

Dynamic mechanism design: dynamic arrivals and changing values*

Daniel F. Garrett

Toulouse School of Economics, University of Toulouse Capitole
daniel.garrett@tse-fr.eu

June 12, 2017

Abstract

We study the optimal mechanism in a dynamic sales relationship where the buyer's arrival date is uncertain, and where his value changes stochastically over time. The buyer's arrival date is the first date at which contracting is feasible and is his private information. To induce immediate participation, the buyer is granted positive expected rents even if his value at arrival is the lowest possible. The buyer is punished for arriving late; i.e., he expects to earn less of the surplus. Optimal allocations for a late arriver are also further distorted below first-best levels. Conditions are provided under which allocations converge to the efficient ones long enough after contracting, and this convergence occurs irrespective of the time the contract is initially agreed (put differently, the so-called "principle of vanishing distortions" introduced by Battaglini (2005) continues to apply irrespective of the buyer's arrival date).

JEL Classification: D82

Keywords: dynamic mechanism design, dynamic arrivals, stochastic process

*This paper adapts the second part of my earlier work "Durable Goods Sales with Dynamic Arrivals and Changing Values" to a simpler setting where the good is not durable. The paper has benefited from detailed comments from Simon Board, Rahul Deb, Jeff Ely, Igal Hendel, Konrad Mierendorff, Alessandro Pavan, Bill Rogerson, Bruno Strulovici and Rakesh Vohra. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 714147).

1 Introduction

Markets for most goods are highly dynamic. Buyers may become interested in acquiring goods at different times, such as when they first encounter advertisements for the product. Once in the market, their preferences can be expected to change. Buyers' eagerness to consume often hinges on their own circumstances. Purchasers of cellular telephone plans or wireless internet packages, for instance, have preferences that fluctuate with their available leisure time and contact with friends. Commercial buyers' needs may change in long-term supply relationships. For instance, a restaurant's preferences for acquiring high-quality ingredients from a supplier may vary with changes in its menu, which may come at the whim of the chef.

This paper studies the optimal mechanism for a buyer who has vertical preferences over the quality levels that the seller can supply. The buyer's arrival date to the market (which is the first date he can contract with the seller) is uncertain and, having arrived, his preferences evolve stochastically with time. The key difficulty for designing the profit-maximizing mechanism in such a setting is that the buyer is strategic, and can "hide" his (privately known) arrival to the market. That is, he may participate in the mechanism only at the moment of his choice. In particular, the buyer may prefer to wait to learn if his preferences will change before participating.

That the revelation principle applies in our set-up means that there is never any loss in inducing the buyer to participate in the mechanism immediately upon arrival; in fact, such a policy is typically strictly the most profitable. Because a buyer who has not yet participated in the mechanism retains the ability to wait and participate at a later date, he enjoys a positive outside option. This outside option is *endogenous*, since it depends on the seller's choice of mechanism for later participation dates. An optimal mechanism therefore punishes late participation: If the buyer participates later, then he faces worse terms of trade, purchases quality levels which are distorted further below their efficient levels, and expects to earn less rent. By lowering the option value of waiting, the seller extracts more of the surplus for herself. Our finding thus contrasts with the much simpler case of constant values, where the optimal mechanism involves a repetition of the static optimum, and where the buyer therefore receives the same treatment irrespective of the participation date. Because values are persistent in our setting, how the buyer fares if delaying participation depends on his current value for quality, and this means the value of his outside option is type dependent.¹

The quality levels supplied under a contract signed at a given date τ depend critically on the ratio between the probability of arrival at date τ and the probability of arrival at any earlier date. A smaller ratio implies that the seller cares relatively less about efficiency at date τ and more about limiting the rents available in case of arrival before τ . When the ratio decreases with time, the optimal quality allocations thus become increasingly (downward) distorted at later contracting dates. When the horizon is infinite, and when the buyer arrives at each date with positive probability, the

¹See Jullien (2000) for a study of (static) mechanism design with type-dependent outside options.

ratio necessarily converges to zero with time. The rents the buyer expects for an optimal mechanism then converge to zero as the participation date goes to infinity.

Although the buyer receives lower qualities if he arrives late, it is often still the case that quality prescriptions converge to their first-best levels after a sufficiently long relationship. Put differently, the “principle of vanishing distortions” first described by Battaglini (2005) and adapted to richer settings by Pavan, Segal and Toikka (2014) continues to hold. The reason is that quality choices at dates long after the relationship has commenced have little effect on the information rents that the buyer expects, as perceived at the time of contracting. This is familiar from the existing literature: loosely, the result is driven by an assumption that a buyer’s initial value for quality is a poor predictor of his value far in the future. Of course, many stochastic processes fail such a restriction; existing work (see, e.g., Boleslavsky and Said, 2013, Pavan, Segal and Toikka, 2014, Skrzypacz and Toikka, 2015, and Bergemann and Strack, 2015) shows how distortions can fail to vanish if the type process is sufficiently persistent.

Outline. The rest of the paper is as follows. The remainder of this section provides a review of related literature. Section 2 then introduces a model in which the buyer can have at most two values for quality (building on work by Battaglini, 2005, as explained below). Section 3 provides a detailed analysis of the two-value model and Section 4 concludes. Appendix A provides proofs of all results relating to the two-value model, while Appendix B provides additional results (with proofs) for the model with a continuum of values (which builds, especially, on work by Pavan, Segal and Toikka, 2014).

1.1 Related literature

This paper connects two distinct lines of research in dynamic mechanism design. One strand considers profit-maximizing mechanisms for agents whose preferences evolve stochastically with time and who are available to participate at the date the principal fixes the mechanism (see, e.g., Baron and Besanko (1984), Besanko (1985), Courty and Li (2000), Battaglini (2005), Eso and Szentes (2007), and Pavan, Segal and Toikka (2014)). The other considers dynamic mechanisms when agents arrive over time but preferences do not change (see, e.g., Conlisk, Gerstner and Sobel (1984), Board (2008), Gershkov and Moldovanu (2009), Said (2012), Pai and Vohra (2013) and Board and Skrzypacz (2016)).² While these strands have mainly developed independently (see Bergemann and Said (2011) for a summary), combining features from both is an important step towards realism and allows us to uncover new tradeoffs.

