

Optimal Dynamic Mechanism Design and the Virtual Pivot Mechanism

Sham M. Kakade* Ilan Lobel† Hamid Nazerzadeh‡§

March 25, 2011

Abstract

We consider the problem of designing optimal mechanisms for settings where agents have dynamic private information. We present the VIRTUAL-PIVOT Mechanism, that is optimal in a large class of environments that satisfy a separability condition. The mechanism satisfies a rather strong equilibrium notion (it is periodic ex-post incentive compatible and individually rational). We provide both necessary and sufficient conditions for immediate incentive compatibility for mechanisms that satisfy periodic ex-post incentive compatibility in future periods. The result also yields a strikingly simple mechanism for selling a sequence of items to a single buyer. We also show the allocation rule of the VIRTUAL-PIVOT Mechanism has a very simple structure (a Virtual Index) in multi-armed bandit settings. Finally, we show through examples that the relaxation technique we use does not produce optimal dynamic mechanisms in general non-separable environments.

1 Introduction

We study the problem of designing optimal mechanisms for environments with dynamic private information and propose a mechanism that is profit-maximizing in a class of environments that we call separable.

In a separable environment, the valuation function of an agent can be decomposed as the product (or sum) of a function of the agent's first signal and another function of the agent's future signals. As an example, consider a manufacturer that sells to several retailers. Each retailer has two pieces of private information: her profit margin (which she knows a priori) and the demand she faces (which she learns over time and could potentially be time-varying).

Another compelling application is that of online advertisement auctions where a publisher sells the space on her website to advertisers. The advertisers usually know a-priori their profit margins on

*Department of Statistics, Wharton School, University of Pennsylvania – skakade@wharton.upenn.edu. Work performed in part at Microsoft Research.

†Stern School of Business, New York University – ilobel@stern.nyu.edu. Work performed in part at Microsoft Research.

‡Microsoft Research, New England Lab – hamidnz@microsoft.com

§We would like to thank Maher Said, Gregory Lewis, Daron Acemoglu, Susan Athey, Markus Mobius, Mallesh Pai, Andrzej Skrzypacz, and Rakesh Vohra for many insightful suggestions and helpful comments.

each sale, but they estimate over time the conversion rates (fraction of ads that turn into sales). This is also a separable environment.

The optimal mechanism we propose, the VIRTUAL-PIVOT Mechanism, is quite intuitive — it combines ideas based on the “virtual value” formulation of Myerson [1981] for static revenue-optimal mechanism design and the dynamic “pivot” mechanism proposed by Bergemann and Välimäki [2010] for maximizing social welfare. The mechanism essentially maximizes an affine transformation of the social welfare which corresponds to a certain virtual surplus. Furthermore, the mechanism satisfies strong (periodic ex-post) notions of incentive compatibility and individual rationality.

One notable special case of our results is the setting with only one buyer. Namely, consider a setting where the mechanism at each period has one item to sell to a single buyer. The mechanism has a fixed production cost γ for the item. Under separability assumptions, the optimal mechanism in this setting has a surprising simple form (with a simple indirect implementation which we present later) — the mechanism offers the agent a “menu” of contracts, of the form $(p, M(p))$ to the agent. If an agent chooses a contract, she will be charged an upfront payment of $M(p)$ and afterwards the mechanism posts a price of $p > \gamma$ at each time step — the agent has the option to pay more upfront for cheaper prices in the future. Note that even if the agent’s valuation is increasing (or decreasing) over time and the seller is fully aware of this fact, the optimal mechanism involves offering the item at all periods at a constant price p .

In the general solution with multiple buyers, the VIRTUAL-PIVOT Mechanism still retains this flavor. Roughly speaking, each agent, based on her initial type, is assigned a certain weight function in an affine transformation of the social welfare that is maximized by the mechanism, see Section 5.1. The more the agent pays up front the higher her importance will be in the social welfare function (leading to more allocations to her in the future).

Our setting considers a mechanism which allows agents to report their type every round. In particular, this implies that they are able to *re-report* all of their historical private information that has bearing on the current and future values. Allowing re-reporting of private signals is a crucial step in obtaining periodic ex-post incentive guarantees. Once we obtain periodic ex-post incentive compatibility for all future periods, we are able to provide necessary and sufficient conditions for incentive compatibility at the first period. We directly show these conditions are satisfied for our optimal mechanism.

Finally, we provide examples of how the standard “relaxation” approach to dynamic mechanism design will not succeed without certain additional assumptions.

1.1 Related Work

Two natural objectives in the dynamic mechanism design are maximizing the long term social welfare of all buyers (*efficiency*) and maximizing the long term revenue or profit of a seller (*optimality*). With regards to maximizing the long term social welfare, there are elegant extensions of the efficient (VCG) mechanism to quite general dynamic settings, including the dynamic pivot mechanism of Bergemann and Välimäki [2010] and the dynamic team mechanism of Athey and Segal [2007] (see also Cavallo et al. [2007], Bapna and Weber [2008], Nazerzadeh et al. [2008]).

The literature on dynamic revenue-optimal mechanism has been primarily focused on settings

where the agents arrive and depart dynamically over time, but their private information remains *fixed*, see Vulcano et al. [2002], Pai and Vohra [2008], Said [2009], Gershkov and Moldovanu [2009], Skrzypacz and Board [2010]. In these setting, the mechanism designer faces a dynamic problem, but the incentive constraints of each agents are essentially static because agents do not obtain any “new” private information over the course of the mechanism. For surveys on dynamic mechanism design see Bergemann and Said [2011], Parkes [2007].

We consider a setting where the private information of the agents changes over time, a line of research that was pioneered by Baron and Besanko [1984]. Courty and Li [2000] provide an optimal mechanism for an environment where agents have private information about the future distribution of their valuations. Battaglini [2005] studies a setting with a single agent whose private information is given by a 2-state Markov Chain and shows that the optimal allocation converges over time to the efficient allocation. In contrast to the results in Battaglini [2005], in the settings we consider, the allocation distortion generated by the agents’ initial private information does not disappear with time. For a more detailed discussion, see Subsection 5.1.

A closely related work to ours is that of Eso and Szentes [2007] who study a two-period model where each agent receives a signal at the first period and the seller can also allow each agent to receive an additional private signal at the second period. Under certain concavity and monotonicity conditions on the signals, they show that the optimal mechanism allows the agents to receive their second signals; however, agents do not obtain any rents from the fact second period signal is private. They also propose a ‘handicap’ auction for the case where the agents’ valuations are given by the sum of the first and second period signals. We use similar ideas and show that for a broad class of environments, the seller is able to extract the information rent associated with all signals except the initial one, even if the seller does not control the agents’ ability to obtain further private signals. However, as we show in Section 7, there exist dynamic settings where the seller cannot extract the entire information rent from future signals. We also note the work in Deb [2008], which provides an optimal mechanism in a setting with only one buyer where the value is Markovian in the previous value, among other technical conditions.

Another closely related work to ours is by Pavan et al. [2008, 2009], which presents a comprehensive characterization of *necessary* conditions for incentive compatibility in both finite horizons (Pavan et al. [2008]) and infinite horizons (Pavan et al. [2009]). Finding sufficient conditions for incentive compatibility turns out to be a difficult challenge, as Pavan et al. [2008] acknowledge.¹ A setting where they show incentive compatibility (Proposition 12) is one where types evolve according to an $AR(k)$ auto-regressive stochastic process with non-negative coefficients, the set of feasible actions of the mechanism is a lattice (therefore excluding problems such as allocating private goods between two or more agents) and actions do not affect the evolution of types.²

See Subsection 3.3 for further discussion on how our methodology relates to prior work.

¹“As for incentive compatibility in period one, we were only able to check it application-by-application, but we have been able to verify it in a few special settings.” (Pavan et al. [2008], page 4)

²See Assumptions SCP and DNOT, both used in proving Proposition 12 in Pavan et al. [2008].

1.2 Organization

We organize our paper as follows. In Section 2, we formalize our model, define concepts such as incentive compatibility and optimality of mechanisms. In Section 3, we discuss our approach for designing optimal mechanisms. In particular, both necessary and sufficient constraints for incentive compatibility are provided here. We introduce the notion of separability in Section 4 and provide an upper bound on the revenue of any mechanism in a separable environment. In Section 5, we propose our mechanism and states our main optimality result. Special cases (including the setting with only one buyer) are considered in Section 6. Section 7 provides simple examples showing how the usual incentive constraints from static mechanism are insufficient for the dynamic case. It also shows that without our separability assumptions, the particular relaxation approach we take is insufficient. The Appendix contains all the proofs.

2 Preliminaries

In this section, we formalize our model and define concepts such as incentive compatibility and optimality of mechanisms.

2.1 The Dynamic Environment

We consider a discrete-time, δ -discounted infinite-horizon ($t = 0, 1, 2, \dots$) model that consists of one seller and n agents (buyers). The seller decides upon an action a_t at each period t among the feasible set of actions \mathcal{A}_t , at a cost of $c_t(a^t)$ to the seller, where $a^t = (a_0, a_1, \dots, a_t)$ represents all the actions taken by the mechanism up to time t .

At every period, each agent $i \in \{1, \dots, n\}$ receives a private signal $s_{i,t} \in S_{i,t}$. In particular, we make the following assumption about the first signal $s_{i,0}$ throughout the paper:

Assumption 2.1. *For each agent i , $s_{i,0} \in [0, 1]$ is real valued and distributed according to F_i . Furthermore, assume that F_i is strictly increasing and has a density, which we denote by f_i .*

This first signal summarizes all the initial private information of the agent (which has bearing on her entire stream of valuations). Furthermore, for all $t \geq 1$, each agent also receive a private signal $s_{i,t} \in S_{i,t}$ — here we not concerned with whether or not these future signals are real or not (the set $S_{i,t}$ is arbitrary for $t \geq 1$).

The *type* of agent i at time t is the sequence of signals of the buyer i up to (and including) time t , which is denoted by $s_i^t = (s_{i,0}, \dots, s_{i,t})$. The type provides a summary of all the agent's private information which has bearing on all her current and future valuations. For notational convenience, we let vector $s^t = \{s_i^t\}_{i \in [n]}$ denote the (joint) types of all agents at time t . At each period t , agent i obtains value $v_{i,t}(a^t, s_i^t)$, which is a function of her type and the seller's past and current actions. We assume quasi-linear utilities and denote the payment of agent i at time t by $p_{i,t}$, so that the (instantaneous) utility of agent i at time t is given by

$$u_{i,t} = v_{i,t}(a^t, s_i^t) - p_{i,t}.$$

We also assume throughout the following regularity condition.

Assumption 2.2. *The partial derivative $\frac{\partial v_{i,t}(a^t, s_{i,0}, \dots, s_{i,t})}{\partial s_{i,0}}$ exists for all i, t, a^t , and s_i^t , and it is bounded by $\bar{V} < \infty$.*

We now specify the stochastic process over the signals. The signal $s_{i,t}$ that agent i receives at time t may be correlated to her previous signals $s_{i,0}, \dots, s_{i,t-1}$ and the past actions of the seller a_0, \dots, a_{t-1} , but it is independent (conditionally on the seller's actions) of all signals of the other agents. Formally, the stochastic signal $s_{i,t}$ is determined by the stochastic kernel $K_{i,t}(s_{i,t}|a^{t-1}, s_i^{t-1})$. Without loss of generality, we make the assumption that the first signal is independent of the future signals:

Assumption 2.3. *For each agent i , the distribution of the initial signal $s_{i,0}$ is independent of the future signals $s_{i,t}$ for $t \geq 1$.*

In fact, without loss of generality, one can assume that all the signals are independent (as noted by Pavan et al. [2009]³). Even under this assumption, importantly, note that the value of agent i at any future period ($t \geq 1$) may still be correlated with the signal $s_{i,0}$. Here, we only explicitly assume $s_{i,0}$ to be independent of the future — arbitrary dependencies among future signals are permitted. While this assumption is without loss of generality, an example in Section 7 suggests that there are technical reasons for which this is formulation natural (in particular, this example shows how the relaxation approach we take may fail in a different representation).

We also assume the mechanism has the ability to exclude agents from the system at time $t = 0$. That is, it can select a subset of the agents that will obtain no value (and will not make payments) at any period $t \geq 0$. The exclusion of an agent from the system does not impact the value obtained by the other agents if the mechanism still takes the same sequence of actions a_1, \dots, a_t .

Assumption 2.4. *The set of feasible actions \mathcal{A}_0 at time $t = 0$ is equal to $2^{\{1, \dots, n\}}$, that is, the set of all subsets of $\{1, \dots, n\}$. If $i \notin a_0$, then agent i is excluded from the system, i.e., $p_{i,t} = 0$ and $v_{i,t}(a^t, s_i^t) = 0$ for all t, a^t , and s_i^t . No agent obtains immediate value from the choice of a_0 , i.e., $v_{i,0}(a_0, s_{i,0}) = 0$ irrespective of whether $i \in a_0$ or not. Also, the value obtained by each agent does not depend on the exclusion of other agents. In addition, the cost incurred by the mechanism only depends on the actions not on the excluded agents.*

The assumption implies that for any pair of actions a_0, a'_0 in \mathcal{A}_0 such that $i \in a_0$ and $i \in a'_0$, the value $v_{i,t}(a_0, a_1, \dots, a_t, s_i^t) = v_{i,t}(a'_0, a_1, \dots, a_t, s_i^t)$ for all t, a_1, \dots, a_t , and s_i^t . Also, $c_t(a_0, a_1, \dots, a_t) = c_t(a'_0, a_1, \dots, a_t)$ for all t — of course, exclusion of an agent may change the choice of the actions taken by the mechanism. The assumption that the agents do not obtain value at $t = 0$ is made without loss of generality and for simplicity of presentation. Nevertheless, the mechanism may charge the agents $p_{i,t} \neq 0$ at that time. The above assumption simplifies satisfying the participation constraints. For example, if an agent only obtains negative values from the actions, she would be excluded from the mechanism. Observe that if the actions taken by the mechanism correspond to allocations of items to agent, this assumption can be simply satisfied.

³ The argument essentially follows from this observation: any stochastic process $\{X_t\}_{t \in \mathbb{N}}$ may be simulated by a function f such that $x_{t+1} = f(x_1, \dots, x_t, z_t)$ where the z_t 's are i.i.d. and uniform on $[0, 1]$ (here f is constructed based on the distribution $\Pr(X_{t+1}|X_1, \dots, X_t)$); see Pavan et al. [2009].

At each period $t \geq 0$,

1. Each agent i receives her private signal $s_{i,t} \sim K_{i,t}(\cdot | a^{t-1}, s_i^{t-1})$.
2. Each agent i provides a report, \hat{s}_i^t , of her current type, $s_i^t = (s_{i,0}, \dots, s_{i,t})$, as determined by her private history $h_{i,t}$. In particular, $\hat{s}_i^t = R_i(h_{i,t})$.
3. As a function of the public history, h_t , and the current reports, \hat{s}^t , the mechanism determines the action $a_t \in \mathcal{A}_t$ and the payments $p_{i,t}$ for each agent i . In particular, $a_t = q(h_t, \hat{s}^t)$ and the joint prices are $\{p_{i,t}\}_{i \in [n]} = p(h_t, \hat{s}^t)$.

Figure 1: A generic mechanism

For our theorems to hold, we make further restrictions on the functional form of the values. In particular, we will place a *separability* assumption on the environment, which we precisely state in Section 4.

Throughout the paper, suppose Assumptions 2.1, 2.2, 2.3 and 2.4 hold.

