, \theta_{-i}^s) = \chi_{is-1}(\hat{\theta}_i^{s-1}, \theta_{-i}^{s-1}) \geq \chi_{s-1}\left(\theta_i^{s-1}, \theta_{-i}^{s-1}\right) = \chi_s\left(\theta_i^s, \theta_{-i}^s\right)$.

(ii) $\chi_{0,s-1}(\hat{\theta}_i^{s-1}, \theta_{-i}^{s-1}) = \chi_{0,s-1}\left(\theta_i^{s-1}, \theta_{-i}^{s-1}\right) = 1$. In this case, the allocation is given by the above maximization program. Note that for any $\hat{\chi}$, the expected benefit to selling to agent $i \geq 1$ rather than either selling to someone else or waiting is strictly increasing in θ_i^s. Thus, by the aggregation method and the monotone selection theorem, we have $\chi_{is}(\hat{\theta}_i^s, \theta_{-i}^s) \geq \chi_{is}\left(\theta_i^s, \theta_{-i}^s\right)$.

(iii) $\chi_{0,s-1}(\hat{\theta}_i^{s-1}, \theta_{-i}^{s-1}) = 0$ while $\chi_{0,s-1}\left(\theta_i^{s-1}, \theta_{-i}^{s-1}\right) = 1$. In this case, by properties (a),(b) in period $s-1$, we must have $\chi_{is}(\hat{\theta}_i^s, \theta_{-i}^s) = \chi_{i,s-1}(\hat{\theta}_i^{s-1}, \theta_{-i}^{s-1}) = 1 \geq \chi_{is}\left(\theta_i^s, \theta_{-i}^s\right)$.

(iv) $\chi_{0,s-1}(\hat{\theta}_i^{s-1}, \theta_{-i}^{s-1}) = 1$ while $\chi_{0,s-1}\left(\theta_i^{s-1}, \theta_{-i}^{s-1}\right) = 0$. This case is inconsistent with properties (a),(b) in period $s-1$ and feasibility.

This establishes the result. ■

References

[1] Albanesi, S. and C. Sleet, (2006), "Dynamic Optimal Taxation with Private Information," *Review of Economic Studies*, 73(1), 1-30.

[2] Angus, J. (1994), "The Probability Integral Transform and Related Results," *SIAM Review*, 36(4), 652-654.

[3] Athey, S., and I. Segal (2007), "An Efficient Dynamic Mechanism," Mimeo, Stanford University.

[4] Atkeson, A. and R. E. Jr Lucas, (1992), "On efficient Distribution with Private Information," *Review of Economic Studies*, 59(3):427-453.

[5] Baron, D. and D. Besanko, (1984), "Regulation and Information in a Continuing Relationship," *Information Economics and Policy*, 1 447-470.

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

[7] Battaglini, M. (2007), "Optimality and Renegotiation in Dynamic Contracting," *Games and Economic Behavior*, 60(2), 213-246.

[8] Battaglini, M. and S. Coate (2008), "Pareto Efficient Income Taxation with Stochastic Abilities," *Journal of Public Economics*, 92(3-4), 844-868.

[9] Besanko, David (1985) "Multiperiod Contracts Between Principal and Agent with Adverse Selection," *Economics Letters*, 17, 33-37.

[10] Bergemann, D., and M. Said, (2011), "Dynamic Auctions," *Wiley Encyclopedia of Operations Research and Management Science.*

[11] Bergemann, D., and J. Välimäki, (2008), "Bandit Problems," New Palgrave Dictionary of Economics, Durlauf and Blume Editors.

[12] Bergemann, D., and J. Välimäki, (2010), "The Dynamic Pivot Mechanism," *Econometrica*, 78(2), 771-789.

[13] 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.

[14] Billingsley, P. (1995), *Probability and Measure*, 3rd ed., Wiley.

[15] Board, S., (2007), "Selling Options," *Journal of Economic Theory* 136, 324-340.