The analysis in the main text focuses on a highly tractable framework that builds on Battaglini (2005). Battaglini studies a dynamic contracting setting with a fixed and commonly known partic-

²There is also a literature with dynamic arrivals but without commitment; examples include Conlisk, Gerstner and Sobel (1984) and Dilme and Li (2016).

ipation date, and where the buyer can have only two values for quality. The buyer's value evolves according to a first-order Markov process. The key novelty in our set-up is that the buyer's arrival date to the market – which is the first date at which a contract can be agreed – is instead uncertain and the buyer's private information. The optimal mechanism must therefore permit participation at each possible arrival date (unlike in Battaglini's model, where the buyer is simply excluded forever if he does not participate at the first opportunity). An important difficulty, new relative to Battaglini's paper, is that design of the mechanism is non-separable across different participation dates. For example, adjustments to how the buyer is treated when participating at any given date typically necessitate adjustments also to the treatment for earlier dates, in order to preserve the buyer's incentive constraints (in particular, to guarantee timely participation in the mechanism).

It is perhaps worth emphasizing that the present paper builds on an earlier working paper, Garrett (2011), which was the first to address dynamic mechanism design with both dynamic arrivals and stochastic evolution of preferences. A key aim of the present paper is to summarize insights from the earlier work, but in a simplified setting.³

Deb and Said (2015) have since provided an analysis of a two-period model, where a unit is allocated in the second period and there is no competition among buyers. Deb and Said solve the case where the seller fully commits (as in the present work), but then focus on relaxing this commitment ability. That the seller allocates a single homogeneous unit to a single buyer simplifies the analysis, facilitating a characterization of the optimal mechanism under quite weak restrictions on the distribution of buyer information and values (in particular, see the full-commitment case). The present paper instead considers variable quality in a repeated Mussa-Rosen framework. The approach in this paper necessarily differs from Deb and Said and appears more readily applicable to a broad class of contracting problems (including multi-agent settings).

Ely, Garrett and Hinnosaar (forthcoming) also provide an analysis of a two-period problem, with allocation of the good in the second period. The focus there, however, is on a restricted "simple" mechanism (where early ticket sales are made at a single price, but auctions are permitted to reallocate capacity). Unlike the two-period settings of Deb and Said and Ely, Garrett and Hinnosaar, the present paper analyzes an arbitrary horizon length. It therefore elucidates how optimal mechanisms evolve when agent participation can take place over longer horizons, and shows how the distortions in optimal mechanisms tend to accumulate over time, so that agents who arrive later receive more distorted allocations. For instance, studying longer horizons allows us to examine issues such as the limiting behavior of mechanisms as the arrival time becomes arbitrarily late, as well as the applicability of the so-called "principle of vanishing distortions" for relationships that have lasted a sufficiently long time.

³The earlier (unpublished) paper considers a durable goods setting. The dynamic optimization problem there is more complex, motivating the simpler environment of the present paper.

Perhaps the most important antecedents to Garrett (2011), and hence the present paper, are Deb (2011, 2014) and Nocke, Peitz and Rosar (2011).⁴ Deb studies a seller's optimal price path in an infinite-horizon setting where a buyer arrives at a fixed date (date zero), and his value changes at a single random time. Deb finds that the optimal price path often features low introductory pricing. Given the restriction to a price path, a buyer often chooses not to participate in the mechanism at first instance; that is, he chooses to delay his purchase decision. The seller may choose a higher price at later dates (after date zero) precisely to deter delay in purchasing. Nocke, Peitz and Rosar study a two-period model where the buyer learns his value only at the second date. They also find that introductory pricing can be optimal (again, this can reflect the seller's aim to punish delayed purchase). The optimal price path in their paper turns out to implement also the optimal mechanism (as chosen without any restrictions). Notably, this means that the seller sometimes finds it optimal to induce participation at a date after the buyer is initially available, although by the revelation principle she could achieve the same outcome by always inducing participation at the initial date. In our setting (as noted above), the seller instead typically finds it *strictly* more profitable to induce buyer participation in the mechanism at the first possible instant. Participation then occurs at different dates in our optimal mechanism only because the buyer's arrival date is random/heterogeneous.

There are still further papers highlighting the value to a seller of deterring delayed purchase. Garrett (2016) studies the optimal price path in an infinite-horizon setting where buyers arrive over time, and where values then change randomly over time. Armstrong and Zhou (2015) study commitments a seller may make to deter buyers from searching for a better product and then returning to purchase. While these papers focus on particular applications and selling formats, the present paper focuses on developing a mechanism design approach that can be applied quite generally in settings where agents arrive over time and have preferences that change randomly.

2 Model and Preliminaries

Basics. We consider a repeated buyer-seller relationship in discrete time, which starts at date $t = 1$ and lasts until the end of period $T \in \{2, \dots, \infty\}$. The buyer values consumption of a non-durable good, which can be provided by the seller in each period. Both buyer and seller have a common discount factor $\delta \leq 1$ ($\delta < 1$ in case $T = \infty$).

⁴Other antecedents include Mierendorff (2016) and Deb and Pai (2013). These papers consider buyers who are heterogeneous not only in their value for the good (which is constant across time), but also in the deadline they face for purchasing it. The optimal mechanism can be implemented as a "biased" auction on a final date, with what Mierendorff terms an (endogenous) "outside option" to purchase the good at a posted price before the auction. Under this implementation, buyers may again face a choice between participating in the mechanism (by purchasing the good at the posted price) or delaying participation until a later date (say participating in the biased auction).

The buyer arrives at some date $\tau \in \{1, \dots, T\}$. This is the first date at which the buyer is available to communicate with the seller; contracting is impossible before this date. Inability to contract at earlier dates may reflect a range of reasons: the buyer may be entirely unaware of the seller's existence before encountering an advertisement which explains the mechanism, or he may be aware of the seller's offer but unable to communicate until the opportunity arises to physically meet.⁵

Payoffs. At each date after arrival, if the buyer purchases a good of quality q at date t , paying p , then he earns a date- t payoff

$$\theta_{\tau,t}q - p$$

where $\theta_{\tau,t} \in \Theta$ is his value at date t , and where Θ is a bounded subset of \mathbb{R}_+ representing the set of possible buyer values for quality. If the buyer does not purchase the good at date t and makes no payment, then he earns a payoff equal to zero in that period.