2.2 Mechanisms, Incentive Constraints, and Optimality

A mechanism $\mathcal{M}(q, p)$ is defined by a pair of an allocation rule $q(\cdot)$ and a payment rule $p(\cdot)$. We let \mathcal{Q} denote the set of all allocation rules. By the Revelation Principle (cf. Myerson [1986]), without loss of generality, we focus on (dynamic) direct mechanisms.⁴

At each period t , each agent i , makes a *report*, denoted by \hat{s}_i^t , of her type s_i^t . Using our standard shorthand notation, we denote the joint reports of all agents by $\hat{s}^t = \{\hat{s}_i^t\}_{i \in [n]}$. Note that since $s_i^t = (s_{i,0}, \dots, s_{i,t})$ includes the set of all signals that each agent has received, each agent *re-reports* all of their previous signals at every period. The report of an agent can be conditioned on the history, which we now specify.

The public history at time t , denoted by h_t , is the sequence of reports and actions of the mechanism until period $t - 1$; namely, $h_t = (\hat{s}_0, a_0, \hat{s}^1, a_1, \dots, \hat{s}^{t-1}, a_{t-1})$. The private history of agent i at time t , denoted by $h_{i,t}$, includes the public history and her current type (sequence of signals she received up to, and including, time t), i.e., $h_{i,t} = (s_{i,0}, \hat{s}_0, a_0, s_{i,1}, \hat{s}^1, a_1, \dots, s_{i,t-1}, \hat{s}^{t-1}, a_{t-1}, s_{i,t})$.

The allocation and payment rules are functions of the public history at time t , h_t , and the reports of all agents at time t , \hat{s}^t . The allocation rule determines the action taken by the mechanism and the payment rule determines the payment of each agent.

The reporting strategy of agent i , denoted by R_i , is a mapping from her private history $h_{i,t}$ to a report of her current type \hat{s}_i^t . Mechanism \mathcal{M} and the reporting strategy profile $R = \{R_i\}_{i \in [n]}$ determine a stochastic process which is described in Figure 1.

We now define the incentive constraints of the mechanism. Denote the expected (discounted) future

⁴ The Revelation Principle implies that an equilibrium outcome in any indirect mechanism can also be induced as an equilibrium outcome of an (incentive compatible) direct mechanism.

value of agent i under the (joint) reporting strategy R in mechanism \mathcal{M} by:

$$V_i^{\mathcal{M},R} = \mathbb{E} \left[\sum_{t=0}^{\infty} \delta^t v_{i,t}(a^t, s_i^t) \right]$$

and the expected (discounted) future utility (of i under R in \mathcal{M}) as:

$$U_i^{\mathcal{M},R} = \mathbb{E} \left[\sum_{t=0}^{\infty} \delta^t (v_{i,t}(a^t, s_i^t) - p_{i,t}) \right]$$

where the expectation is with respect to the stochastic process induced by the reporting strategy and the mechanism. Similarly, for the expected value and utility of agent i , conditioned on a private history $h_{i,t}$ and type of the other agents s_{-i}^t , we have:

$$\begin{aligned} V_i^{\mathcal{M},R}(h_{i,t}, s_{-i}^t) &= \mathbb{E} \left[\sum_{\tau=t}^{\infty} \delta^\tau v_{i,\tau}(a^\tau, s_i^\tau) \middle| h_{i,t}, s_{-i}^t \right] \\ U_{i,t}^{\mathcal{M},R}(h_{i,t}, s_{-i}^t) &= \mathbb{E} \left[\sum_{\tau=t}^{\infty} \delta^\tau (v_{i,\tau}(a^\tau, s_i^\tau) - p_{i,\tau}) \middle| h_{i,t}, s_{-i}^t \right] \end{aligned}$$

Note that this expectation is well defined (even on private histories which have probability 0 under R), since the reporting strategies are mappings from *all* possible private histories of agent i (and we have conditioned on the public history and current joint type).

Roughly speaking, the notion of incentive compatible is one in which no agent wants to deviate from the truthful strategy, as long as all other agents are truthful. This involves a somewhat delicate quantification with regards to the history. Our (weaker and stronger) notions of incentive compatibility are identical to those in Bergemann and Välimäki [2010].

Definition 2.1. (*Incentive Compatibility*) Let \mathcal{T} denote the (joint) truthful reporting strategy.

- *Dynamic mechanism \mathcal{M} is (Bayesian) incentive compatible (IC) if, for each agent i , truthfulness is a best response to the truthful strategy of other agents — precisely, if for each i and R_i ,*

$$U_i^{\mathcal{M},\mathcal{T}} \geq U_i^{\mathcal{M},(R_i,\mathcal{T}_{-i})}$$

- *Dynamic mechanism \mathcal{M} is periodic ex-post incentive compatible if, for each agent i and at any time t , truthfulness is a best response to the truthful strategy of other agents — precisely, if for each i and time t , reporting strategy R_i , private history $h_{i,t}$, and current type of the other agents s_{-i}^t :*

$$U_{i,t}^{\mathcal{M},\mathcal{T}}(h_{i,t}, s_{-i}^t) \geq U_{i,t}^{\mathcal{M},(R_i,\mathcal{T}_{-i})}(h_{i,t}, s_{-i}^t) \quad (1)$$

Note that the (weaker) Bayesian notion of IC implies that the truthful reporting strategy is a best response from a private history that is generated under \mathcal{T} with probability 1. In contrast, the (stronger) periodic ex-post notion demands that the truthful strategy is a best response on *any*

private history, even those which have probability 0 under \mathcal{T} (e.g. those histories where agents mis-reported in the past). See Bergemann and Välimäki [2010] for further discussion.

The notion of individual rationality is one, where at the equilibrium, the agents choose to participate (as it demands that the agents’ utilities be non-negative). Precisely,

Definition 2.2. (*Individual Rationality*) Let \mathcal{T} denote the (joint) truthful reporting strategy.

- Mechanism \mathcal{M} is (Bayesian) individually rational (IR) if, for each agent i , the expected future utility under the truthful strategy is non-negative, i.e., $U_i^{\mathcal{M},\mathcal{T}} \geq 0$.
- Mechanism \mathcal{M} is periodic ex-post individually rational if the expected future utility is non-negative for each agent i and time t , private history, $h_{i,t}$, and joint type of the other agents s_{-i}^t , i.e., $U_{i,t}^{\mathcal{M},\mathcal{T}}(h_{i,t}, s_{-i}^t) \geq 0$.

The *expected profit* of a mechanism \mathcal{M} is the discounted sum of all payments of the agents minus the cost of the actions

$$\text{Profit}^{\mathcal{M}} = \mathbb{E} \left[\sum_{t=0}^{\infty} \delta^t \left(-c_t(a^t) + \sum_{i=1}^n p_{i,t} \right) \right] \quad (2)$$

under the (joint) truthful reporting strategy \mathcal{T} . The objective of the seller is to maximize this expected profit, subject to both the incentive compatibility and individual rationality constraints. Precisely,

Definition 2.3. (*Optimality*) A Bayesian individually rational and Bayesian incentive compatible mechanism is optimal if it maximizes the expected profit among all Bayesian individually rational and Bayesian incentive compatible mechanisms.

Note the optimal mechanism is only required to satisfies the weaker Bayesian incentive constraints. This definition of optimality guarantees that the mechanism obtains an expected profit higher than (or, at least, equal to) any other mechanism that is incentive compatible and individually rational. Ideally, we might hope for an optimal mechanism which also satisfies the stronger (periodic ex-post) incentive constraints, which ensures truthfulness is a best response even if agents have deviated in the past. As we show, the mechanism we propose, the VIRTUAL-PIVOT Mechanism, enjoys these stronger guarantees.

3 A Relaxation Approach

We now provide a methodology for optimal dynamic mechanism design. The relaxation approach we take is the standard one also used in Eöso and Szentes [2007], Pavan et al. [2008], Deb [2008]. The difficulty is in “un-relaxing”, i.e., showing that a candidate for the optimal policy satisfies the more stringent dynamic IC constraints.

Here, we are able to provide both necessary and sufficient conditions for dynamic IC. In particular, the use of the periodic ex-post notion of incentive compatibility is critical in this characterization.

3.1 Relaxing

In this section, we consider a simpler, yet closely related, problem where we can utilize known static mechanism design techniques to design an optimal mechanism — these techniques are also used in [Eso and Szentes \[2007\]](#), [Pavan et al. \[2008\]](#), [Deb \[2008\]](#). The idea is to relax the optimization problem (of finding the optimal mechanism) by only imposing certain incentive constraints that arise in a simpler version of the problem. Roughly speaking, we attempt to solve a (simpler) less-constrained optimization problem. The critical issue is in showing that the solution to this less-constrained problem is also the optimal solution for the original problem.

Definition 3.1. (*Relaxed Environment*) Consider an environment where only the initial type $s_{i,0}$ is private to each agent i , while all her future signals are observed by the mechanism. We define this to be the relaxed environment and refer to our original environment as the dynamic environment.

While the mechanism in the relaxed environment has full information with regards to the agents signals from $t \geq 1$, note that $s_{i,0}$ may affect all the future values of the agent. Observe that any direct mechanism in the dynamic environment induces a mechanism in the relaxed environment in a natural way: for $t \geq 1$, simply use the agents actual signals $s_{i,1}, \dots, s_{i,t}$ as well as the reported initial signal $\hat{s}_{i,0}$ as the reported type $\{\hat{s}_i^t\}$ (as the input to the allocation and payment rules of the mechanism).

The following lemma is a rather straightforward observation.

Lemma 3.1. Let \mathcal{E} be a dynamic environment and $\mathcal{E}^{\text{relaxed}}$ be the corresponding relaxed environment. We have that:

- If \mathcal{M} is an incentive compatible and individually rational mechanism in \mathcal{E} , then it is an incentive compatible and individually rational mechanism in $\mathcal{E}^{\text{relaxed}}$.
- Let R^* be the optimal revenue in $\mathcal{E}^{\text{relaxed}}$. Suppose a (Bayesian) incentive compatible and individually rational mechanism \mathcal{M} in \mathcal{E} has revenue R^* , then \mathcal{M} is optimal for both \mathcal{E} and $\mathcal{E}^{\text{relaxed}}$.

This lemma suggest a natural optimal mechanism design approach: first, find an allocation rule q^* of an optimal mechanism in the relaxed environment $\mathcal{E}^{\text{relaxed}}$; then determine if there exists a pricing rule for p^* such that: 1) the mechanism (q^*, p^*) is IC and IR in the dynamic environment \mathcal{E} ; 2) the expected revenue it achieves is R^* . If such a pricing is possible, then (q^*, p^*) is optimal in \mathcal{E} . In our separable environments, we show this approach is applicable. Furthermore, in [Section 7](#), we discuss the limitations of this approach, where we provide certain non-separable environments for which the optimal revenue in \mathcal{E} is strictly less than the optimal revenue in $\mathcal{E}^{\text{relaxed}}$.

Envelope and Revenue Lemmas

Since in the relaxed environment the only piece of private information for each agent i is $s_{i,0}$, using the standard approach from static mechanism design (see [Myerson \[1981\]](#), [Milgrom and Segal \[2002\]](#)), we provide the following lemma.

Lemma 3.2. (*Envelope Condition*) Suppose that the mechanism \mathcal{M} is IC in the relaxed environment. Then for all i , $s_{i,0}$ and $s'_{i,0}$,

$$U_i(s_{i,0}, s_{-i,0}) - U_i(s'_{i,0}, s_{-i,0}) = \int_{s'_{i,0}}^{s_{i,0}} \mathbb{E} \left[\sum_{t=0}^{\infty} \delta^t \frac{\partial}{\partial s_{i,0}} v_{i,t}(a^t, s_{i,0}, s_{i,1}, \dots, s_{i,t}) \Big|_{s_{i,0}=z} \Big| s_{i,0} = z, s_{-i,0} \right] dz. \quad (3)$$

where $U_i(s_{i,0}, s_{-i,0})$ is utility of agent i under the truthful strategy in \mathcal{M} , with the initial types are $s_{i,0}$ for i and $s_{-i,0}$ for the other agents.

Again using standard techniques from static mechanism design, we can use the envelope condition above to establish the profit of any IC mechanism in the relaxed environment.

Lemma 3.3. (*Expected Profit*) Suppose that the mechanism \mathcal{M} is IC in the relaxed environment. Then, the expected profit obtained by the mechanism, $\text{Profit}^{\mathcal{M}}$, is equal to:

$$\mathbb{E} \left[\sum_{t=1}^{\infty} \delta^t \left(\sum_{i=1}^n \left(v_{i,t}(a^t, s_i^t) - \frac{1 - F_i(s_{i,0})}{f_i(s_{i,0})} \frac{\partial v_{i,t}(a^t, s_{i,0}, s_{i,1}, \dots, s_{i,t})}{\partial s_{i,0}} \right) - c_t(a^t) \right) - \sum_{i=1}^n U_i^{\mathcal{M}}(0, s_{-i,0}) \right] \quad (4)$$

where the expectation is taken over $s_{i,0}$ and $s_{-i,0}$.

This lemma can be used to derive a candidate for the optimal allocation rule: if we pick an allocation rule that maximizes the equation above and pick a payment rule that makes it both IC and IR, then we will have an optimal mechanism.

3.2 Un-Relaxing

From the relaxed environment, we can find a candidate for an optimal allocation rule. The main challenge here is how to find a payment rule and show that such a mechanism satisfies dynamic IC constraints. It turns out that it is natural to break this into two stages.

The first step is understanding how to ensure IC for $t \geq 1$. Here, there seems to be no general methodology in the literature (note that we are not assuming any structure on the stochastic process for the signals s_t , for $t \geq 1$). Our approach involves going one step further and trying to insure periodic ex-post IC for periods $t \geq 1$. Recent work by Bergemann and Välimäki [2010] show how to guarantee periodic ex-post IC in the context of maximizing social welfare. Our results make use of this, but to do so, a critical conceptual step is to allow agents to re-report their entire type at every period. This way, we are able to obtain periodic ex-post IC for $t \geq 1$.

For $t = 0$, where $s_{i,0}$ is real valued, we explicitly characterize the necessary and sufficient conditions for dynamic IC based on the fact that we have a periodic ex-post IC mechanism for periods $t \geq 1$. This is a key technical step in our proof.

Re-Reporting and Periodic Ex-Post IC

Recall that each agent i reports her entire type $s_i^t = (s_{i,0}, \dots, s_{i,t})$ at each period t , not just her most recent private signal $s_{i,t}$. At the first glance, it may seem that this re-reporting of past private

signals is redundant. However, there are a few of reasons why this approach is quite natural, both conceptually and technically.

Re-reporting simplifies the task of obtaining periodic ex-post IC guarantees. It gives an opportunity for agents that have reported untruthfully in the past to correct their past mis-reports and, in this way, return to truthful reporting course. In fact, it is unclear how to obtain such a guarantee for a mechanism which does not allow re-reporting — such a mechanism must incentivize an agent who has mis-reported in the past to report truthfully in the future, without the ability to query the agents about their previous signals.

For such mechanisms (which do not allow re-reporting), often the best that the mechanism can do is offer the agent an opportunity to misreport again in order to course-correct. In essence, these are the techniques used in [Eso and Szentes \[2007\]](#), [Pavan et al. \[2008\]](#), [Deb \[2008\]](#) to obtain IC (not periodic ex-post IC) by restricting the value functions or the stochastic process.