[16] Board, S. (2008), "Durable-Goods Monopoly with Varying Demand," *Review of Economic Studies*, 75(2) 391-413.

[17] Board, S. and A. Skrzypacz, (2010), "Revenue Management with Forward-Looking Buyers," mimeo UCLA and Stanford.

[18] Calzolari, G. and A. Pavan, (2006a), "On the Optimality of Privacy in Sequential Contracting," *Journal of Economic Theory*, Volume 130-1.

[19] Calzolari, G. and A. Pavan, (2006b), "Monopoly with Resale," *Rand Journal of Economics*, Volume 37-2.

[20] Cremer, J.and R. P. McLean, (1988), "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions," *Econometrica*, 56(6):1247-1257.

[21] Courty, P. and H. Li, (2000), "Sequential Screening," *Review of Economic Studies*, 67, 697-717.

[22] Dizdar, D., Gershkov, A. and B. Moldovanu (2011), "Revenue Maximization in the Dynamic Knapsack Problem," *Theoretical Economics*, 6(2) 157-184.

[23] Dudley, R. (2002), *Real Analysis and Probability*, Cambridge University Press.

[24] Edmans, A. and X. Gabaix (2010), "Tractability in Incentive Contracting," *Review of Financial Studies*, forthcoming.

[25] Edmans, A. X. Gabaix, T. Sadzik, and Y. Sannikov, (2010), "Dynamic Incentive Accounts," Mimeo NYU Stern, Princeton and Wharton.

[26] Eso, P. and B. Szentes, (2007), "Optimal Information Disclosure in Auctions and the Handicap Auction," *Review of Economic Studies*, 74(3), 705-731.

[27] Farhi, E. and I. Werning, (2010), "Insurance and Taxation over the Life Cycle," Mimeo, Harvard and MIT.

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

[29] Fudenberg, D. and J. Tirole, (1991), *Game Theory*, MIT Press.

[30] Galperti, S., (2011), "Commitment, Flexibility, and Optimal Screening of Time-Inconsistency," Mimeo, Northwestern University.

[31] Garcia, D. (2005), "Monotonicity in Direct Revelation Mechanisms," *Economics Letters*, 88(1), 21.26.

[32] Garrett, D. (2011), "Durable-good Sales with Changing Values and Unobservable Arrival Dates," Mimeo, Northwestern University.

[33] Garrett, D. and A. Pavan (2011a), "Managerial Turnout in a Changing World," mimeo Northwestern University.

[34] Garrett, D. and A. Pavan (2011b), "Dynamic Managerial Compensation: On the Optimality of Seniority-Based Schemes," mimeo Northwestern University.

[35] Gershkov, A. and B. Moldovanu, (2009a), "Dynamic Revenue Maximization with Heterogenous Objects: A mechanism Design Approach," *American Economic Journal: Microeconomics*, 1(2), 168-198.

[36] Gershkov, A. and B. Moldovanu, (2009b), "Learning About the Future and Dynamic Efficiency," *American Economic Review*, 99(4), 1576-1587.

[37] Gershkov, A. and B. Moldovanu, (2009c), "Optimal Search, Learning and Implementation," Mimeo, University of Bonn.

[38] Gershkov, A. and B. Moldovanu, (2010a), "Efficient Sequential Assignment with Incomplete Information," *Games and Economic Behavior*, 68(1), 144-154.

[39] Gershkov, A. and B. Moldovanu, (2010b), "Demand Uncertainty and Dynamic Allocation with Strategic Agents," Mimeo University of Bonn.

[40] Gershkov, A. and B. Moldovanu, (2010c), "Non-Bayesian Optimal Search and Dynamic Implementation," Mimeo University of Bonn.

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

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

[43] Golosov, M. M. Troshkin, and A. Tsyvinski, (2010), "Optimal Dynamic Taxes," Mimeo Yale University.