The seller has a (per-period) cost of producing q units equal to $c(q)$, where $c(\cdot)$ is a continuously-differentiable cost function defined on $[0, \bar{q}]$. The cost function $c(\cdot)$ is strictly increasing, strictly convex, and satisfies $c'(0) = c(0) = 0$ and $c'(\bar{q}) > \sup \Theta$.⁶ (A helpful example of such a cost function puts $c(q) = \frac{q^2}{2}$ up to some bound $\bar{q} > \sup \Theta$; this renders closed-form solutions for optimal qualities in Proposition 1 below.) The seller then earns a period- t payoff from selling q units at price p equal to $p - c(q)$ (the seller's date- t payoff in case making no sale at date t , and receiving no payment, is equal to zero).

Distribution of buyer arrival dates. The probability of arrival at each date τ is $\rho_\tau \geq 0$, with $\sum_{\tau=1}^T \rho_\tau \leq 1$. For each τ , let $\beta_\tau = \sum_{s=1}^{\tau-1} \rho_s$ be the probability that the buyer arrives before date τ . (It is worth emphasizing that, in case $\rho_1 = 1$ and $\rho_t = 0$ for $t \geq 2$, then the model will be identical to that in Battaglini, 2005. As indicated above, the key novelty in our setting will be that the buyer's arrival time is instead random.)

Process for values: In the main text, we consider the case where the buyer has two possible values for the good, leaving the case of a continuum of values to Appendix B. Hence, $\Theta = \{\theta_L, \theta_H\}$, with $0 < \theta_L < \theta_H$ (θ_L is the "low value", while θ_H is the "high value"). We adopt the following notation: if the buyer arrives at date τ , then a sequence of values from date t to $t' > t$, with $t \geq \tau$, is denoted $\theta_{\tau,t}^{t'} = (\theta_{\tau,t}, \dots, \theta_{\tau,t'})$.

The processes we consider satisfy the following restriction, which is particularly important for keeping the seller's problem tractable. The distribution of the buyer's value at each date after his arrival depends only on his value in the previous period, and neither on his earlier values nor on his

⁵Our notion of "arrival" is distinct from other notions that one might be tempted to use, such as the date a buyer first learns his value for the good. See Akan, Ata and Dana (2015) for a model where a buyer learns his value at different dates.

⁶The latter assumption will guarantee that the solutions to the optimal quality schedules that we derive below remain strictly below \bar{q} .

arrival date.⁷ This implies that, at any date t , the period- t value $\theta_{\tau,t}$ is a sufficient statistic for later values.

If the buyer arrives at date τ , then he draws a value $\theta_{\tau,\tau} \in \{\theta_L, \theta_H\}$. The probability of drawing a high date- τ value is given by $\Pr(\tilde{\theta}_{\tau,\tau} = \theta_H) = \mu \in (0, 1)$. Values at each date $t > \tau$ are determined by the transitions $\Pr(\tilde{\theta}_{\tau,t} = \theta_H | \tilde{\theta}_{\tau,t-1} = \theta_L) = \alpha_L \in (0, 1)$ and $\Pr(\tilde{\theta}_{\tau,t} = \theta_H | \tilde{\theta}_{\tau,t-1} = \theta_H) = \alpha_H \in (0, 1)$, with $\alpha_L < \alpha_H$. Thus, a high value at any date implies a greater likelihood of high values at future dates (put differently, the process satisfies the first-order stochastic dominance assumption which is common in the literature).

Mechanisms. Both the buyer's arrival time and the evolution of his value are his private information. The seller can fully commit to a dynamic mechanism. By the revelation principle, we restrict attention to incentive-compatible direct mechanisms. The buyer is asked to report his arrival date τ and initial value $\theta_{\tau,\tau}$, and then to report his subsequent values $\theta_{\tau,t}$ in each period $t > \tau$. If the buyer arrives at date τ , then he can report to the mechanism at any moment from that date onwards.

A mechanism $\Omega = \langle \mathbf{q}, \mathbf{p} \rangle$ is a collection of allocation rules $\mathbf{q} = \langle q_{\tau,t} \rangle_{1 \leq \tau \leq t}$ and payments $\mathbf{p} = \langle p_{\tau,t} \rangle_{1 \leq \tau \leq t}$. If the buyer reports to the mechanism at date τ , and then reports a sequence of values $\hat{\theta}_{\tau,\tau}^t = (\hat{\theta}_{\tau,\tau}, \dots, \hat{\theta}_{\tau,t}) \in \Theta^{t-\tau}$, then he receives the quality $q_{\tau,t}(\hat{\theta}_{\tau,\tau}^t) \in [0, \bar{q}]$ and pays $p_{\tau,t}(\hat{\theta}_{\tau,\tau}^t) \in \mathbb{R}$ at date t . A buyer who reports to the mechanism at date τ is deemed to accept the offer and binds himself to participate at all future dates. A buyer who never participates in the mechanism earns a total discounted payoff equal to zero.

As is the case elsewhere in the literature, given that we impose no cash constraints, our assumption that the buyer can fully commit comes at no loss of generality. Indeed, by appropriately structuring the timing of payments so that the buyer's continuation payoff in the mechanism is never negative, the buyer can always be induced to continue participating at every subsequent date, irrespective of his realized values.