Necessary and Sufficient Conditions for IC

In the previous subsection, we argued that re-reporting simplifies the task of constructing a periodic ex-post IC mechanism. We postpone the discussion of how we can use re-reporting to actually construct a periodic ex-post IC mechanism until Section 5.

For now, assume that a mechanism \mathcal{M} is periodic ex-post IC for all periods $t \geq 1$. That is, for any period $t \geq 1$, any agent i , private history $h_{i,t}$, types of other agents s_{-i}^t , and reporting strategy R_i , Eq. (1) is satisfied. We now provide necessary and sufficient conditions for such a mechanism to be IC (at period $t = 0$).

Consider a subset of an agent's reporting strategies that we denote by $x' \rightarrow x$. Define $x' \rightarrow x$ as the reporting strategy in which the agent reports x' as her first type $s_{i,0}$ (at $t = 0$), and subsequently (re-)reports it as x in all future periods ($t \geq 1$). Furthermore, under the strategy $x' \rightarrow x$, all other signals $s_{i,t}$ (for $t \geq 1$) are truthfully reported. In others words, at $t = 0$, she initially reports $\hat{s}_{i,0} = x'$, and, for $t \geq 1$, she reports $\hat{s}_i^t = (x, s_{i,1}, s_{i,2}, \dots, s_{i,t})$. In $x' \rightarrow x$, we also allow x' and x to be functions of $s_{i,0}$. For example, the truthful strategy \mathcal{T}_i can be represented as $s_{i,0} \rightarrow s_{i,0}$.

The expected utility of agent i under mechanism \mathcal{M} and reporting strategy $x' \rightarrow x$ given her initial type $s_{i,0}$ is $U^{\mathcal{M},(x' \rightarrow x, \mathcal{T}_{-i})}(s_{i,0})$. For notational convenience, we drop the explicit dependence on the mechanism and the other agents' playing the truthful strategy and denote this by:

$$U^{x' \rightarrow x}(s_{i,0}) = U^{\mathcal{M},(x' \rightarrow x, \mathcal{T}_{-i})}(s_{i,0}). \quad (5)$$

Similarly, we define the expected value of agent i under strategy $x' \rightarrow x$, assuming other agents are truthful by:

$$V^{x' \rightarrow x}(s_{i,0}) = V^{\mathcal{M},(x' \rightarrow x, \mathcal{T}_{-i})}(s_{i,0}). \quad (6)$$

We also use the notation $U^{x' \rightarrow x}(s_{i,0}, s_{-i,0})$ and $V^{x' \rightarrow x}(s_{i,0}, s_{-i,0})$, when we condition on the initial types of the other agents $s_{-i,0}$.

Suppose the mechanism \mathcal{M} is one which is periodic ex-post IC for periods $t \geq 1$. Under such a mechanism, if agent i deviates at period $t = 0$, while all other agents are truthful, agent i 's best response strategy at all future periods $t \geq 1$ is to reveal her true type. Therefore, if her true first

type is $s_{i,0}$, then to verify if truthfulness is a best response, we only need to verify that the truthful policy provides more utility than all misreporting strategies of the form $s'_{i,0} \rightarrow s_{i,0}$. Therefore, if mechanism \mathcal{M} is periodic ex-post IC for periods $t \geq 1$, then it is also IC at $t = 0$ if, and only if, for any true type x and time 0 report x' ,

$$U_i^{x \rightarrow x}(x) \geq U_i^{x' \rightarrow x}(x).$$

Subtracting $U_i^{x' \rightarrow x'}(x')$ from both sides, we get the following characterization: the mechanism \mathcal{M} is IC if, and only if, for all x and x' ,

$$U_i^{x \rightarrow x}(x) - U_i^{x' \rightarrow x'}(x') \geq U_i^{x' \rightarrow x}(x) - U_i^{x' \rightarrow x'}(x'). \quad (7)$$

Furthermore \mathcal{M} is periodic ex-post IC if the above holds where we condition on the other types $s_{-i,0}$. That is, the mechanism is periodic ex-post IC if for all x , x' and $s_{-i,0}$,

$$U_i^{x \rightarrow x}(x, s_{-i,0}) - U_i^{x' \rightarrow x'}(x', s_{-i,0}) \geq U_i^{x' \rightarrow x}(x, s_{-i,0}) - U_i^{x' \rightarrow x'}(x', s_{-i,0}). \quad (8)$$

These observations are useful in that it we can use envelope conditions to precisely characterize incentive compatibility in terms of the expected values of the agents. First, we obtain that periodic ex-post IC for $t \geq 1$ implies the following lemma.

Lemma 3.4. *(Periodic Ex-Post IC) Suppose that mechanism \mathcal{M} satisfies the periodic ex-post IC conditions for all $t \geq 1$. Then, for all x and x' in $[0, 1]$, we have*

$$U_i^{x' \rightarrow x}(x) - U_i^{x' \rightarrow x'}(x') = \int_{x'}^x \frac{\partial V_i^{x' \rightarrow z}(s)}{\partial s} \Big|_{s=z} dz \quad (9)$$

It is straightforward to show that the partial derivative exists and, for any x , y and z , is given by:

$$\frac{\partial V^{x \rightarrow y}(s)}{\partial s} \Big|_{s=z} = \mathbb{E} \left[\sum_{t=0}^{\infty} \delta^t \frac{\partial}{\partial s_{i,0}} v_{i,t}(a^t, s_{i,0}, s_{i,1}, \dots, s_{i,t}) \Big|_{s_{i,0}=s} \Big| s_{i,0} = z \right] \quad (10)$$

where the expectation is under joint strategy $(x \rightarrow y, \mathcal{T}_{-i})$ in \mathcal{M} (see Lemma A.1 in the Appendix).

The following lemma uses the characterization above to obtain both necessary and sufficient conditions for incentive compatibility (at $t = 0$).

Lemma 3.5. *(Necessary and Sufficient Conditions for IC) Suppose that the mechanism \mathcal{M} satisfies the periodic ex-post IC conditions for all $t \geq 1$. Then, \mathcal{M} is IC for all $t \geq 0$ if, and only if, both conditions below are satisfied:*

- *(Envelope Condition) For all x and x' ,*

$$U_i^{x \rightarrow x}(x) - U_i^{x' \rightarrow x'}(x') = \int_{x'}^x \frac{\partial V_i^{z \rightarrow z}(s)}{\partial s} \Big|_{s=z} dz. \quad (11)$$

- *(Interval Dominance) For all x and x' ,*

$$\int_{x'}^x \frac{\partial V_i^{z \rightarrow z}(s)}{\partial s} \Big|_{s=z} dz \geq \int_{x'}^x \frac{\partial V_i^{x' \rightarrow z}(s)}{\partial s} \Big|_{s=z} dz. \quad (12)$$

Furthermore, \mathcal{M} is ex-post periodic IC if and only if the previous two conditions are satisfied when we condition on every possible other initial types $\hat{s}_{-i,0}$.

The result above is analogous to the characterization of incentive compatibility in standard single-parameter settings, where an envelope condition and monotonicity are used to characterize IC (see Myerson [1981]). The envelope condition above is a standard one, but interval dominance replaces monotonicity in a dynamic setting. It compares the utility obtained by the truthful strategy (left-hand side) with other strategies of the form $x' \rightarrow s_{i,0}$ (right-hand side), as these are the only plausible candidate strategies when the mechanism is ex-post IC for periods $t \geq 1$.

3.3 On Our Methodology

Although previous papers in the literature (see $\ddot{\text{E}}\text{so}$ and Szentes [2007], Pavan et al. [2008]) also provide optimal mechanisms using the relaxation approach, we emphasize that our construction and results do *not* immediately follow from them. The key challenge we address in our paper is showing that the allocation rule generated by the relaxation has an associated payment rule that makes the mechanism IC and IR in the dynamic setting. Our solution requires a combination of using the re-reporting technique, with constructing payments based on $\text{Bergemann and V\ddot{a}lim\ddot{a}ki}$ [2010] to obtain periodic ex-post IC for periods $t \geq 1$, as well as proving IC (at $t = 0$) by using our characterization of IC under the assumption of periodic ex-post IC for $t \geq 1$.

Furthermore, we show in Section 7 that the relaxation approach does not work in every setting. In fact, the second example provides a simple dynamic environment in which the usual notions of monotonicity hold for the optimal allocation in the relaxed environment, and yet, this same allocation rule is not optimal in the dynamic environment (clearly showing how static notions of monotonicity are insufficient). Although we are not able to address the challenging problem of explicitly characterizing the necessary and sufficient properties of an environment for which this relaxation approach will succeed, we do provide environments in which both: the relaxation approach fails and various assumptions of our separable environment are violated. Roughly speaking, these show that at least some variant of our assumptions are required for the relaxation approach to be successful.

4 Separability

We now define a class of environments where the optimal allocation in the relaxed environment is closely related to an (affinely transformed) efficient allocation. In the next section, we provide an optimal mechanism for this class.

To be able to construct an optimal dynamic mechanism, we need to assume some structure on how the agents' values relate to their signals. The next property specifies two natural relationships between the signals and the values.

Property 4.1. (*Functional Separation*) *An environment satisfies functional separation if the value function of each agent is either multiplicatively or additively separable:*

- The value function of agent i is **multiplicatively-separable** if there exists functions A_i and $B_{i,t}$ such that:

$$v_{i,t}(a^t, s_i^t) = A_i(s_{i,0})B_{i,t}(a^t, s_{i,1}, \dots, s_{i,t}) \quad (13)$$

- The value function of agent i is **additively-separable** if there exists A_i , $B_{i,t}$, $C_{i,t}$ such that:

$$v_{i,t}(a^t, s_i^t) = A_i(s_{i,0})C_{i,t}(a^t) + B_{i,t}(a^t, s_{i,1}, \dots, s_{i,t}) \quad (14)$$

Definition 4.1. We call an environment separable if Assumption 2.3 and Property 4.1 hold.⁵

Note that an environment may not be separable at the first glance, but there might exist a transformation of the signals and value functions which makes the environment separable.

In the following section, unless otherwise stated, we assume the environment is separable.

4.1 Examples of Separable Environments

Sponsored Search A prominent example of multiplicatively separable value functions arises in the setting of online advertising. Consider a sponsored search auction for a keyword that corresponds to a certain product. Suppose agent i is an online retailer of such a product who participates in the corresponding auction. Every time a user types in the keyword, the ad spaces are allocated to the retailers. Every time a user purchases the product from them, the retailer i obtains a profit of $s_{i,0}$ (and 0 otherwise). The type of each agent i (besides $s_{i,0}$) would represent the Bayesian belief about the probability of a purchase occurring given the retailer’s ad is shown. Therefore, $v_{i,t}(s_i^t) = s_{i,0} \times \Pr[\text{purchase}|s_{i,1}, \dots, s_{i,t}]$. After each time the ad of retailer i is shown to a user, the retailer updates her belief about probability of a purchase.

Supply Chain A widget manufacturer supplies one or more retailers who sells these products to consumers. The widgets have some associated production cost $c_t(\cdot)$ that are borne by the manufacturer. The retailers have two pieces of private information: their profit margin on the widgets (which they know upfront and is captured by $s_{i,0}$) and the demand they face for the widgets (which they learn over time and is captured by $s_{i,1}, \dots, s_{i,t}$). Each retailer’s value is then given by $v_{i,t}(s_i^t) = s_{i,0} \times \text{Demand}(s_{i,1}, \dots, s_{i,t})$.

AR(k) For an example of additively separable value function, consider auto-regressive (AR) processes. One example of an $AR(k)$ model for the evolution of the valuation of each agent i is as follows: the initial value of agent i is given by $v_{i,1} = s_{i,0}$, and every time the item is allocated to agent i her valuation is updated according to $v_{i,t} = \sum_{\tau=1}^k \gamma_{i,t,\tau} v_{i,t-\tau} + \eta_{i,t}(a^t, s_{i,1}, \dots, s_{i,t})$, where $\gamma_{i,t,\tau}$ are constants and η is a noise process. It is straightforward to use functions $A_i(s_{i,0})$, $C_{i,t}(a^t)$, and $B_{i,t}(a^t, s_{i,1}, \dots, s_{i,t})$ to model this process as an additively separable value function.

⁵ We do assume that Assumption 2.3 holds throughout the paper, but we state the definition above as a combination of Property 4.1 and Assumption 2.3 in order to clearly state that for an environment to be separable, the value function of each agent must satisfy both a functional and a statistical (independence of first signal) separation.

4.2 The Relaxed Environment and the Virtual Welfare

In the relaxed environment, we can use the standard techniques of static mechanism design (Myerson [1981], Milgrom and Segal [2002]) to establish an upper bound on the profit of the optimal mechanism. The next lemma establishes that the profit of any IC mechanism is an ‘affine transformation’ of the social welfare of the agents. The affine factors are given by the functions α and β in the lemma. Note that they only depend on the initial signals (and the actions of the mechanism) and do not explicitly depend on the signals from $t \geq 1$. This observation underlies our construction of the optimal mechanism.

Lemma 4.1. *Consider the relaxed environment and an incentive compatible mechanism \mathcal{M} . Suppose the environment is separable (as in Definition 4.1), and A_i , $B_{i,t}$ and $C_{i,t}$ are uniformly bounded. Then, under the stochastic process induced by \mathcal{M} and the truthful reporting strategy, the expected discounted sum of payments by each agent i is equal to*

$$\mathbb{E} \left[\sum_{t=0}^{\infty} \delta^t p_{i,t} \right] = \mathbb{E} \left[\sum_{t=0}^{\infty} \delta^t (\alpha_i(s_{i,0}) v_{i,t}(a^t, s_i^t) + \beta_{i,t}(a^t, s_{i,0})) \right] - \mathbb{E} \left[U_i^{\mathcal{M}, \mathcal{T}}(s_{i,0} = 0, s_{-i,0}) \right] \quad (15)$$

where the functions α_i and $\beta_{i,t}$ are given by:

- For multiplicatively-separable values,

$$\alpha_i(s_{i,0}) = 1 - \frac{1 - F_i(s_{i,0})}{f_i(s_{i,0})} \frac{A'_i(s_{i,0})}{A_i(s_{i,0})} \quad \beta_{i,t}(a^t, s_{i,0}) = 0$$

- For additively-separable values,

$$\alpha_i(s_{i,0}) = 1 \quad \beta_{i,t}(a^t, s_{i,0}) = -\frac{1 - F_i(s_{i,0})}{f_i(s_{i,0})} A'_i(s_{i,0}) C_{i,t}(a^t)$$

The lemma above yields a bound on the profit of the optimal mechanism for the relaxed environment. Recall that Lemma 3.1 established that the profit for the dynamic environment is bounded by the profit from the relaxed one. Combining these two lemmas and the fact that an IR mechanism must satisfy $U_i^{\mathcal{M}, \mathcal{T}}(s_{i,0} = 0) \geq 0$, we obtain the following profit bound.

Corollary 4.1. *Under the assumptions in Lemma 4.1, for both the relaxed and the dynamic environments, the Profit ^{\mathcal{M}} of any incentive compatible and individually rational mechanism \mathcal{M} is bounded as follows:*

$$\text{Profit}^{\mathcal{M}} \leq \max_{q \in \mathcal{Q}} \mathbb{E} \left[\sum_{t=1}^{\infty} \delta^t \left(\sum_{i=1}^n (\alpha_i(s_{i,0}) v_{i,t}(a^t, s_i^t) + \beta_{i,t}(a^t, s_{i,0})) - c_t(a^t) \right) \right], \quad (16)$$

where \mathcal{Q} is the set of all allocation rules.