[44] Green, E. J., (1987), "Lending and the Smoothing of Uninsurable Income," in Prescott E. and Wallace N., editors, *Contractual Agreements for Intertemporal Trade*, University of Minnesota Press.

[45] Hendel, I. and A. Lizzeri, (2003)., "The Role Of Commitment In Dynamic Contracts: Evidence From Life Insurance," *Quarterly Journal of Economics*,118(1), 299-327.

[46] Kakade, S., Lobel, I. and H. Nazerzadeh, (2011), "Optimal Dynamic Mechanism Design and the Virtual Pivot Mechanism," mimeo Wharton School, University of Pennsylvania.

[47] Kapicka, M. (2010), "Efficient Allocations in Dynamic Private Information Economies with Persistent Shocks: A First Order Approach," Mimeo, University of California, Santa Barbara.

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

[49] Kukushkin, N.S. (2009), "On the Existence of Monotone Selections," working paper.

[50] Laffont, J.-J. and J. Tirole (1988), "The Dynamics of Incentive Contracts," *Econometrica*, 56(5), 1153-1175.

[51] Laffont, J.-J. and J. Tirole (1990), "Adverse Selection and Renegotiation in Procurement," *Review of Economic Studies*, 57, 597-625.

[52] Milgrom, P., and I. Segal (2002), "Envelope Theorems for Arbitrary Choice Sets," *Econometrica*, 70(2), 583.601.

[53] Milgrom, P., and C. Shannon (1994), "Monotone Comparative Statics," *Econometrica*, 62(1), 157-180.

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

[55] Pai, M. and R. Vohra (2008), "Optimal Dynamic Auctions," Mimeo, Northwestern University.

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

[57] Pollard, D. (2002), *A User's Guide to Measure Theoretic Probability*, Cambridge University Press.

[58] Rahman, D. (2010), "Dynamic Implementation," mimeo, University of Minnesota.

[59] Riordan, M., and D. Sappington (1987), "Awarding Monopoly Franchises," *American Economic Review*, 77(3),375-387.

[60] Rochet, J. C. (1987), "A Necessary and Sufficient Condition for Rationalizability in a Quasi-Linear Context," *Journal of Mathematical Economics*, 16, 191–200.

[61] Rochet, J-C., and L. Stole, (2003), "The Economics of Multidimensional Screening," in M. Dewatripont, L.P. Hansen, and S. Turnovsky, editors, *Advances in Economics and Econometrics: Theory and Applications - Eight World Congress of Econometric Society*, 150-197. Cambridge University Press.

[62] Said, M. (2011), "Auctions with Dynamic Populations: Efficiency and Revenue Maximization," Mimeo, Washington University in St. Louis.

[63] Skreta, V. (2006), "Sequentially Optimal Mechanisms," *Review of Economic Studies*, 73(4), 1085-1111.

[64] Strausz, R. (2006), "Deterministic versus Stochastic Mechanisms in Principal-Agent Models," *Journal of Economic Theory*, 127(1), 306-314.

[65] Strulovici, B., (2011), "Renegotiation-proof Contracts with Moral Hazard and Persistent Private Information," Mimeo, Northwestern University.

[66] Thomas, J. and T. Worrall (1990), "Income Fluctuation and Asymmetric Information: An Example of a Repeated Principal-Agent Problem," *Journal of Economic Theory*, 51, 367-390.

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

[68] Topkis, D. (1998), *Supermodularity and Complementarity*, Princeton University Press.

[69] Townsend, R. (1982), "Optimal Multiperiod Contracts and the Gain from Enduring Relationships under Private Information," *Journal of Political Economy* 90, 1166-86.

[70] Zhang, Y. (2009), "Dynamic Contracting, Persistent Shocks and Optimal Taxation," *Journal of Economic Theory*, 144(2), 635-675.

[71] Whittle, P. (1982), *Optimization Over Time, vol. 1*, Wiley, Chichester.

[72] Williams, N. (2011), "Persistent Private Information," *Econometrica*, forthcoming.