3 Analysis and results

Consider the process defined above where the buyer has two possible values for the good. Fix a mechanism $\Omega = \langle \mathbf{q}, \mathbf{p} \rangle$ and consider a buyer who reports to the mechanism at date τ , makes reports $\hat{\theta}_{\tau,\tau}^{t-1}$ up to date $t-1$ (if any) and has a date- t valuation $\theta_{\tau,t}$. The expected continuation payoff of this buyer if he plans to report truthfully at all future dates is

$$V_{\tau,t}^{\Omega}(\theta_{\tau,t}; \hat{\theta}_{\tau,\tau}^{t-1}) \equiv \mathbb{E} \left[\sum_{s=t}^T \delta^{s-t} \left(\theta_{\tau,s} q_{\tau,s}(\hat{\theta}_{\tau,\tau}^{t-1}, \tilde{\theta}_{\tau,t}^s) - p_{\tau,s}(\hat{\theta}_{\tau,\tau}^{t-1}, \tilde{\theta}_{\tau,t}^s) \right) \mid \tilde{\theta}_{\tau,t} = \theta_{\tau,t} \right].$$

⁷As a result of this restriction, our analysis cannot be directly applied to settings where the evolution of preferences is different for different cohorts of buyers.

Using the same arguments as in Battaglini (2005), we can establish the following useful result concerning how the buyer's continuation payoff at any date t depends on his date- t value. To state it, we introduce the following notation: for any $k \in \{1, \dots, T\}$, $\theta_L^k = (\theta_L, \dots, \theta_L)$ is a sequence of low values of length k .

Lemma 1 (Battaglini, 2005) *Fix an incentive-compatible mechanism Ω , and consider a buyer who first reports at date τ , and then reports a sequence $\hat{\theta}_{\tau, \tau}^{t-1}$ up to date $t - 1$ (or makes no reports in case $t = \tau$). The buyer's expected payoff satisfies*

$$V_{\tau, t}^{\Omega}(\theta_H; \hat{\theta}_{\tau, \tau}^{t-1}) - V_{\tau, t}^{\Omega}(\theta_L; \hat{\theta}_{\tau, \tau}^{t-1}) \geq (\theta_H - \theta_L) \sum_{s=t}^T \delta^{s-t} (\alpha_H - \alpha_L)^{s-t} q_{\tau, s}(\hat{\theta}_{\tau, \tau}^{t-1}, \theta_L^{s-t+1}). \quad (1)$$

Lemma 1 provides a lower bound on the additional payoff the buyer expects when his value is high rather than low at a given date t . One way to interpret the condition is as follows. First, suppose we adjust payments at dates $t + 1$ onwards so that the payoffs satisfy (1) with equality at all such dates, and for all histories. Assuming the new mechanism is incentive compatible at dates $t + 1$ onwards, the buyer is then willing to always report a low value at all such dates. We can then evaluate the buyer's incentive to misreport at date t under the assumption that he always reports a low value in future. A necessary condition for incentive compatibility at date t is then that the buyer expects a payoff from a high value at date t (i.e., $V_{\tau, t}^{\Omega}(\theta_H; \hat{\theta}_{\tau, \tau}^{t-1})$) which exceeds that for a low value (i.e., $V_{\tau, t}^{\Omega}(\theta_L; \hat{\theta}_{\tau, \tau}^{t-1})$) by more than the expected difference in valuations under a strategy of always reporting a low value. The right-hand side of (1) is this difference.

Next, we must go beyond Battaglini's (2005) analysis to deduce a lower bound on $V_{\tau, \tau}^{\Omega}(\theta_L; \emptyset)$ (i.e., on the buyer's payoff when participating at date τ with a low value) in an incentive-compatible and individually-rational mechanism.⁸ Here, we use the following requirement. A buyer who arrives at date τ with a value θ_L must prefer to participate at date τ rather than to delay participation until date $\tau + 1$, then reporting truthfully at all future dates. That is, for all dates τ ,

$$V_{\tau, \tau}^{\Omega}(\theta_L; \emptyset) \geq \delta \left((1 - \alpha_L) V_{\tau+1, \tau+1}^{\Omega}(\theta_L; \emptyset) + \alpha_L V_{\tau+1, \tau+1}^{\Omega}(\theta_H; \emptyset) \right). \quad (2)$$

Note that the simplicity of this participation constraint follows from our Markovian assumption that the buyer's current value is a sufficient statistic for the evolution of his future values. In particular, a buyer who arrives at date τ but participates at date $\tau + 1$ faces the same problem as a buyer who in fact arrives at date $\tau + 1$ (conditional on his date- $\tau + 1$ value being either low or high).

Condition (2), together with the one given in Lemma 1, yields the following result.

⁸In Battaglini's (2005) paper, the only relevant participation date is $\tau = 1$, and the optimal mechanism sets $V_{1,1}^{\Omega}(\theta_L; \emptyset) = 0$.

Lemma 2 Fix an incentive-compatible mechanism $\Omega = \langle \mathbf{q}, \mathbf{p} \rangle$. The expected payoff of a buyer who arrives at date τ with a low value must satisfy

$$V_{\tau,\tau}^{\Omega}(\theta_L; \emptyset) \geq \alpha_L (\theta_H - \theta_L) \sum_{i=1}^{T-\tau} \sum_{s=\tau+i}^T \delta^{s-\tau} (\alpha_H - \alpha_L)^{s-\tau-i} q_{\tau+i,s} (\theta_L^{s-\tau-i+1}). \quad (3)$$

Lemma 2 provides a lower bound on payoffs that will turn out to be tight in the optimal mechanism (under a certain regularity condition to be specified momentarily). To begin understanding this expression, it is simplest to consider the case where T is finite. Since the buyer must be willing to participate if he arrives at date T with a low value, we have $V_{T,T}^{\Omega}(\theta_L; \emptyset) \geq 0$. Now consider the buyer who arrives at date $\tau = T - 1$ with value θ_L . If the buyer chooses not to participate at $T - 1$, then he will have the option to participate at date T with a high value with probability α_L . In this case, he earns a positive rent $V_{T,T}^{\Omega}(\theta_H; \emptyset)$, which is at least $(\theta_H - \theta_L) q_{T,T}(\theta_L)$ by (1). Hence, we have $V_{T-1,T-1}^{\Omega}(\theta_L; \emptyset) \geq \alpha_L \delta (\theta_H - \theta_L) q_{T,T}(\theta_L)$, which is (3) evaluated at $\tau = T - 1$. We can then work recursively backwards to deduce lower bounds on the rents at earlier dates. For instance, to deduce a lower bound $V_{T-2,T-2}^{\Omega}(\theta_L; \emptyset)$, we observe that, if the buyer delays participation until the subsequent period, then he earns a rent $V_{T-1,T-1}^{\Omega}(\theta_L; \emptyset)$ (which is at least $\alpha_L \delta (\theta_H - \theta_L) q_{T,T}(\theta_L)$) in case his value remains low (with probability $1 - \alpha_L$), or this plus an additional rent $V_{T-1,T-1}^{\Omega}(\theta_H; \emptyset) - V_{T-1,T-1}^{\Omega}(\theta_L; \emptyset)$ satisfying (1) if his value turns high (with probability α_L).