The bound above determines an upper bound on the profit of any optimal dynamic mechanism. In the next section, we provide a dynamic mechanism that satisfies IC and IR and achieves this upper bound.

5 The VIRTUAL-PIVOT Mechanism

We now present the VIRTUAL-PIVOT mechanism an optimal dynamic mechanism in separable environments.

The key insight from Section 4 is that the profit of a dynamic mechanism is bounded by an affine transformation of the social welfare of the agents, where the affine parameters are given by the functions α_i and $\beta_{i,t}$ in Lemma 4.1.

We define an affine weight function through a pair of vectors $(\hat{\alpha}, \hat{\beta})$, such that $\hat{\alpha} = (\hat{\alpha}_1, \dots, \hat{\alpha}_n) \in \mathbb{R}^n$ and $\hat{\beta} = (\hat{\beta}_1, \dots, \hat{\beta}_n) \in (\mathcal{A} \times \mathbb{R})^n$, where \mathcal{A} includes all possible action vectors a^t for any t . In particular, $\hat{\beta}$ is allowed to depend action a^t , so that $\hat{\beta}(a^t) = (\hat{\beta}_1(a^t), \dots, \hat{\beta}_n(a^t)) \in \mathbb{R}^n$. For any $(\hat{\alpha}, \hat{\beta})$, time t , and vectors of actions a^t and types s^t , the weighted social welfare with respect to $(\hat{\alpha}, \hat{\beta})$ is defined as

$$W^{(\hat{\alpha}, \hat{\beta})}(a^{t-1}, s^t) \triangleq \max_{q \in \mathcal{Q}} \mathbb{E} \left[\sum_{\tau=t}^{\infty} \delta^\tau \left(\sum_{i=1}^n \left(\hat{\alpha}_i v_{i,\tau}(a^\tau, s_i^\tau) + \hat{\beta}_i(a^\tau) \right) - c_\tau(a^\tau) \right) \middle| s^t, a^{t-1} \right], \quad (17)$$

where the max is over all the possible allocation rules. Using a standard dynamic programming argument, the weighted social welfare satisfies the following (Bellman) equations:

$$W^{(\hat{\alpha}, \hat{\beta})}(a^{t-1}, s^t) = \max_{a_t \in \mathcal{A}_t} \mathbb{E} \left[\sum_{i=1}^n \left(\hat{\alpha}_i v_{i,t}(a^t, s_i^t) - \hat{\beta}_i(a^t) \right) - c_t(a^t) + \delta W^{(\hat{\alpha}, \hat{\beta})}(a^t, s^{t+1}) \middle| s^t, a^{t-1} \right] \quad (18)$$

where s_i^{t+1} is the next (random) type when conditioned on s^t and a^t .

Note, however, that the affine parameters $(\hat{\alpha}, \hat{\beta})$ we need to use to achieve the bound from Corollary 4.1 are not numbers (or, in the case of β , functions of the sequence of actions), but functions of the first signal $s_{i,0}$ of each agent i . An important challenge in implementing an IC mechanism is eliciting $s_{i,0}$ in an incentive compatible way in order to obtain the desired $(\hat{\alpha}, \hat{\beta})$. An important design choice in the VIRTUAL-PIVOT Mechanism is to use the first report of $s_{i,0}$ to determine the affine parameters $(\hat{\alpha}, \hat{\beta})$ and maintain those affine parameters fixed for all periods, irrespective of future re-reports of $s_{i,0}$.

The VIRTUAL-PIVOT mechanism is presented in Figure 2. The mechanism consists of two stages:

- (Subscription Phase) At time 0, each agent i , reports her initial type, $\hat{s}_{i,0}$. Then, the mechanism assigns affine parameters $(\hat{\alpha}_i = \alpha_i(\hat{s}_{i,0}), \hat{\beta}_i(\cdot) = \beta_{i,t}(\cdot, \hat{s}_{i,0}))$ to each agent i , where the functions α_i and β_i are given in Lemma 4.1. Then, the mechanism excludes the agents whose expected discounted payments would be negative (or zero). If $p_i^*(\hat{s}_0) \leq 0$ (see definition in Eq. (24)), then $i \notin a_0$. Otherwise, agent $i \in a_0$ and pays $p_{i,0}(\hat{s}_0)$ (see definition in Eq. (25)).
- (Allocation Phase) For $t \geq 1$, the VIRTUAL-PIVOT mechanism is equivalent to an affine dynamic pivot mechanism. The affine parameters are fixed and the mechanism solicits reports from the agents in order to choose actions that maximize the affinely transformed social welfare $W^{(\hat{\alpha}, \hat{\beta})}$.

The VIRTUAL-PIVOT Mechanism:

(Subscription Phase) At time $t = 0$, for each agent i ,

She reports $\hat{s}_{i,0}$.

Let $\hat{\alpha}_i \leftarrow \alpha_i(\hat{s}_{i,0})$, $\hat{\beta}_i(a^\tau) \leftarrow \beta_{i,\tau}(a^\tau, \hat{s}_{i,0})$ for all $\tau \geq 1$ and $a^\tau \in \mathcal{A}_\tau$.

If $p_i^*(\hat{s}_0) \leq 0$ (see Eq. (24)), then $i \notin a_0$ (agent i is excluded).

If $p_i^*(\hat{s}_0) > 0$, then let $i \in a_0$ and charge her $p_{i,0}(\hat{s}_0)$, see Eq. (25).

(Allocation Phase) At each time $t = 1, 2, \dots$

Each agent i reports \hat{s}_i^t .

Let a_t^* be an action that maximizes $W^{(\hat{\alpha}, \hat{\beta})}(a^{*t}, \hat{s}^t)$, see Eq. (19).

Let $m_{i,t}$ be the flow marginal contribution of agent i , see Eq. (21).

The payment of each agent i is equal to $p_{i,t}(\hat{s}_t) \leftarrow v_{i,t}(a^{*t}, \hat{s}_i^t) - \frac{m_{i,t}}{\hat{\alpha}_i}$.

Figure 2: **The VIRTUAL-PIVOT Mechanism**

To gain some intuition, let us consider the multiplicative-separable case. Roughly speaking, an agent with a higher initial signal $s_{i,0}$, would be assigned a larger $\hat{\alpha}_i$. A larger $\hat{\alpha}_i$ increases the weight of the agent in the affine transformation and hence increases the value obtained by the agent.

We discuss the allocation and payment rules in more details in Section 5.2. Before that, we present our main results.

5.1 Optimality

We make the following assumptions.

Assumption 5.1. (*Monotone Hazard Rate*) Assume that $\frac{f_i(s_{i,0})}{1-F_i(s_{i,0})}$ is strictly increasing.

Assumption 5.2. Assume that:

- (*Multiplicative Case*) If the value function of agent i is multiplicatively separable, then $A_i(s_{i,0})$ is strictly increasing, twice differentiable, and concave in $s_{i,0}$.
- (*Additive Case*) If the value function of agent i is additively separable, then, for all $a^t \in \mathcal{A}$, $A_i(a^t)$ is strictly increasing, twice differentiable and concave in $s_{i,0}$. Also, $C_{i,t}(a^t)$ is positive.

The function $A_i(s_{i,0}) = s_{i,0}$ is an example of a function that satisfies Assumption 5.2. These assumptions imply that α_i is strictly increasing for multiplicatively separable value functions and that $\beta_{i,t}$ is differentiable and strictly increasing for additively separable value functions (see Lemma A.2).

Theorem 5.1. *(Optimality) Suppose that the environment is separable and that Assumptions 5.1, and 5.2 hold. Then, the VIRTUAL-PIVOT mechanism is optimal in both the relaxed and the dynamic environments. In addition, the VIRTUAL-PIVOT mechanism is periodic ex-post individually rational and periodic ex-post incentive compatible.*

The proof of this theorem is presented in Subsection 5.3.

The assumptions above allow us to satisfy the dynamic IC condition from Lemma 3.5. For optimality of the mechanism in the relaxed environment, a weaker set of assumptions could potentially be sufficient.

The VIRTUAL-PIVOT Mechanism is optimal for both the relaxed and dynamic environments and the profit obtained by the mechanism, as well as the utility obtained by the agents, are identical in both environments. Therefore, the agents obtain no “information rent” for periods $t \geq 1$. That is, the agents are not able to obtain any benefit from the fact that signals $s_{i,1}, \dots, s_{i,t}$ are private. This no-information-rent property was noted in a two-period model by Eöso and Szentes [2007] where the mechanism is able to control whether or not agents obtain a second private signal. Theorem 5.1 implies that the no-information-rent property holds even in infinite-horizon problems where the seller has partial control (or even no control) over what private signals agents obtain over time (signals evolve according to a stochastic kernel $K_{i,t}(s_{i,t}|a^{t-1}, s_i^{t-1})$), as long as the environment is separable. We show in Section 7 that this property does not extend to general non-separable settings.

Since there is no information rent for periods $t \geq 1$, there is no allocation distortion associated with signals $s_{i,t}$ for $t \geq 1$. The initial signal $s_{i,0}$, however, causes distortion from the efficient allocation at *every* period as if the mechanism design problem was a static one. To see this easily, consider a setting where each agent i has a multiplicatively separable valuation and $A_i(s_{i,0}) = s_{i,0}$, i.e., the value function of agent i is $v_{i,t} = s_{i,0} \times B_{i,t}(a^t, s_{i,1}, \dots, s_{i,t})$. The VIRTUAL-PIVOT Mechanism allocates in order to maximize the “virtual valuations” of

$$\left(s_{i,0} - \frac{1 - F_i(s_{i,0})}{f_i(s_{i,0})} \right) B_{i,t}(a^t, s_{i,1}, \dots, s_{i,t}).$$

That is, the first signal $s_{i,0}$ is replaced at every period by the virtual value $s_{i,0} - \frac{1 - F_i(s_{i,0})}{f_i(s_{i,0})}$ of static mechanism design (see Myerson [1981]). Our results contrast to the ones of Battaglini [2005], Zhang [2011], where the impact of the first signal $s_{i,0}$ on the value $v_{i,t}$ is transient (it disappears as t grows) and, therefore, the allocation distortion is also transient.

5.2 The Allocation and Payment Rules

We first discuss the allocation rule of the mechanism. At each time t , the mechanism chooses allocation a_t^* that maximizes $W^{(\hat{\alpha}, \hat{\beta})}(a^{*t-1}, \hat{s}^t)$ where $a^{*t-1} = (a_0^*, \dots, a_{t-1}^*)$ represents the past

actions of the mechanism. From Eq. (18), we have

$$a_t^* \in \operatorname{argmax}_{\{a_t \in \mathcal{A}_t\}} \left\{ \sum_{i=1}^n \left(\hat{\alpha}_i v_{i,t}((a^{*t-1}, a_t), \hat{s}_i^t) + \hat{\beta}_i(a^{*t-1}, a_t) \right) - c_t(a^{*t-1}, a_t) \right. \\ \left. + \delta \mathbb{E} \left[W^{(\hat{\alpha}, \hat{\beta})}((a^{*t-1}, a_t), s_i^{t+1}) \mid s^t = \hat{s}_t \right] \right\} \quad (19)$$

Note that only reports from two time periods (0 and t) are used to determine a_t^* . That is, \hat{s}_0 is used to determine the affine parameters and \hat{s}_t is used to determine the agents' types at period t . At time t , the mechanism does not use the agents' reports between times 1 to time $t-1$ (for the allocation or payments).

We now show how the payments are determined. We start from the payments $p_{i,t}$ for $t \geq 1$ and then use those to construct $p_{i,0}$. To make the mechanism incentive compatible, $p_{i,t}$ is determined such that the (instantaneous) utility of agent i at time t is proportional to her flow marginal contribution to the affinely transformed social welfare, denoted by $m_{i,t}$.

$$m_{i,t} = W^{(\hat{\alpha}, \hat{\beta})}(a^{*t-1}, \hat{s}^t) - \delta \mathbb{E} \left[W^{(\hat{\alpha}, \hat{\beta})}(a^{*t}, s^{t+1}) \mid s^t = \hat{s}_t, a_t^* \right] \\ - W_{-i}^{(\hat{\alpha}, \hat{\beta})}(a^{*t-1}, \hat{s}^t) + \delta \mathbb{E} \left[W_{-i}^{(\hat{\alpha}, \hat{\beta})}(a_{-i}^{*t}, s^{t+1}) \mid s^t = \hat{s}_t, a^{*t}, a_{-i,t}^* \right] \quad (20)$$

where $W_{-i}^{(a,b)}$ is the affinely transformed social welfare obtained in the absence of agent i

$$W_{-i}^{(\hat{\alpha}, \hat{\beta})}(a^{t-1}, s^t) \triangleq \max_{q \in \mathcal{Q}} \mathbb{E} \left[\sum_{\tau=t}^{\infty} \delta^\tau \left(\sum_{j: j \neq i} \left(\hat{\alpha}_j v_{j,\tau}(a^\tau, s_j^\tau) + \hat{\beta}_j(a^\tau) \right) - c_\tau(a^\tau) \right) \mid s^t, a^{t-1} \right]$$

and $a_{-i,t}^*$ is the action that maximizes $W_{-i}^{(\hat{\alpha}, \hat{\beta})}(a^{*t-1}, s^t)$ at time t .

Equivalently, we have

$$m_{i,t} = \sum_{j=1}^n \left(\hat{\alpha}_j v_j(a^{*t}, \hat{s}^{j,t}) + \hat{\beta}_j(a^{*t}) \right) - c_t(a^{*t}) \\ - W_{-i}^{(\hat{\alpha}, \hat{\beta})}(a^{*t-1}, \hat{s}^t) + \delta \mathbb{E} \left[W_{-i}^{(\hat{\alpha}, \hat{\beta})}(a_{-i}^{*t}, s^{t+1}) \mid s^t = \hat{s}_t, a^{*t}, a_{-i,t}^* \right] \quad (21)$$

The payment by agent i at time t is then given by

$$p_{i,t}(\hat{s}_t) = v_{i,t}(a^{*t}, \hat{s}_i^t) - \frac{m_{i,t}}{\hat{\alpha}_i}. \quad (22)$$

In Bergemann and Välimäki [2010], the idea of such a payment based on flow marginal contributions was introduced and shown to establish incentive compatibility for the welfare maximizing allocation rule. Similarly, the payments that we use (which are scaled versions of the flow marginal contributions) establish incentive compatibility for the affinely transformed welfare maximizing allocation rule.

We now construct the payment at time 0. Consider the allocation rule q^* that maximizes the weighted social welfare conditioned on the reports at time 0, i.e.

$$q^* \in \operatorname{argmax}_{q \in \mathcal{Q}} \mathbb{E} \left[\sum_{t=1}^{\infty} \delta^t \left(\sum_{i=1}^n (\hat{\alpha}_i v_{i,t}(q^t, s_i^t) + \hat{\beta}_i(q^t)) - c_t(q^t) \right) \middle| s_0 = \hat{s}_0 \right] \quad (23)$$

where $q_t = q(h_t, s^t)$ and $q^t = (q_0, \dots, q_t)$. We drop the (explicit) dependence of q_t on h_t and s^t to simplify the presentation. Note that if the agents are truthful, then q^* and a^* correspond to the same allocation rule. Define $p_i^*(\hat{s}_0)$ as follows:

$$p_i^*(\hat{s}_0) = V_i(\hat{s}_0) - \int_0^{\hat{s}_{i,0}} \frac{\partial V_i^{z \rightarrow z}(s_{i,0}, \hat{s}_{-i,0})}{\partial s_{i,0}} \Big|_{s_{i,0}=z} dz \quad (24)$$

where:

$$\frac{\partial V_i^{z \rightarrow z}(s_{i,0}, \hat{s}_{-i,0})}{\partial s_{i,0}} \Big|_{s_{i,0}=z} = \mathbb{E} \left[\sum_{t=1}^{\infty} \delta^t \frac{\partial v_{i,t}(q^{*t}, s_{i,0}, s_{i,1}, \dots, s_{i,t})}{\partial s_{i,0}} \Big|_{s_{i,0}=z} \middle| s_{i,0} = z, s_{-i,0} = \hat{s}_{-i,0} \right].$$

The value $p_i^*(\hat{s}_0)$ is the payment of agent i in the relaxed environment, given by the envelope condition. If $p_i^*(\hat{s}_0) \leq 0$, then the mechanism excludes agent i (that is, $i \notin a_0^*$).

The total expected discounted sum of payments in the relaxed and dynamic environments must match in order to achieve our optimality bound. Therefore, $p_i^*(\hat{s}_0)$ must be equal to expected discounted sum of payments from agent i . Hence, the payment of agent i at time 0 equals

$$p_{i,0}(\hat{s}_0) = p_i^*(\hat{s}_0) - \mathbb{E} \left[\sum_{t=1}^{\infty} \delta^t p_{i,t}(s_i^t) \middle| s_0 = \hat{s}_0 \right]. \quad (25)$$

5.3 Un-Relaxing: Proof of Theorem 5.1

In this subsection, we present the three steps of the proof of Theorem 5.1. The proofs of the following lemmas are given in the appendix.

The first step is to show that the mechanism, if incentive compatible, does indeed yield the profit from the upper bound in Corollary 4.1. The argument used to prove this lemma is a standard one from Myerson [1981]. We also show that the VIRTUAL-PIVOT Mechanism is periodic ex-post individually rational.

Lemma 5.1. *If the VIRTUAL-PIVOT mechanism is incentive compatible, then it is optimal. Moreover, it is periodic ex-post individually rational at $t = 0$.*

The lemma below guarantees that, under the VIRTUAL-PIVOT mechanism, it is always a best response for agents to report their types truthfully regardless of the history, at any time $t \geq 1$ (assuming that other agents will be truthful in the future but not necessarily in the past). This lemma follows the technique of Bergemann and Välimäki [2010], except that it maximizes an affine transformation of the social welfare, instead of the social welfare itself.

Lemma 5.2. *The VIRTUAL-PIVOT mechanism is periodic ex-post incentive compatible and periodic ex-post individually rational for all periods $t \geq 1$.*

The lemma above not only rules out deviations at periods $t \geq 1$, but it also rules out combined deviations at period $t = 0$ and future periods. That’s because if an agent deviates at period 0, she still wants to truthfully report her type at future period (the mechanism is periodic ex-post IC).

Therefore, we need only concern ourselves with period $t = 0$ deviations from the truthful strategy. The proof of Theorem 5.1 is completed by the following lemma.

Lemma 5.3. *Suppose the assumption of Theorem 5.1 hold. Then the VIRTUAL-PIVOT mechanism satisfies the conditions provided by Lemma 3.5 (i.e., Eqs. (11) and (12)). These conditions are satisfied for all agents conditioned on any initial type $s_{-i,0}$ of the other agents and, therefore, the mechanism is periodic ex-post incentive compatible.*

This is a key technical result in our paper. Proving this lemma involves addressing the key difference between the dynamic and the static setting, as we explicitly show the conditions of Lemma 3.5 hold. The separability assumption is central here.

6 Special Cases of the VIRTUAL-PIVOT Mechanism

In this section, we show that the VIRTUAL-PIVOT mechanism can be simply implemented in some natural special cases where it enjoys additional guarantees. First, we present an indirect implementation of the mechanism in an environment with a single agent. Then, we look at environments where the evolution of the types of the agents is either fully dependent or fully independent of the actions of the mechanism.

6.1 The OPTIMAL-CONTRACTING Mechanism for a Single Agent

We now consider the case where there is only a single agent. In this case, the optimal mechanism can be implemented as remarkably simple *indirect* mechanism.

In particular, the indirect OPTIMAL-CONTRACTING Mechanism is presented in Figure 3. The mechanism works as follows. The Subscription Phase is the only period at which the agent ever makes a report of her type. In particular, the agent just makes a report \hat{s}_0 of s_0 .⁶ In the Posted-Price-Phase, the mechanism simply posts a price for every possible action; the agent decides upon the action; the agent pays the respective price for this action; the mechanism executes this chosen action. These prices may vary as a function of time, as they depend on her previous purchases. After $t \geq 1$, the mechanism does not solicit reports from the agent.

Corollary 6.1. *Suppose the assumptions of Theorem 5.1 hold and that there is only one agent. Then OPTIMAL-CONTRACTING is an optimal mechanism.*

⁶ Observe that the subscription phase can be implemented in an indirect manner by offering a menu of contracts at time 0. However, for the simplicity of presentation, we assume the agent reports her initial type.

The OPTIMAL-CONTRACTING Mechanism for One Agent

(Subscription Phase) At time $t = 0$,

The agent reports \hat{s}_0 .

If $p^*(\hat{s}_0) \leq 0$, then terminate the process, (see Eq. (24)).

Otherwise, charge the agent $p_0(\hat{s}_0)$ and continue, (see Eq. (25)).

(The Posted Price Phase) At each time $t = 1, 2, \dots$

The mechanism informs the agent of the price of each possible action.

$$p_t(a^t, \hat{s}_0) = \frac{c_t(a^t) - \beta_t(a^t, \hat{s}_0)}{\alpha(\hat{s}_0)}$$

The agent chooses an action a_t , pays $p_t(a^t, \hat{s}_0)$, and the mechanism takes this action.

Figure 3: **The OPTIMAL-CONTRACTING Mechanism for a Single Agent**

In indirect mechanisms, we need to concern ourselves with what equilibrium we are implementing since agents are no longer simply reporting their types. The corollary above refers to the equilibrium where ties are broken as in the VIRTUAL-PIVOT Mechanism.

To observe how simple the OPTIMAL-CONTRACTING Mechanism is, consider a scenario where the mechanism is considering selling a stream of items to an agent. At each time period, the seller has two possible actions: allocate an item to the agent at a production cost $\gamma \geq 0$ or not (at no cost). The agents' valuation is multiplicative separable (hence, $\beta_t(a^t, \hat{s}_0) = 0$).

The OPTIMAL-CONTRACTING Mechanism can be implemented as follows: the seller offers a family of contracts to the agent of the form $(p, M(p))$. The agent either leaves (and the process terminates) or she picks a price p . If she picks a price p she is immediately charged $M(p)$. At every period $t \geq 1$, the agent will be offered to buy the item at the constant price p .

The value $M(p)$ the mechanism selects is

$$M(p) = p_0 \left(\alpha^{-1} \left(\frac{\gamma}{p} \right) \right),$$

for each possible positive value of $p_0(s_0)$. In equilibrium, the agent will either leave (if $p^*(s_0) \leq 0$) or will pick price $p = \frac{\gamma}{\alpha(s_0)}$.

This mechanism is optimal regardless of the value function of the agent, as long as it is multiplicatively separable. Even if the agent's value v_t is increasing or decreasing over time and the seller knows about it, it is still optimal for the seller to offer a family of contracts of the form $(p, M(p))$ which includes a constant price for every item ($t \geq 1$).

6.2 Controlled and Uncontrolled Environments

There are two natural extremes for how the stochastic process of the environment evolves. At one extreme is the fully *uncontrolled* environment, where the evolution of the agents' signals has no dependence on the action taken by the mechanism. Here, we show the VIRTUAL-PIVOT mechanism enjoys a much stronger incentive compatibility notion. At the other extreme is a *multi-armed bandit* process (which can be considered a fully controlled environment). Here, the type of an agent only evolves if the agent was allocated the item (and no evolution occurs otherwise), and the optimal allocation rule has a particularly simple form.

6.2.1 Fully Uncontrolled Environments

Define an *uncontrolled* environment to be one in which the stochastic process of each agent is independent of the actions taken by the mechanism, i.e.,

$$K_{i,t}(s_{i,t}|a^t, s_i^{t-1}) = K_{i,t}(s_{i,t}|s_i^{t-1}).$$

In this environment the allocation rule of the VIRTUAL-PIVOT Mechanism is myopic, in that the mechanism's decision is to maximize the *instantaneous* weighted social welfare (as opposed to considering how this impacts future decisions). In particular, we have that:

$$\begin{aligned} & \operatorname{argmax}_{\{a_t \in \mathcal{A}_t\}} \mathbb{E} \left[\sum_{i=1}^n \left(\hat{\alpha}_i v_{i,t}(a^t, \hat{s}_i^t) + \hat{\beta}_i(a^t) \right) - c_t(a^t) \middle| \hat{s}^t, a^{t-1} \right] \\ &= \operatorname{argmax}_{\{a_t \in \mathcal{A}_t\}} \mathbb{E} \left[\sum_{i=1}^n \left(\hat{\alpha}_i v_{i,t}(a^t, \hat{s}_i^t) + \hat{\beta}_i(a^t) \right) - c_t(a^t) + \delta W^{(\hat{\alpha}, \hat{\beta})}(a^t, \hat{s}^{t+1}) \middle| \hat{s}^t, a^{t-1} \right] \end{aligned}$$

This is a straightforward corollary of the uncontrolled assumption. ¹

Corollary 6.2. (*A Dominant Strategy IC*) *Suppose the assumptions of Theorem 5.1 hold. The VIRTUAL-PIVOT Mechanism has the property that for every timestep $t \geq 1$, (e.g. after timestep $t = 0$), the truthful reporting strategy is a dominant strategy.*

This guarantee is immediate since each allocation from $t \geq 1$ is just instantly maximizing a social welfare function (and the action taken by the mechanism and the reports provided by the agents have no effect on the future evolution of signals). Hence, periodic ex-post IC for periods $t \geq 1$ immediately implies ex-post IC (and, hence, dominant strategies implementation) for periods $t \geq 1$. Hence, if agents knew their own (and other agents') past, present and future signals, they would still report truthfully at all histories after $t = 0$. Note, however, that at period $t = 0$, the mechanism is still periodic ex-post IC (not ex-post IC).

6.2.2 Fully Controlled (Multi-Armed Bandit) Environments

We now consider the setting where there is only one item to sell every round — so the action space for the mechanism at each period $t \geq 1$ consists of choosing which agent should receive the item

(or choosing not to allocate the item). The environment now considered is one where the type of an agent changes evolves only if the mechanism takes an action. Namely, the type of an agent only changes when the mechanism allocates the item to the agent. We call this environment *controlled*; the underlying stochastic process corresponds to multi-armed bandits where each arm is mapped to an agent. A special case of this setting is also considered in Pavan et al. [2009], in what they call “bandit auctions”. In their model, the types of agents evolve, in an additive manner, according to restricted multi-armed bandits processes.

In a multi-armed bandit process, there is a “state” of each arm and this only evolves if the arm was “pulled”. In our setting, *fully controlled* environment is one where if on any round $t - 1$ where agent i is not allocated the item, the signal $s_{i,t}$ is irrelevant. Precisely, we have that if i is not allocated at time $t - 1$, then we have that: 1) all current and future values do not depend on $s_{i,t}$. 2) the distribution of all future signals are independent of $s_{i,t}$. We also assume, for simplicity, that there are no costs associated with actions in the fully controlled setting.

An notable feature of this environment is that the optimal allocation is an index-based policy (a Gittins-type index, see Gittins [1989], Whittle [1982]). Namely, we can assign a number to each agent, independent of the other agents, and the optimal allocation rule is to give the item to the agent with the highest positive index. In the fully controlled environment, the optimal allocation can be implemented using virtual indices.

Definition 6.1 (Virtual Index). *For each agent i , the virtual index is defined as:*

$$\mathcal{G}_i^{(\hat{\alpha}, \hat{\beta})}(s_{i,t}) = \max_{\tau_i} \mathbb{E} \left[\frac{\sum_{t=t'}^{\tau_i} \delta^t \left(\hat{\alpha}_i v_{i,t'}(a^{t'}, s_{i}^{t'}) + \hat{\beta}_i(a^{t'}) \right)}{\sum_{t=t'}^{\tau_i} \delta^t} \middle| s_{i,t} \right] \quad (26)$$

where the maximum is taken over all stopping times τ_i .

The optimal allocation rule is to give the item to the agent with the highest positive virtual index. The virtual index can be computed individually for each agent and, therefore, it decouples the n -agent problem into n single-agent problems.

The payments, however, cannot be computed separately for each agent as they dependent on the externalities created by the agent receiving an item. The agents who do not receive an item at time t do not cause externalities and, therefore, do not make payments at time t (other than time $t = 0$). For the agent i that does get the item at time t , $W_{-i}^{(\hat{\alpha}, \hat{\beta})}(a^{*t}, \hat{s}^t) = \mathbb{E} \left[W_{-i}^{(\hat{\alpha}, \hat{\beta})}(\hat{s}^{t+1}) \middle| a^{*t} \right]$. Hence, we obtain the following corollary.

Corollary 6.3. *(The Virtual Index Mechanism) Consider the fully controlled environment defined above and suppose the assumptions of Theorem 5.1 hold. The allocation rule of the VIRTUAL-PIVOT Mechanism is to simply allocate to the agent with the highest virtual index. Moreover, for $t \geq 1$,*

$$p_{i,t}(\hat{s}^t) = \frac{1}{\hat{\alpha}_i} \left((1 - \delta) W_{-i}^{(\hat{\alpha}, \hat{\beta})}(a^{*t}, \hat{s}^t) - \hat{\beta}_i(a^{*t}) \right).$$

To gain some intuition, consider the multiplicative-separable case. An agent with a higher initial type $s_{i,0}$, would be assigned a larger $\hat{\alpha}_i$. A larger $\hat{\alpha}_i$ increases agent i 's virtual index and, therefore,

increases the expected discounted value that agent i obtains. Moreover, she pays a lower payment at each period $t \geq 1$. However, for these privileges, she will be required to make a higher upfront payment (at $t = 0$).

7 Limitations of the Relaxation Approach

In this section, we provide examples where the optimal mechanisms in the dynamic and relaxed environments obtain different revenues. Our first example shows that if s_1 is correlated with the future signals then the relaxation approach may fail. This suggests that our assumption of independence of s_1 (which is without loss of generality) is potentially a natural representation. Our second example provides a simple, yet non-separable, value function in which the relaxation approach fails.

The examples are two period environments with one agent (e.g. future values can be considered to be 0 and we can set $\delta = 1$ without loss of generality). The agent receives signals s_0 and s_1 at times 0 and 1. At the end of the period $t = 1$, the mechanism takes an action $a \in \{0, 1\}$, corresponding to an allocation of an item. The agent obtains a value of $a \times v(s_0, s_1)$ – no value is obtained at $t = 0$.

Correlated Signals

Suppose the value of the agent is equal to her second signal, namely, $v(s_0, s_1) = s_1$. Assume $s_0 \in [0, 1]$ and $s_1 \in [0, 1]$ are correlated. In the relaxed environment, the optimal mechanism is trivial: observe s_1 , and take action $a = 1$, at the price equal to s_1 . Hence, the optimal mechanism extracts the whole surplus which is equal to $\mathbb{E}[s_1]$.