Expression (3) is central to our analysis, for it shows how the rent that must be promised to ensure agent participation accumulates with time. When T is finite, the agent, in deciding whether to participate at date T , faces the usual outside option of not purchasing (which has payoff zero). As just explained, ensuring participation at date $T - 1$ requires ceding larger rents because the agent can wait for his value to increase at date T . In turn, this raises the rents that must be promised to ensure participation at $T - 2$.

In terms of characterizing the optimal mechanism, the value of Lemma 2 is that it allows us to find a convenient lower bound on buyer payoffs as a function of the quality allocations \mathbf{q} . In particular, Lemmas 1 and 2 together allow us to provide an upper bound on the achievable profit in an incentive-compatible mechanism, as stated in the next result. This bound coincides with the seller's profits in case all the inequalities in (1) and (3) hold as equalities. This upper bound, analogous to the “virtual surplus” in static mechanism design, turns out to be achievable under a mild condition on the arrival probability, which we describe below.

Lemma 3 Suppose that Ω is an incentive-compatible, individually-rational mechanism implementing an allocation \mathbf{q} . Then expected profits are no greater than

$$\mathbb{E} \left[\sum_{s=\tilde{\tau}}^T \delta^{s-1} \left(m_{\tilde{\tau}}^s \left(\tilde{\theta}_{\tilde{\tau},\tilde{\tau}}^s \right) q_{\tilde{\tau},s} \left(\tilde{\theta}_{\tilde{\tau},\tilde{\tau}}^s \right) - c \left(q_{\tilde{\tau},s} \left(\tilde{\theta}_{\tilde{\tau},\tilde{\tau}}^s \right) \right) \right) \right], \quad (4)$$

where, for all τ , and all $s \geq \tau$,

$$\begin{aligned} m_\tau^s(\theta_L^{s-\tau+1}) &= \theta_L - \left(\frac{\beta_\tau \alpha_L}{\rho_\tau (1-\mu)} + \frac{\mu}{1-\mu} \right) \left(\frac{\alpha_H - \alpha_L}{1 - \alpha_L} \right)^{s-\tau} (\theta_H - \theta_L), \text{ and} \\ m_\tau^s(\theta_{\tau,\tau}^s) &= \theta_{\tau,s} \text{ for all } \theta_{\tau,\tau}^s \neq \theta_L^{s-\tau+1}, \end{aligned} \quad (5)$$

and where expectations are taken over the arrival time $\tilde{\tau}$, as well as the realized values $\tilde{\theta}_{\tilde{\tau},\tilde{\tau}}^T$.

It will be helpful to understand the virtual values m_τ^s in (5) as the surplus due to awarding additional quality at date s to a buyer who arrived at date τ , less the effect on the lowest feasible values of buyer rents. Condition (1) shows that the (lower bound on) additional rent a buyer expects when arriving with a high rather than a low value, i.e. $V_{\tau,\tau}^\Omega(\theta_H; \emptyset) - V_{\tau,\tau}^\Omega(\theta_L; \emptyset)$ for arrival date τ , depends only on the quality at histories where the buyer's value remains low. In turn, the bound on rents in (3) also depends only on the quality at these histories. Hence, at any history $\theta_{\tau,\tau}^s \neq \theta_L^{s-\tau+1}$, the virtual value corresponds simply to the buyer's value for quality $\theta_{\tau,s}$ (note that the same observation is made also in Battaglini's, 2005, paper; although, as noted, that paper considers only a single participation date, say $\tau = 1$).

For any history $\theta_{\tau,\tau}^s$ where the buyer arrives at date τ and his value remains low until s , the virtual value of incremental quality is the buyer's value θ_L less a quantity that can be rewritten as

$$\frac{\beta_\tau \alpha_L (\alpha_H - \alpha_L)^{s-\tau} (\theta_H - \theta_L) + \rho_\tau \mu (\alpha_H - \alpha_L)^{s-\tau} (\theta_H - \theta_L)}{\rho_\tau (1 - \mu) (1 - \alpha_L)^{s-\tau}}. \quad (6)$$

Analogous to the distortion term in the virtual values of static mechanism design, this expression is the ratio of the effect of date- s quality $q_{\tau,s}(\theta_L^{s-\tau+1})$ on buyer rents to the probability this quality is awarded. The second term in the numerator (i.e., $\rho_\tau \mu (\alpha_H - \alpha_L)^{s-\tau} (\theta_H - \theta_L)$) corresponds to the additional expected rents if the buyer happens to arrive at date τ with a high (rather than a low) value, an event which occurs with probability $\rho_\tau \mu$. The first term in the numerator (i.e., $\beta_\tau \alpha_L (\alpha_H - \alpha_L)^{s-\tau} (\theta_H - \theta_L)$) corresponds to the effect of increasing $q_{\tau,s}(\theta_L^{s-\tau+1})$ on the rents earned in case of arrival at date $\tau - 1$ or earlier (the probability of such an arrival time is β_τ , and how much rent the buyer expects at such dates depends on the rate at which a low value turns high, α_L , as explained in relation to Lemma 2). The denominator in (6) (i.e., $\rho_\tau (1 - \mu) (1 - \alpha_L)^{s-\tau}$) is simply the probability that the history $\theta_{\tau,\tau}^s = \theta_L^{s-\tau+1}$ occurs.

One can now choose the qualities which maximize the expression (4), and then verify that these qualities can be implemented as part of an incentive-compatible mechanism. This leads to the following result.