We now show that under weak assumptions, the revenue of any dynamic mechanism that cannot observe the second signal is less than $\mathbb{E}[s_1]$. Consider an incentive compatible and individually rational mechanism \mathcal{M} . Note that due to individual rationality constraints, a mechanism cannot extract more revenue than $\mathbb{E}[s_1 \times a^{\mathcal{M}}(s_0, s_1) | s_0] \leq \mathbb{E}[s_1 | s_0]$ from an agent of type s_0 where $a^{\mathcal{M}}(s_0, s_1)$ represents the mechanism's action (i.e., the probability of allocation). Thus, \mathcal{M} can extract a revenue of $\mathbb{E}[s_1]$ only if $a = 1$ with probability 1.

On the other hand, if a mechanism chooses $a = 1$ with probability 1, then the expected payment at time $t = 0$, $\mathbb{E}[p_0 + p_1 | s_0]$, should be identical for all possible first period types s_0 with probability 1, by Lemma 3.2 (if not, then the agent would misreport her type as the type with the minimum expected payment). Hence, the expected payment of the agent is less than or equal to $\inf_{s_0} \mathbb{E}[s_1 | s_0]$. Suppose s_0 and s_1 are correlated such that for an set τ of non-zero measure, if $s_0 \in \tau$ then, $\mathbb{E}[s_1 | s_0] < \mathbb{E}[s_1]$. In this case, the revenue of \mathcal{M} is strictly less than $\mathbb{E}[s_1]$.

Non-separable Value Functions

Now assume s_0 to be uniformly drawn from $[0, 1]$ and let s_1 be drawn independently and uniformly from the set $\{+, \times\}$. The value at time 1 is:

$$\begin{aligned} v(s_0, +) &= s_0 + c_+ \\ v(s_0, \times) &= s_0 c_\times \end{aligned}$$

For all future times, assume the value is 0. Here, we assume c_+ is a constant greater than 1 and we later set c_\times to be a large positive constant.

Note that this value function is of the form:

$$v(s_0, s_1) = A(s_1)s_0 + B(s_1)$$

and does not satisfy our separability assumptions.

We observe that by Eq. (4), there is a unique optimal allocation in the relaxed environment. This optimal allocation corresponds to the two static optimal auctions for the special cases where $s_1 = +$ and $s_1 = \times$. In particular, the allocation for $q(s_0, s_1 = +)$ is one which always allocates (because c_+ is greater than 1). The allocation for $q(s_0, s_1 = \times)$ occurs only if $s_0 \geq 0.5$. This allocation uniquely maximizes Eq. (4), under the assumption that $U(0) = 0$.⁷ To see this, note that for each setting of s_1 , we have a static problem of optimal auction design with one item and one buyer. Furthermore, as the values are 0 at $s_1 = 0$ we have $U(0) = 0$.

It is interesting to note the following rather natural monotonicity properties:

- The value $v(s_0, s_1)$ is monotone (and linear) in s_0 .
- The optimal (relaxed) utility is $U(s_0)$ is monotone in s_0 .
- The future value $V(s_0)$ under the optimal allocation is monotone in s_0 .

Nonetheless, we show that dynamic IC is more stringent and that the optimal revenue in the dynamic environment is less. Let r^* be this optimal revenue in the relaxed environment. Now observe that if r^* is achievable in the dynamic environment, then it must be due to this allocation rule — Eq. (4) also specifies the expected payments in the dynamic environment. As a proof by contradiction, let us suppose that this allocation rule could be implemented in an IC manner in the dynamic environment.

Since the allocation does not change between 0 and 0.5, Lemma 3.2 implies:

$$U(s_0 = 0.5) - U(s_0 = 0) = \frac{1}{2}v(0.5, +) - \frac{1}{2}v(0, +)$$

Hence, the average revenue at s_0 is:

$$\begin{aligned} E[p_0 + p_1 | s_0 = 0] &= V(s_0 = 0) - U(s_0 = 0) \\ &= \frac{1}{2}v(0, +) - U(s_0 = 0) \\ &= \frac{1}{2}v(0.5, +) - U(s_0 = 0.5) \end{aligned}$$

⁷Again, technically, there are a family of maximizers, which agree with probability one. The argument holds for any of these maximizers.

Now consider the misreporting strategy R of using $\hat{s}_0 = 0$ when $s_0 = 0.5$ and then reporting $\hat{s}_1 = \times$ when $s_1 = +$ and reporting $\hat{s}_1 = +$ when $s_1 = \times$. Here, the agent obtains the item when $(s_0, s_1) = (0.5, \times)$ (since $(\hat{s}_0, \hat{s}_1) = (0, +)$ is reported which leads to an allocation). The value under this strategy is:

$$V^R(s_0 = 0.5) = \frac{1}{2}v(0.5, \times)$$

(since with a $1/2$ probability the agent obtains $s_1 = \times$). Also, note that the distribution of misreports \hat{s}_1 is uniform under R , so that the expected payments under R at $s_0 = 0.5$ are identical to those at $s_0 = 0$. Hence,

$$\begin{aligned} U^R(s_0 = 0.5) &= V^R(s_0 = 0.5) - E[p_0 + p_1 | s_0 = 0] \\ &= \frac{1}{2}v(0.5, \times) - \frac{1}{2}v(0.5, +) + U(s_0 = 0.5) \\ &= \frac{1}{2}(0.5c_\times - 0.5 - c_+) + U(s_0 = 0.5) \end{aligned}$$

Thus, for sufficiently large c_\times , we have that this misreporting strategy obtains strictly greater utility than that of the truthful strategy. Furthermore, by a continuity argument, for a neighborhood $[0.5, 0.5 + \epsilon]$ this misreporting strategy will also provide strictly more revenue (since the allocation rule does not change above $s_1 \geq 0.5$). Thus, we have a contradiction — there is a misreporting strategy resulting in strictly greater (unconditional) expected utility.

8 Concluding Remarks

In this work, we propose an optimal dynamic mechanism, the VIRTUAL-PIVOT Mechanism, for separable environments. Our methodology is as follows: we first find a candidate allocation rule by solving the mechanism design problem in a relaxed environment, as is standard in this literature. The key challenge we address is how to find a (dynamic) payment rule that makes this candidate allocation rule incentive compatible. Our solution methodology involves aiming for a bigger goal: finding a payment rule that makes the candidate allocation rule periodic ex-post incentive compatible. We show that this is possible for periods after the initial one if we allow the agent to “re-report” their entire history of signals at each period. In particular, the payment rule we need is constructed by mapping the candidate allocation rule to an affine transformation of the social welfare function. We find necessary and sufficient conditions for incentive compatibility at the initial periods for mechanisms that satisfy periodic ex-post incentive compatibility for periods after the first one. Finally, we show that the VIRTUAL-PIVOT Mechanism satisfies these conditions and is, therefore, incentive compatible.

The VIRTUAL-PIVOT mechanism is quite simple and could be implemented in settings such as selling online advertisement (see Section 4). The variant of this mechanism specialized to one-buyer settings, the OPTIMAL-CONTRACTING Mechanism, is even simpler and shows that the structure of the optimal mechanism can be quite counterintuitive.

We show in Section 7 that this relaxation approach will not work in designing optimal mechanisms for general non-separable settings. The precise extent to which our technique works in non-separable

settings and what methodology could be used in designing optimal mechanisms when the relaxation method fails are promising areas for future research.

References

- Susan Athey and Ilya Segal. An efficient dynamic mechanism. *Working paper*, 2007.
- Abishak Bapna and Tomas Weber. Efficient dynamic allocation with uncertain valuations. *Working Paper*, 2008.
- David P. Baron and David Besanko. Regulation and information in a continuing relationship. *Information Economics and Policy*, 1(3):267–302, 1984.
- Marco Battaglini. Long-term contracting with markovian customers. *American Economic Review*, 95(3):637–658, 2005.
- Dirk Bergemann and Maher Said. Dynamic auctions: A survey. *Wiley Encyclopedia of Operations Research and Management Science*, 2011.
- Dirk Bergemann and Juuso Välimäki. The dynamic pivot mechanism. *Econometrica*, 78:771–789, 2010.
- Ruggiero Cavallo, David C. Parkes, and Satinder Singh. Efficient online mechanisms for persistent, periodically inaccessible self-interested agents. *Working paper*, 2007.
- Pascal Courty and Hao Li. Sequential screening. *Review of Economic Studies*, 2000.
- Rahul Deb. Optimal contracting of new experience goods. *Working paper*, 2008.
- Peter  eso and Balazs Szentes. Optimal information disclosure in auctions and the handicap auction. *Review of Economic Studies*, 74(3):705–731, 2007.
- Alex Gershkov and Benny Moldovanu. Dynamic revenue maximization with heterogeneous objects: A mechanism design approach. *American Economic Journal: Microeconomics*, 1(2):168–198, 2009.
- J. C. Gittins. *Allocation Indices for Multi-Armed Bandits*. Wiley, London, 1989.
- Paul Milgrom and Ilya Segal. Envelope theorems for arbitrary choice sets. *Econometrica*, 70(2): 583–601, 2002.
- Roger Myerson. Optimal auction design. *Mathematics of Operations Research*, 1981.
- Roger Myerson. Multistage games with communications. *Econometrica*, 54(2):323–358, 1986.
- Hamid Nazerzadeh, Amin Saberi, and Rakesh Vohra. Dynamic cost-per-action mechanisms and applications to online advertising. In *Proceedings of the 17th International Conference on World Wide Web*, pages 179–188, 2008.
- Mallesh Pai and Rakesh Vohra. Optimal dynamic auctions. *Working paper*, 2008.

- David Parkes. Online mechanisms. *Algorithmic Game Theory (Nisan et al. eds.)*, 2007.
- Alessandro Pavan, Ilya Segal, and Juuso Toikka. Dynamic mechanism design: Incentive compatibility, profit maximization and information disclosure. *Working paper*, 2008.
- Alessandro Pavan, Ilya Segal, and Juuso Toikka. Infinite-horizon mechanism design: The independent-shock approach. *Preliminary Draft*, 2009.
- Maher Said. Auctions with dynamic populations: Efficiency and revenue maximization. *Working paper*, 2009.
- Andrzej Skrzypacz and Simon Board. Optimal dynamic auctions for durable goods: Posted prices and fire-sales. *Working paper*, 2010.
- Gustavo Vulcano, Garrett van Ryzin, and Costis Maglaras. Optimal dynamic auctions for revenue management. *Management Science*, 48(11):1388–1407, 2002.
- Peter Whittle. *Optimization Over Time*, volume 1. Wiley, Chichester, 1982.
- Hao Zhang. Structural analysis of a dynamic adverse-selection model. *Working paper*, 2011.

A Appendix

A.1 Proofs for the Relaxed Environment

Lemma A.1. *For any reporting strategy $y \rightarrow z$ and initial type x , the partial derivative of the expected value of agent i $V_i^{y \rightarrow z}(x)$ (see definition in Eq. (6)) with respect to x exists and is:*

$$\frac{\partial V^{y \rightarrow z}(x)}{\partial x} = \mathbb{E} \left[\sum_{t=0}^{\infty} \delta^t \frac{\partial}{\partial s_{i,0}} v_{i,t}(a^t, s_{i,0}, s_{i,1}, \dots, s_{i,t}) \Big|_{s_{i,0}=x} \right]$$

(where the expectation is under $y \rightarrow z$ and \mathcal{T}_{-i}). Furthermore, it is bounded by

$$\left| \frac{\partial V_i^{y \rightarrow z}(x)}{\partial x} \right| \leq \frac{\bar{V}}{1 - \delta}.$$

Proof. From Assumption 2.2, we have that for all i, t, a, x and $s_{i,1}, \dots, s_{i,t}$,

$$\left| \frac{\partial}{\partial x} v_{i,t}(a^t, x, s_{i,1}, \dots, s_{i,t}) \right| \leq \bar{V} < \infty.$$

Therefore, by Lebesgue’s Dominated Convergence Theorem, the partial derivative $\frac{\partial \bar{V}^{y \rightarrow z}(x)}{\partial x}$ exists,

$$\frac{\partial V^{y \rightarrow z}(x)}{\partial x} = \frac{\partial}{\partial x} \mathbb{E} \left[\sum_{t=0}^{\infty} \delta^t v_{i,t}(a^t, x, s_{i,1}, \dots, s_{i,t}) \right] = \mathbb{E} \left[\sum_{t=0}^{\infty} \delta^t \frac{\partial}{\partial x} v_{i,t}(a^t, x, s_{i,1}, \dots, s_{i,t}) \right]$$

and $\left| \frac{\partial V^{y \rightarrow z}(x)}{\partial x} \right| \leq \mathbb{E} \left[\sum_{t=0}^{\infty} \delta^t \bar{V} \right] = \frac{\bar{V}}{1 - \delta}$. □

Proof of Lemma 3.1. Any strategy available to the agents in the relaxed environment is a feasible strategy in the dynamic environment. Therefore, if all other agents are truthful, any profitable deviation from the truthful strategy in the relaxed environment implies a profitable deviation in the dynamic environment. Since no such profitable deviations exist in the dynamic environment, we obtain that the mechanism \mathcal{M} is incentive compatible in the relaxed environment. Therefore, the optimal revenue in the relaxed environment provides an upper bound on the revenue in the dynamic environment. \square

Proof of Lemma 3.2. For consistency with the notation used in the rest of the paper, we represent the utility of agent i with initial type $s_{i,0} = z'$ and reporting his initial type as $\hat{s}_{i,0} = z$ by $U_i^{z \rightarrow z}(z')$, assuming all other agents are truthful. Respectively, $V_i^{z \rightarrow z}(z')$ and $P_i^{z \rightarrow z}(z')$ represent the expected discounted value and payment of agent i under initial type z' and reported initial type z' .

The expected utility of agent i under reporting strategy $z \rightarrow z$ and initial type x is

$$U_i^{z \rightarrow z}(z) = V_i^{z \rightarrow z}(z) - P_i^{z \rightarrow z}(z). \quad (27)$$

Under the same reporting strategy $z \rightarrow z$, but under initial type z' , the utility of agent i is

$$U_i^{z \rightarrow z}(z') = V_i^{z \rightarrow z}(z') - P_i^{z \rightarrow z}(z'). \quad (28)$$

The payments are functions only of reported types, not true types, and therefore, $P_i^{z \rightarrow z}(z) = P_i^{z \rightarrow z}(z')$. Therefore, for any $z \neq z'$, combining Eqs. (27) and (28) yields

$$\frac{U_i^{z \rightarrow z}(z) - U_i^{z \rightarrow z}(z')}{z - z'} = \frac{V_i^{z \rightarrow z}(z) - V_i^{z \rightarrow z}(z')}{z - z'}.$$

At the same time, if $z > z'$, incentive compatibility yields $U_i^{z' \rightarrow z'}(z') \geq U_i^{z \rightarrow z}(z')$, hence

$$\frac{U_i^{z \rightarrow z}(z) - U_i^{z' \rightarrow z'}(z')}{z - z'} \leq \frac{U_i^{z \rightarrow z}(z) - U_i^{z \rightarrow z}(z')}{z - z'}.$$

Since the partial derivative $\frac{\partial V_i^{z \rightarrow z}(x)}{\partial x}$ exists for all x (see Lemma A.1), we can take the limit as $z' \uparrow z$ and obtain that the left-hand side derivative of $U_i^{z \rightarrow z}(z)$ satisfies

$$\frac{d_- U_i^{z \rightarrow z}(z)}{dz} \leq \left. \frac{\partial V_i^{z \rightarrow z}(s)}{\partial s} \right|_{s=z}.$$

Using the same argument for $z' > z$, we obtain that the right-hand side derivative of $U_i^{z \rightarrow z}(z)$ satisfies

$$\frac{d_+ U_i^{z \rightarrow z}(z)}{dz} \geq \left. \frac{\partial V_i^{z \rightarrow z}(s)}{\partial s} \right|_{s=z}.$$

Since $\left| \frac{\partial V_i^{z \rightarrow z}(s)}{\partial s} \right|$ is bounded by $\frac{\bar{V}}{1-\delta}$ by Lemma A.1, we get that the absolute value of both the left-hand and right-hand side derivatives of $U_i^{z \rightarrow z}(z)$ are also bounded by $\frac{\bar{V}}{1-\delta}$. The function $U_i^{z \rightarrow z}(z)$ is, therefore, $\frac{\bar{V}}{1-\delta}$ -Lipschitz-continuous and, thus, differentiable almost everywhere. At all points where the derivative exists, $\frac{dU_i^{z \rightarrow z}(z)}{dz} = \left. \frac{\partial V_i^{z \rightarrow z}(s)}{\partial s} \right|_{s=z}$. Therefore, the envelope condition follows:

$$U_i^{x \rightarrow x}(x) - U_i^{x' \rightarrow x'}(x') = \int_{x'}^x \frac{dU_i^{z \rightarrow z}(z)}{dz} dz = \int_{x'}^x \left. \frac{\partial V_i^{z \rightarrow z}(s)}{\partial s} \right|_{s=z} dz. \quad (29)$$

Plugging in the result from Lemma A.1, we obtain the desired result. \square

Proof of Lemma 3.3. For notational convenience, we write:

$$\frac{\partial v_{i,t}(a^t, s_{i,0}, s_{i,1}, \dots, s_{i,t})}{\partial s_{i,0}} \Big|_{s_{i,0}=s_{i,0}} = \frac{\partial v_{i,t}(a^t, s_i^t)}{\partial s_{i,0}}$$

where the s_i^t implicitly depends on the first signal.