Proposition 1 *Suppose that, for all $\tau \leq T-1$, $m_{\tau+1}^{\tau+1}(\theta_L) \leq m_\tau^{\tau+1}(\theta_L^2)$. Profit-maximizing qualities $q_{\tau,s}^*$ are given, for each arrival date τ , each date $s \geq \tau$, and each $\theta_{\tau,\tau}^s$, by*

$$c'(q_{\tau,s}^*(\theta_{\tau,\tau}^s)) = \max \{m_\tau^s(\theta_{\tau,\tau}^s), 0\}. \quad (7)$$

The proof proceeds by constructing a mechanism with allocations given by (7), such that all of the inequalities in (1) and (3) hold with equality. Buyer rents are then as small as possible for an incentive-compatible individually-rational mechanism implementing these allocations (implying that the upper bound on profits in (4) is attained). Given that the allocations (7) maximize the expression in (4), the mechanism must maximize profits provided it is incentive compatible.⁹

Optimal qualities balance the cost of providing a given quality level against the “virtual value” of provision introduced in Lemma 3. As discussed above, following sequences of low values, virtual values are less than the value to the buyer, capturing ex-ante expected buyer rents.

The allocation $q_{1,t}^*$ which applies when the buyer arrives at date $\tau = 1$ is exactly the allocation that the seller would optimally choose in a problem where the buyer is *known* to arrive at date 1. Hence the allocation for a date-1 arrival is precisely the same as in Battaglini’s (2005) paper, which did not study uncertain arrival times. This result is to be expected, since the allocation for date-1 arrival does not affect the rents that must be left in case of arrival after date 1. The only difference between the mechanism for $\tau = 1$ in the present setting, and the one studied by Battaglini, lies in the prices paid (equivalently, the rent obtained) by the buyer. In the present setting, the buyer’s payments must be lower so that the buyer is willing to participate at date 1 rather than delaying participation.

For arrival at dates $\tau > 1$, the optimal qualities at histories of low values, i.e. $q_{\tau,t}^* (\theta_L^{t-\tau+1})$, are further distorted below first-best values. In particular, $q_{\tau,\tau+k}^* (\theta_L^{k+1}) \leq q_{1,1+k}^* (\theta_L^{k+1})$ for all $\tau \geq 2$ and $k \in \{0, 1, \dots, T - \tau\}$. These additional distortions reflect the seller’s goal of reducing rents in case of arrival at a later date, in turn permitting a reduction of rents in case of earlier arrival (including possible arrival at date 1).

Note that Proposition 1 provides a sufficient condition for the incentive compatibility of our candidate mechanism in terms of the primitives of the problem. In particular, the assumption that $m_{\tau+1}^{\tau+1}(\theta_L) \leq m_{\tau}^{\tau+1}(\theta_L^2)$ is equivalent to

$$\frac{\beta_{\tau+1}}{\rho_{\tau+1}} \frac{\alpha_L}{1-\mu} + \frac{\mu}{1-\mu} \geq \frac{\alpha_H - \alpha_L}{1-\alpha_L} \left(\frac{\beta_{\tau}}{\rho_{\tau}} \frac{\alpha_L}{1-\mu} + \frac{\mu}{1-\mu} \right). \quad (8)$$

This condition guarantees that, for all $\tau \leq T - 1$ and all $s \in \{\tau + 1, \dots, T\}$,

$$q_{\tau,s}^* (\theta_L^{s-\tau+1}) \geq q_{\tau+1,s}^* (\theta_L^{s-\tau}).$$

In other words, the assumption guarantees that a buyer receives a higher quality allocation if he participates in the mechanism one period earlier, even if his values turn out to remain low from

⁹The proof of incentive compatibility can be viewed in two parts: (i) verifying the optimality of truthful reporting of values after participating in the mechanism, and (ii) verifying that the buyer is always willing to participate in the mechanism upon arrival. While the first part is familiar from Battaglini’s work, the second part is new to the present paper.

the date of arrival. This ensures that, if the buyer's value is high, he expects a higher rent from immediate participation than by delaying (that this is true also when the buyer's value is low follows because the mechanism is constructed to satisfy (3); i.e., the buyer is precisely indifferent between participating and instead waiting one more period when his value is low).¹⁰

From (8), it is easy to see that our sufficient condition is more likely to hold in case $\frac{\beta_\tau}{\rho_\tau}$ does not decrease too fast in τ or if the process is not too persistent, i.e. if $\frac{\alpha_H - \alpha_L}{1 - \alpha_L}$ is small. It is always enough that $\frac{\beta_\tau}{\rho_\tau}$ is increasing in τ . It is therefore worth emphasizing that $\frac{\beta_\tau}{\rho_\tau}$ being increasing in τ is a condition that holds for many natural distributions of arrival times. Since the probability of earlier arrival β_τ is increasing in τ , it suffices that ρ_τ is non-increasing. For instance, if T is finite with $\rho_\tau = \frac{1}{T}$ for all τ , then $\frac{\beta_\tau}{\rho_\tau} = \tau - 1$. If arrivals are geometrically distributed with parameter $\lambda \in (0, 1)$, i.e. $\rho_\tau = (1 - \lambda)^{\tau-1} \lambda$ for all τ , then $\frac{\beta_\tau}{\rho_\tau} = \sum_{s=1}^{\tau-1} (1 - \lambda)^{s-\tau}$.

If $\frac{\beta_\tau}{\rho_\tau}$ decreases too fast in τ , then profits equal to the the maximum of (4) may not be attainable in an incentive-compatible mechanism. In this case, one must resort to “ironing” to derive the optimal allocation. Roughly speaking, this requires raising the quality after histories of low values for earlier arrivals, and lowering them for later arrivals, as compared to the quality levels specified in (7). We do not study the ironed solution, but expect our key qualitative insights to carry over to settings where ironing is needed.

If $\frac{\beta_\tau}{\rho_\tau}$ is increasing in τ , then the weight the seller attributes to reducing the rent of earlier arrivals increases over time relative to the weight she assigns to the surplus generated in case of arrival at date τ . Since buyer rents are determined by the qualities allocated in case the buyer's value remains low, this implies that these qualities are distorted downward more at later dates relative to the first-best levels. In particular, conditional on the buyer being in the relationship for the same length of time, distortions are greater if the buyer arrives later. Formally, we find the following.