Consider first the utility $U_i^{\mathcal{M}}(s)$ of an agent i under an initial type profile s , which is given by

$$U_i^{\mathcal{M}}(s_{i,0}, s_{-i,0}) - U_i^{\mathcal{M}}(0, s_{-i,0}) = \int_0^{s_{i,0}} \mathbb{E} \left[\sum_{t=1}^{\infty} \delta^t \frac{\partial v_{i,t}(a^t, s_i^t)}{\partial s_{i,0}} \Big|_{s_{i,0}=z, s_{-i,0}} \right] dz.$$

from Lemma 3.2. Taking the expectation of this term over all possible first period signals $s_{1,0}, \dots, s_{n,0}$, we obtain

$$\mathbb{E}[U_i^{\mathcal{M}}(s_{i,0}, s_{-i,0}) - U_i^{\mathcal{M}}(0, s_{-i,0})] = \int_0^1 \left(\int_0^{s_{i,0}} \mathbb{E} \left[\sum_{t=1}^{\infty} \delta^t \frac{\partial v_{i,t}(a^t, s_i^t)}{\partial s_{i,0}} \Big|_{s_{i,0}=z} \right] dz \right) f_i(s_{i,0}) ds_{i,0}.$$

Inverting the order of integration,

$$\begin{aligned} \mathbb{E}[U_i^{\mathcal{M}}(s_{i,0}, s_{-i,0}) - U_i^{\mathcal{M}}(0, s_{-i,0})] &= \int_0^1 \int_z^1 \mathbb{E} \left[\sum_{t=1}^{\infty} \delta^t \frac{\partial v_{i,t}(a^t, s_i^t)}{\partial s_{i,0}} \Big|_{s_{i,0}=z} \right] f_i(s_{i,0}) ds_{i,0} dz \\ &= \int_0^1 \mathbb{E} \left[\sum_{t=1}^{\infty} \delta^t \frac{\partial v_{i,t}(a^t, s_i^t)}{\partial s_{i,0}} \Big|_{s_{i,0}=z} \right] (1 - F_i(z)) dz. \end{aligned}$$

By multiplying and dividing the right-hand side of the equation above by the density $f_i(z)$ we obtain an unconditional expectation,

$$\mathbb{E}[U_i^{\mathcal{M}}(s_{i,0}, s_{-i,0}) - U_i^{\mathcal{M}}(0, s_{-i,0})] = \mathbb{E} \left[\sum_{t=1}^{\infty} \delta^t \frac{1 - F_i(s_{i,0})}{f_i(s_{i,0})} \frac{\partial v_{i,t}(a^t, s_i^t)}{\partial s_{i,0}} \right].$$

Now note that the discounted sum of payments $\mathbb{E}[\sum_{t=1}^{\infty} \delta^t p_{i,t}]$ is equal to the expected discounted valuation of agent i - $\mathbb{E}[\sum_{t=1}^{\infty} \delta^t v_{i,t}(a^t, s_i^t)]$ - minus her utility, which yields the claim. \square

Proof of Lemma 3.4. The expected utility of agent i under reporting strategy $x' \rightarrow z$ and initial type z is

$$U_i^{x' \rightarrow z}(z) = V_i^{x' \rightarrow z}(z) - P_i^{x' \rightarrow z}(z), \quad (30)$$

where $P_i^{x' \rightarrow z}(z)$ is the expected discounted sum of payments of agent i under reporting strategy $x' \rightarrow z$ and initial type z (see similar definitions of $U_i^{x' \rightarrow z}(z)$ and $V_i^{x' \rightarrow z}(z)$ in Eqs. (5) and (6)). Under the same reporting strategy $x' \rightarrow z$, but under initial type z' , the utility of agent i is

$$U_i^{x' \rightarrow z}(z') = V_i^{x' \rightarrow z}(z') - P_i^{x' \rightarrow z}(z'). \quad (31)$$

The payments are functions only of reported types, not true types, and therefore, $P_i^{x' \rightarrow z}(z) = P_i^{x' \rightarrow z}(z')$. Therefore, for any $z \neq z'$, combining Eqs. (30) and (31) yields

$$\frac{U_i^{x' \rightarrow z}(z) - U_i^{x' \rightarrow z}(z')}{z - z'} = \frac{V_i^{x' \rightarrow z}(z) - V_i^{x' \rightarrow z}(z')}{z - z'}.$$

Periodic ex-post IC guarantees that $U_i^{x' \rightarrow z'}(z') \geq U_i^{x' \rightarrow z}(z')$. Therefore, for any $z > z'$,

$$\frac{U_i^{x' \rightarrow z}(z) - U_i^{x' \rightarrow z'}(z')}{z - z'} \leq \frac{U_i^{x' \rightarrow z}(z) - U_i^{x' \rightarrow z}(z')}{z - z'}.$$

Since the partial derivative $\frac{\partial V_i^{z \rightarrow z}(x)}{\partial x}$ exists for all x (see Lemma A.1), we can take the limit as $z' \uparrow z$ and obtain that the left-hand side derivative of $U_i^{x' \rightarrow z}(z)$ for any constant x' satisfies

$$\frac{d_- U_i^{x' \rightarrow z}(z)}{dz} \leq \left. \frac{\partial V_i^{x' \rightarrow z}(s)}{\partial s} \right|_{s=z}.$$

Using the same argument for $z' > z$, we obtain that the right-hand side derivative of $U_i^{x' \rightarrow z}(z)$ satisfies

$$\frac{d_+ U_i^{x' \rightarrow z}(z)}{dz} \geq \left. \frac{\partial V_i^{x' \rightarrow z}(s)}{\partial s} \right|_{s=z}.$$

Since $\left| \frac{\partial V_i^{x' \rightarrow z}(s)}{\partial s} \right|$ is bounded by $\frac{\bar{V}}{1-\delta}$ by Lemma A.1, we get that the absolute value of both the left-hand and right-hand side derivatives of $U_i^{x' \rightarrow z}(z)$ are also bounded by $\frac{\bar{V}}{1-\delta}$. The function $U_i^{x' \rightarrow z}(z)$ is, therefore, $\frac{\bar{V}}{1-\delta}$ -Lipschitz-continuous and, thus, differentiable almost everywhere. At all points where the derivative exists, $\frac{dU_i^{x' \rightarrow z}(z)}{dz} = \left. \frac{\partial V_i^{x' \rightarrow z}(s)}{\partial s} \right|_{s=z}$. Therefore, the envelope condition follows:

$$U_i^{x' \rightarrow x}(x) - U_i^{x' \rightarrow x'}(x') = \int_{x'}^x \frac{dU_i^{x' \rightarrow z}(z)}{dz} dz = \int_{x'}^x \left. \frac{\partial V_i^{x' \rightarrow z}(s)}{\partial s} \right|_{s=z} dz.$$

□

Proof of Lemma 3.5. The envelope condition from the relaxed environment (see Lemma 3.2) also applies to this setting since a deviation that is feasible in the relaxed environment (that is, using reporting strategy $z \rightarrow z$ for an initial type z') is also feasible in the dynamic environment. Therefore, if the mechanism is incentive compatible, then it satisfies Eq. (29), which is identical to Eq. (11).

To see that IC implies the dynamic monotonicity condition in Eq. (12), simply note that IC is equivalent to Eq. (7) and Eqs. (11) and (9) are respectively equal to the left-hand and the right-hand side of Eq. (7). We thus obtain that IC implies Eq. (12).

We now show that if both Eqs. (11) and (12) hold, then the mechanism is IC. If both equations hold, then for all x and x' ,

$$U_i^{x \rightarrow x}(x) - U_i^{x' \rightarrow x'}(x') = \int_{x'}^x \left. \frac{\partial V_i^{z \rightarrow z}(s)}{\partial s} \right|_{s=z} dz \geq \int_{x'}^x \left. \frac{\partial V_i^{x' \rightarrow z}(s)}{\partial s} \right|_{s=z} dz = U_i^{x' \rightarrow x}(x) - U_i^{x' \rightarrow x'}(x'),$$

where the last equality follows from Lemma 3.4. The equation above is equivalent to IC (see Eq. (7)), when the mechanism is periodic ex-post IC for $t \geq 1$. □

Proof of Lemma 4.1. Observe that for multiplicatively-separable value functions

$$\frac{\partial v_{i,t}(a^t, s_i^t)}{\partial s_{i,0}} = A'_i(s_{i,0}) B_{i,t}(a^t, s_{i,1}, \dots, s_{i,t})$$

and, therefore, Eqs. (15) and (4) are identical. Similarly, for additively-separable functions,

$$\frac{\partial v_{i,t}(a^t, s_i^t)}{\partial s_{i,0}} = A'_i(s_{i,0})C_i(a_t)$$

and, therefore, Eqs. (15) and (4) are again identical. \square

Proof of Corollary 4.1. For an IC mechanism \mathcal{M} , the expected discounted sum of payments by agent i is equal to

$$\mathbb{E} \left[\sum_{t=0}^{\infty} \delta^t p_{i,t} \right] = \mathbb{E} \left[\sum_{t=0}^{\infty} \delta^t (\alpha_i(s_{i,0})v_{i,t}(a^t, s_i^t) + \beta_{i,t}(a^t, s_{i,0})) \right] - \mathbb{E} \left[U_i^{\mathcal{M}, \mathcal{T}}(s_{i,0} = 0) \right]$$

by taking expectations over $s_{-i,0}$ (see Eq. (15)). Since the mechanism satisfies IR, $\mathbb{E} \left[U_i^{\mathcal{M}, \mathcal{T}}(s_{i,0} = 0) \right] \geq 0$ and, therefore,

$$\mathbb{E} \left[\sum_{t=0}^{\infty} \delta^t p_{i,t} \right] \leq \mathbb{E} \left[\sum_{t=0}^{\infty} \delta^t (\alpha_i(s_{i,0})v_{i,t}(a^t, s_i^t) + \beta_{i,t}(a^t, s_{i,0})) \right].$$

The profit of \mathcal{M} is given by the sum of payments minus the cost of actions (see Eq. (2)),

$$\text{Profit}^{\mathcal{M}} \leq \mathbb{E} \left[\sum_{t=1}^{\infty} \delta^t \left(\sum_{i=1}^n (\alpha_i(s_{i,0})v_{i,t}(a^t, s_i^t) + \beta_{i,t}(a^t, s_{i,0})) - c_t(a^t) \right) \right].$$

The bound above is valid for all IC and IR mechanisms. By maximizing over the set of all possible allocation rules (payment rules do not enter the equation above), we obtain the desired result. \square

A.2 Proofs of Section 5.3

Lemma A.2. *Suppose Assumptions 5.1 and 5.2 hold. Then α_i is strictly increasing for multiplicatively separable functions and $\beta_{i,t}$ is strictly increasing for additively separable functions.*

Proof. For simplicity of nation, let $s = s_{i,0}$. Also, let $\eta_i(s)$ denote the hazard rate, i.e.,

$$\eta_i(s) = \frac{f_i(s)}{1 - F_i(s)}.$$

In the additive case,

$$\frac{\partial \beta_{i,t}(a^t, s)}{\partial s} = \frac{\eta'_i(s)}{\eta_i^2(s)} A'_i(s) C_{i,t}(a^t) - \frac{1}{\eta_i(s)} A''_i(s) C_{i,t}(a^t)$$

where $(\cdot)'$ denotes a partial derivative with respect to s . By the assumptions that A_i is concave and strictly increasing, and the hazard rate is positive and strictly increasing, we have that the above has the same sign as $C_{i,t}$. In the multiplicative case, first note that $\alpha_i(s) = 1 - \frac{1}{\eta_i(s)} (\log A_i(s))'$. Therefore,

$$\alpha'_i(s) = \frac{\eta'_i(s)}{\eta_i^2(s)} \frac{A'_i(s)}{A_i(s)} - \frac{1}{\eta_i(s)} (\log A_i(s))''$$

which is positive by the assumption. \square

Proof of Lemma 5.1. If agents are truthful, by Eq. (24), the expected payment of each agent i given $s_{i,0}$ is equal to $\max\{p_i^*(s_{i,0}), 0\}$, where 0 occurs if agent i is excluded from the system ($i \notin a_i$). Namely,

$$p_i^*(\hat{s}_0) = V(s_i^t) - \int_0^{\hat{s}_{i,0}} \frac{\partial V_i^{z \rightarrow z}(s_{i,0}, \hat{s}_{-i,0})}{\partial s_{i,0}} \Big|_{s_{i,0}=z} dz \quad (32)$$

where

$$\frac{\partial V_i^{z \rightarrow z}(s_{i,0}, \hat{s}_{-i,0})}{\partial s_{i,0}} \Big|_{s_{i,0}=z} = \mathbb{E} \left[\sum_{t=1}^{\infty} \delta^t \frac{\partial v_{i,t}(q^{*t}, s_{i,0}, s_{i,1}, \dots, s_{i,t})}{\partial s_{i,0}} \Big|_{s_{i,0}=z} \Big| s_{i,0} = z, s_{-i,0} = \hat{s}_{-i,0} \right]$$

For notational convenience, we write:

$$\frac{\partial v_{i,t}(a^t, s_{i,0}, s_{i,1}, \dots, s_{i,t})}{\partial s_{i,0}} \Big|_{s_{i,0}=s_{i,0}} = \frac{\partial v_{i,t}(a^t, s_i^t)}{\partial s_{i,0}}$$

where the s_i^t implicitly depends on the first signal. The expected payment of agent i is equal to:

$$\begin{aligned} & \int_0^1 \max\{p_i^*(s, s_{0,-i}), 0\} f_i(s) ds \\ &= \int_0^1 \left(\mathbb{E} \left[\sum_{t=1}^{\infty} \delta^t v_{i,t}(q^{*t}, s_i^t) \Big|_{s_{i,0} = s, s_{-i,0}} \right] - \int_0^s \mathbb{E} \left[\sum_{t=1}^{\infty} \delta^t \frac{\partial v_{i,t}(q^{*t}, s_i^t)}{\partial s_{i,0}} \Big|_{s_{i,0} = z, s_{-i,0}} \right] dz \right) f_i(s) ds, \end{aligned}$$

where we can drop the max with zero since the agent obtains value zero at all periods when she is excluded from the system. By changing the order of integration, we have

$$\begin{aligned} & \int_0^1 \max\{p_i^*(s, s_{0,-i}), 0\} f_i(s) ds \\ &= \int_0^1 \left(\mathbb{E} \left[\sum_{t=1}^{\infty} \delta^t \left(v_{i,t}(q^{*t}, s_i^t) - \frac{1 - F_i(s)}{f_i(s)} \frac{\partial v_{i,t}(q^{*t}, s_i^t)}{\partial s_{i,0}} \right) \Big|_{s_{i,0} = s, s_{-i,0}} \right] \right) f_i(s) ds \\ &= \int_0^1 \left(\mathbb{E} \left[\sum_{t=1}^{\infty} \delta^t (\alpha_i(s_{i,0}) v_{i,t}(q^{*t}, s_i^t) + \beta_{i,t}(q^{*t}, s_{i,0})) \Big|_{s_{i,0} = s, s_{-i,0}} \right] \right) f_i(s) ds \\ &= \mathbb{E} \left[\sum_{t=1}^{\infty} \delta^t (\alpha_i(s_{i,0}) v_{i,t}(q^{*t}, s_i^t) + \beta_{i,t}(q^{*t}, s_{i,0})) \Big|_{s_{-i,0}} \right] \quad (33) \end{aligned}$$

Therefore, the profit of the mechanism matches the upper-bound provided in Corollary 4.1. Hence, to prove the optimality, it suffices to show that the mechanism is individually rational. By construction, we have the utility of agent i equal to 0 if $s_{i,0} = 0$ for any $s_{-i,0}$. Therefore,

$$U_i(s_0) = \int_0^{s_{i,0}} \mathbb{E} \left[\sum_{t=1}^{\infty} \delta^t \frac{\partial v_{i,t}(q^{*t}, s_{i,0}, s_{i,1}, \dots, s_{i,t})}{\partial s_{i,0}} \Big|_{s_{i,0}=z} \Big| s_{i,0} = z, s_{-i,0} \right] dz.$$

By Assumption 5.2, $\frac{\partial v_{i,t}(q^{*t}, s_{i,0}, s_{i,1}, \dots, s_{i,t})}{\partial s_{i,0}}$ is non-negative. Hence, the mechanism is individually rational. Precisely, periodic ex-post IR at time 0. \square

Proof of Lemma 5.2. Define $u_{i,t}$ to be the instantaneous utility of agent i at time t . We get

$$\begin{aligned}
u_{i,t} &= v_{i,t}(a^{*t}, s_i^t) - p_{i,t} \\
&= (v_{i,t}(a^{*t}, s_i^t) - v_{i,t}(a^{*t}, \hat{s}_i^t)) + \frac{m_{i,t}}{\hat{\alpha}_i} \\
&= v_{i,t}(a^{*t}, s_i^t) + \frac{\hat{\beta}_i(a^{*t})}{\hat{\alpha}_i} \\
&\quad + \frac{1}{\hat{\alpha}_i} \left(\sum_{j \neq i} (\hat{\alpha}_j v_{j,t}(a^{*t}, \hat{s}_j^t)) - c_t(a^{*t}) - W_{-i}^{(\hat{\alpha}, \hat{\beta})}(a^{*t-1}, \hat{s}^t) + \delta \mathbb{E} \left[W_{-i}^{(\hat{\alpha}, \hat{\beta})}(a^{*t}, \hat{s}^{t+1}) \right] \right)
\end{aligned}$$

The last equality follows from Eq. (21). We dropped the conditioning of $W_{-i}^{(\hat{\alpha}, \hat{\beta})}(a^{*t}, \hat{s}^{t+1})$ on $s^t = \hat{s}^t$, a^{*t} , and a_{-i}^{*t} , as it is clear from the context. For ease of notation, let $s = s_0$. Because all agents except i are truthful, we have

$$\begin{aligned}
u_{it} &= \frac{1}{\hat{\alpha}_i} \left(\sum_{j=1}^n (\hat{\alpha}_j v_{j,t}(a^{*t}, s_j^t) + \hat{\beta}_j(a^{*t})) - c_t(a^{*t}) \right. \\
&\quad \left. - W_{-i}^{(\alpha(s), \beta(s))}(a^{*t-1}, s^t) + \delta \mathbb{E} \left[W_{-i}^{(\alpha(s), \beta(s))}(a^{*t}, s^{t+1}) \right] \right)
\end{aligned}$$

If agent i is truthful and other agents are truthful, we have

$$\sum_{t'=t}^{\infty} \delta^{t'} u_{it'} = \frac{1}{\hat{\alpha}_i} \left(W_{-i}^{(\alpha(s), \beta(s))}(a^{*t-1}, s^t) - W_{-i}^{(\alpha(s), \beta(s))}(a^{*t-1}, s^t) \right)$$

Hence, the allocation rule is aligned with the incentive of agent i . She can maximize her utility by reporting truthfully.

Observe that agents with $\hat{\alpha}_i \leq 0$ would have been excluded. Hence, we have $\sum_{t'=t}^{\infty} u_{it'} \geq 0$. Therefore, the mechanism is periodic ex-post IR. \square

Proof of Lemma 5.3. Observe that Eq. (11) is followed from Lemma 5.1 and Eq. (33). To establish Eq. (12), we show that the inequality holds point-wise, i.e., if $x \geq x'$, then

$$\left. \frac{\partial V_i^{x_i \rightarrow x_i}(s)}{\partial s} \right|_{s=x_i} \geq \left. \frac{\partial V_i^{x'_i \rightarrow x_i}(s)}{\partial s} \right|_{s=x_i} \quad (34)$$

By Eq. (10), this is equivalent to

$$\mathbb{E}_{x_i \rightarrow x_i} \left[\sum_{t=0}^{\infty} \delta^t \frac{\partial v_{i,t}(a^{*t}, s_i^t)}{\partial s_{i,0}} \Big|_{s_{i,0}=s} \Big|_{s_{i,0}=x_i} \right] \geq \mathbb{E}_{x'_i \rightarrow x_i} \left[\sum_{t=0}^{\infty} \delta^t \frac{\partial v_{i,t}(a^{*t}, s_i^t)}{\partial s_{i,0}} \Big|_{s_{i,0}=s} \Big|_{s_{i,0}=x_i} \right] \quad (35)$$

where $\mathbb{E}_{x_i \rightarrow x_i}$ is the expectation under the stochastic process determined by agent i reporting according to $x_i \rightarrow x_i$ (while other agents are truthful) and a^{*t} represents the allocation at time t in this case. Similarly, for reporting strategy $x'_i \rightarrow x_i$, we use the notation $\mathbb{E}_{x'_i \rightarrow x_i}$ and represent the allocation at time t by a'^t .

Recall that we have:

$$v_{i,t}(a^t, s_i^t) - \frac{1 - F_i(s_{i,0})}{f_i(s_{i,0})} \frac{\partial v_{i,t}(a^t, s_i^t)}{\partial s_{i,0}} = \alpha_i(s_{i,0})v_{i,t}(a^t, s_i^t) + \beta_{i,t}(a^t, s_{i,0})$$

Hence, we get

$$\frac{\partial v_{i,t}(a^t, s_i^t)}{\partial s_{i,0}} = \frac{f_i(s_{i,0})}{1 - F_i(s_{i,0})} ((1 - \alpha_i(s_{i,0}))v_{i,t}(a^t, s_i^t) - \beta_{i,t}(a^t, s_i^t)) \quad (36)$$

Therefore, by Eq. (36), the inequality below is equivalent to the desired equation, Eq. (34):

$$\begin{aligned} & \mathbb{E}_{x_i \rightarrow x_i} \left[\sum_{t=1}^{\infty} \delta^t ((1 - \alpha_i(x_i))v_{i,t}(a^t, s_i^t) - \beta_{i,t}(a^t, x_i)) \right] \\ & \geq \mathbb{E}_{x'_i \rightarrow x_i} \left[\sum_{t=1}^{\infty} \delta^t ((1 - \alpha_i(x_i))v_{i,t}(a'^t, s_i^t) - \beta_{i,t}(a'^t, x_i)) \right] \end{aligned} \quad (37)$$

In the following we prove the inequality above. For $k \neq i$, define x_k and x'_k to be equal $s_{k,0}$. Because a^* and a' are optimal allocation rules with respect to $(\alpha(x), \beta(x))$ and $(\alpha(x'), \beta(x'))$, we have:

$$\begin{aligned} & \mathbb{E}_{x_i \rightarrow x_i} \left[\sum_{t=1}^{\infty} \delta^t \left(\sum_{j=1}^n (\alpha_j(x_j)v_{j,t}(a^{*t}, s_j^t) + \beta_{j,t}(a^{*t}, x_j)) - c_t(a^{*t}) \right) \right] \\ & \geq \mathbb{E}_{x'_i \rightarrow x_i} \left[\sum_{t=1}^{\infty} \delta^t \left(\sum_{j=1}^n (\alpha_j(x_j)v_{j,t}(a'^t, s_j^t) + \beta_{j,t}(a'^t, x_j)) - c_t(a'^t) \right) \right] \end{aligned}$$

and similarly

$$\begin{aligned} & \mathbb{E}_{x_i \rightarrow x_i} \left[\sum_{t=1}^{\infty} \delta^t \left(\sum_{j=1}^n (\alpha_j(x'_j)v_{j,t}(a^{*t}, s_j^t) + \beta_{j,t}(a^{*t}, x_j)) - c_t(a^{*t}) \right) \right] \\ & \leq \mathbb{E}_{x'_i \rightarrow x_i} \left[\sum_{t=1}^{\infty} \delta^t \left(\sum_{j=1}^n (\alpha_j(x'_j)v_{j,t}(a'^t, s_j^t) + \beta_{j,t}(a'^t, x'_j)) - c_t(a'^t) \right) \right] \end{aligned}$$

Subtracting these inequalities we get:

$$\begin{aligned} & \mathbb{E}_{x_i \rightarrow x_i} \left[\sum_{t=1}^{\infty} \delta^t \sum_{j=1}^n \left((\alpha_j(x_j) - \alpha_j(x'_j))v_{j,t}(a^{*t}, s_j^t) + (\beta_{j,t}(a^{*t}, x_j) - \beta_{j,t}(a^{*t}, x'_j)) \right) \right] \\ & \geq \mathbb{E}_{x'_i \rightarrow x_i} \left[\sum_{t=1}^{\infty} \delta^t \sum_{j=1}^n \left((\alpha_j(x_j) - \alpha_j(x'_j))v_{j,t}(a'^t, s_j^t) + (\beta_{j,t}(a'^t, x_j) - \beta_{j,t}(a'^t, x'_j)) \right) \right] \end{aligned}$$

Because for $k \neq i$, agents are truthful and $x'_k = x_k$, we have

$$\begin{aligned} & \mathbb{E}_{x_i \rightarrow x_i} \left[\sum_{t=1}^{\infty} \delta^t \left((\alpha_i(x_i) - \alpha_i(x'_i))v_{i,t}(a^{*t}, s_i^t) + (\beta_{i,t}(a^{*t}, x_i) - \beta_{i,t}(a^{*t}, x'_i)) \right) \right] \\ & \geq \mathbb{E}_{x'_i \rightarrow x_i} \left[\sum_{t=1}^{\infty} \delta^t \left((\alpha_i(x_i) - \alpha_i(x'_i))v_{i,t}(a'^t, s_i^t) + (\beta_{i,t}(a'^t, x_i) - \beta_{i,t}(a'^t, x'_i)) \right) \right] \end{aligned} \quad (38)$$

Now suppose v_i is multiplicative separable (i.e., $\beta_{i,t}(\cdot, \cdot) = 0$) and Assumption 5.2 holds — we consider the additive valuations later. Because $x \geq x'$, by Assumption 5.2 and Lemma A.2, we have $\alpha_i(x_i) > \alpha_i(x'_i)$; moreover $\alpha_i(x_i)$ is less than 1 for $x \in [0, 1)$. Multiplying both sides of the inequality above by $\frac{1-\alpha_i(x_i)}{\alpha_i(x_i)-\alpha_i(x'_i)}$, yields the following:

$$\mathbb{E}_{x_i \rightarrow x_i} \left[\sum_{t=1}^{\infty} \delta^t (1 - \alpha_i(x_i)) v_{i,t}(a^{*t}, s_i^t) \right] \geq \mathbb{E}_{x'_i \rightarrow x_i} \left[\sum_{t=1}^{\infty} \delta^t (1 - \alpha_i(x_i)) v_{i,t}(a'^t, s_i^t) \right]$$

which is equivalent to Eq. (37) for multiplicative-separable valuations.

Now consider the case of additive-separable value functions. We have $\alpha_i(x) = \alpha_i(x') = 1$. Plugging into Eq. (38) we get

$$\mathbb{E}_{x_i \rightarrow x_i} \left[\sum_{t=1}^{\infty} \delta^t (\beta_{i,t}(a^{*t}, x_i) - \beta_{i,t}(a^{*t}, x'_i)) \right] \geq \mathbb{E}_{x'_i \rightarrow x_i} \left[\sum_{t=1}^{\infty} \delta^t (\beta_{i,t}(a'^t, x_i) - \beta_{i,t}(a'^t, x'_i)) \right]$$

Recall that $\beta_{i,t}(a^t, x_i) = -\frac{1-F_i(x_i)}{f_i(x_i)} A'_i(x_i) C_{i,t}(a^t)$. Because $x \geq x'$, by Assumption 5.2 and Lemma A.2, we have $-\frac{1-F_i(x_i)}{f_i(x_i)} A'_i(x_i) > -\frac{1-F_i(x'_i)}{f_i(x'_i)} A'_i(x'_i)$. By multiplying both sides of the inequality above by $\frac{\frac{1-F_i(x_i)}{f_i(x_i)} A'_i(x_i)}{-\frac{1-F_i(x_i)}{f_i(x_i)} A'_i(x_i) + \frac{1-F_i(x'_i)}{f_i(x'_i)} A'_i(x'_i)}$, we get:

$$-\mathbb{E}_{x_i \rightarrow x_i} \left[\sum_{t=1}^{\infty} \delta^t \beta_{i,t}(a^{*t}, x_i) \right] \geq -\mathbb{E}_{x'_i \rightarrow x_i} \left[\sum_{t=1}^{\infty} \delta^t \beta_{i,t}(a'^t, x'_i) \right]$$

which produces Eq. (37) and, thus, completes the proof. \square

A.3 Proof for the Single Agent Case

Proof of Corollary 6.1. Simply note that under the VIRTUAL-PIVOT Mechanism, if the agent is allocated the item at any time t , the price she pays, under the VIRTUAL-PIVOT Mechanism, is not a function of her report at time t (or any report after $t = 0$). Furthermore, the prices that the agent is charged at $t \geq 1$ are identical to that in the VIRTUAL-PIVOT Mechanism (see Eq. (22)). Also, the prices charged at $t = 0$ is identical to that in the VIRTUAL-PIVOT Mechanism by construction. \square