Corollary 1 *Suppose that $\frac{\beta_\tau}{\rho_\tau}$ is increasing in τ . Consider any two dates τ, τ' , with $\tau < \tau'$, and let $k \in \{0, 1, \dots, T - \tau'\}$. Then*

$$q_{\tau, \tau+k}^* \left(\theta_L^{k+1} \right) \geq q_{\tau', \tau'+k}^* \left(\theta_L^{k+1} \right),$$

with a strict inequality in case $q_{\tau, \tau+k}^ \left(\theta_L^{k+1} \right) > 0$. Moreover, if $\theta_{\tau, \tau} = \theta_{\tau', \tau'}$, then $V_{\tau, \tau}^\Omega \left(\theta_{\tau, \tau}; \emptyset \right) \geq V_{\tau', \tau'}^\Omega \left(\theta_{\tau', \tau'}; \emptyset \right)$, with a strict inequality if $q_{\tau, s}^* \left(\theta_L^{s-\tau+1} \right) > 0$ for some period $s \geq \tau$.*

The result indicates that, when $\frac{\beta_\tau}{\rho_\tau}$ is increasing with τ , late arrivals are punished in that they expect lower rents. This discourages delayed participation in the mechanism, allowing the seller to

¹⁰Formally, the high type's participation constraint (which is the constraint that the mechanism is shown to satisfy in the Appendix) is simply

$$V_{\tau, \tau}^\Omega \left(\theta_H; \emptyset \right) \geq \delta \left((1 - \alpha_H) V_{\tau+1, \tau+1}^\Omega \left(\theta_L; \emptyset \right) + \alpha_H V_{\tau+1, \tau+1}^\Omega \left(\theta_H; \emptyset \right) \right),$$

analogous to the participation constraint (2) for the low type.

extract more rents from early arrivals. In particular, the quality choices are designed to discourage a buyer who arrives with a low value from delaying participation until his value becomes high. It therefore allows the seller to give up less rent in case of earlier arrival while inducing immediate participation.

When $T = +\infty$, and when the buyer arrives with positive probability at each date, the ratio $\frac{\beta_\tau}{\rho_\tau}$ necessarily approaches $+\infty$ with τ . As a consequence, we find that contracts become arbitrarily inefficient as the participation date grows large along histories where the buyer realizes only the low value θ_L .

Corollary 2 *Suppose that $T = +\infty$, with $m_{\tau+1}^{\tau+1}(\theta_L) \leq m_\tau^{\tau+1}(\theta_L^2)$ for all τ (a sufficient condition is that $\frac{\beta_\tau}{\rho_\tau}$ is increasing in τ). For any s , there exists $\bar{\tau}$ sufficiently large that $q_{\tau,t}^*(\theta_L^{t-\tau+1}) = 0$ for all $\tau \geq \bar{\tau}$ and all t such that $t - \tau \leq s$. Hence, buyer rents converge to zero with the participation date.*

What emerges then is a fairly robust principle that optimal mechanisms punish a late arrival. When the horizon is infinite, for instance, buyer rents become arbitrarily small with the participation date. Punishing very late arrivals is beneficial for the seller, since it permits a reduction in rents at all earlier dates, back to date 1.

In contrast, it is worth noting that the “vanishing distortions at the bottom” principle described by Battaglini (2005) continues to hold. In particular, we can show the following.

Corollary 3 *Suppose that $T = +\infty$ with $m_{\tau+1}^{\tau+1}(\theta_L) \leq m_\tau^{\tau+1}(\theta_L^2)$ for all τ (a sufficient condition is that $\frac{\beta_\tau}{\rho_\tau}$ is increasing in τ). For any arrival date τ , $q_{\tau,t}^*(\theta_L^{t-\tau+1})$ converges to its efficient value θ_L as the length of the relationship $t - \tau + 1$ becomes large.*

The reason for this result is that quality choices at dates long after contracting have little effect on the buyer’s rents; choosing qualities close to the efficient ones therefore costs the seller little in terms of the surplus that must be left to the buyer. This is easily seen from the inequality (1) in Lemma 1. According to this lemma, the additional rents that must be given to the buyer in case of arrival with a high value depends on the additional probability that the buyer has of a high value in future. Since $\alpha_H - \alpha_L < 1$, the additional probability of a high value vanishes with time, so later quality allocations affect the buyer’s rents less (see Battaglini, 2005, for a more detailed explanation of the vanishing distortions property).

Finally, it is interesting to consider comparative statics on the transition probabilities. Virtual values and hence qualities for sequences of low values are decreasing in α_H . The reason is that, for higher α_H , a high value persists for a longer time, implying that qualities assigned for sequences of low values have a greater effect on the rents that must be left to the buyer in case his value is initially high (again, see the inequalities in (1) of Lemma 1). The parameter α_L , however, plays two roles. First, a higher value of α_L implies a smaller advantage of high values over low ones, i.e. the opposite

effect as for α_H (see the inequalities in (1)). Second, it increases the likelihood that the buyer's value becomes high if his value is low and he delays participation in the mechanism. This in turn increases the option value of delaying participation. The seller's optimal response to the first effect is to increase qualities, while her optimal response to the second is the opposite. Whether qualities increase or decrease with α_L at any date then depends on parameters and the participation date.

3.1 Experience goods

One subtlety which we have so far overlooked is the possibility that the buyer's value evolves differently when consuming the good compared to when he is without it. For "experience goods", for instance, a buyer may learn about suitability through consumption, but otherwise learn only a little. More generally, the level of excitement a buyer has about a good (and the importance of high quality, in particular) might be expected to fluctuate differently depending on whether the buyer is consuming. In terms of our model, this means values switching at different rates before and after participating in the mechanism.

Let α_H^W and α_L^W , with $\alpha_L^W < \alpha_H^W$, denote the probabilities of a high value at date τ given, respectively, high and low values at date $\tau - 1$ when *not* consuming at $\tau - 1$. Maintain the existing notation for the probability of changes conditional on consumption (i.e., let α_H and α_L denote the probabilities of a high value at date t given high and low values when consuming at date $t - 1$). The previous analysis is easily adapted to this setting, yielding the following result.

Proposition 2 *Suppose that $\frac{\beta_\tau}{\rho_\tau}$ is increasing in τ . Suppose that values evolve differently contingent on past consumption, as described above. Let virtual values be given, for all τ , all $s \geq \tau$, by*

$$\begin{aligned} m_\tau^{W,s}(\theta_L^{s-\tau+1}) &= \theta_L - \left(\frac{\beta_\tau}{\rho_\tau} \frac{\alpha_L^W}{1-\mu} + \frac{\mu}{1-\mu} \right) \left(\frac{\alpha_H - \alpha_L}{1-\alpha_L} \right)^{s-\tau} (\theta_H - \theta_L), \text{ and} \\ m_\tau^{W,s}(\theta_{\tau,\tau}^s) &= \theta_{\tau,s} \text{ for all } \theta_{\tau,\tau}^s \neq \theta_L^{s-\tau+1}. \end{aligned}$$

Profit-maximizing qualities $q_{\tau,s}^W$ are given, for each arrival date τ , each date $s \geq \tau$, and each $\theta_{\tau,\tau}^s$, by

$$c'(q_{\tau,s}^W(\theta_{\tau,\tau}^s)) = \max \{ m_\tau^{W,s}(\theta_{\tau,\tau}^s), 0 \}. \quad (9)$$

A natural assumption is that $\alpha_L^W < \alpha_L$, with the implication that quality allocations are less distorted than those given in Proposition 1. In the extreme case where $\alpha_L^W = 0$, optimal allocations do not depend on the arrival date. The buyer earns no additional rents from his private information about arrival – i.e., a buyer who arrives with a low value expects zero rent irrespective of the arrival date. It is readily checked that the optimal qualities then coincide with those for the optimal mechanism with a known arrival date (as in Battaglini, 2005; equivalently, optimal qualities are the same as for date-1 arrival in Proposition 1). Otherwise, for $\alpha_L^W > 0$, our main qualitative predictions continue to hold; in particular, a buyer whose initial value is low expects positive rent, and (provided that $\frac{\beta_\tau}{\rho_\tau}$ is increasing in τ) quality allocations are less efficient the later the arrival date.

4 Conclusions

This paper considered dynamic mechanism design in a setting where buyers arrive over time and where their preferences evolve stochastically. We provided an approach to fully characterize the optimal mechanism. The key finding was that the mechanism punishes a late participant: late participants face tougher terms of trade and therefore purchase lower qualities and receive less rent. Early arrivals fare better, and buyers earn positive expected rents even if their values are equal to the lowest. Although later arrivals receive less efficient allocations for longer, the “principle of vanishing distortions”, by which allocations converge to first-best levels with time in the relationship, can continue to apply (for instance, it does so in Battaglini’s, 2005, Markovian setting).

Our findings can be expected to have relevance for a broad class of agency relationships beyond the setting of the paper. For instance, a government seeking to procure services at the least cost to tax payers may face new suppliers arising over time whose production costs can be expected to change. Firms seeking to fill top management positions may face potential managers who become available or learn of the position only after time, while their suitability for the job continues to change. Our focus on the seller’s problem with vertical preferences over quality (as in Mussa and Rosen (1978)) was thus only for convenience, and because it allowed us to draw comparisons to the existing literature, especially Battaglini (2005) and subsequent work (e.g., Boleslavsky and Said (2012) and Battaglini and Lamba (2015)). Of course, long-term contracts in various settings may be shaped by considerations specific to each and so different applications may call for further study.

References

- [1] Akan, Mustafa, Baris Ata and James Dana (2011), ‘Revenue Management by Sequential Screening,’ *Journal of Economic Theory*, 159, 728-774.
- [2] Armstrong, Mark and Jidong Zhou (2015), ‘Search Deterrence,’ *Review of Economic Studies*, 83, 26-57.
- [3] Baron, David P. and David Besanko (1984), ‘Regulation and information in a continuing relationship,’ *Information Economics and Policy*, 1, 267-302.
- [4] Battaglini, Marco (2005), ‘Long-Term Contracting with Markovian Consumers’, *American Economic Review*, 95, 637-658.
- [5] Battaglini, Marco and Rohit Lamba (2015), ‘Optimal Dynamic Contracting: the First-Order Approach and Beyond,’ mimeo Princeton and Pennsylvania State University.

- [6] Bergemann, Dirk and Maher Said (2011), ‘Dynamic Auctions,’ in *Wiley Encyclopedia of Operations Research and Management Science*, Wiley, New York, 1511-1522.
- [7] Bergemann, Dirk and Philipp Strack (2015), ‘Dynamic revenue maximization: A continuous-time approach,’ in *Journal of Economic Theory*, 819-853.
- [8] Besanko, David (1985), ‘Multi-period contracts between principal and agent with adverse selection,’ *Economics Letters*, 33-37.
- [9] Blackwell, David (1965), ‘Discounted Dynamic Programming,’ *Annals of Mathematical Statistics*, 36, 2226-2235.
- [10] Board, Simon (2008), ‘Durable-Goods Monopoly with Varying Demand,’ *Review of Economic Studies*, 75, 391-413.
- [11] Board, Simon and Andrzej Skrzypacz (2016), ‘Revenue Management with Forward-Looking Buyers,’ *Journal of Political Economy*, 124, 1046-1087.
- [12] Boleslavsky, Raphael and Maher Said (2013), ‘Progressive Screening: Long-Term Contracting with a Privately Known Stochastic Process,’ *Review of Economic Studies*, 80, 1-34.
- [13] Conlisk, John, Eitan Gerstner and Joel Sobel (1984), ‘Cyclic Pricing by a Durable Goods Monopolist,’ *Quarterly Journal of Economics*, 99, 489-505.
- [14] Courty, Pascal and Li Hao (2000), ‘Sequential Screening,’ *Review of Economic Studies*, 67, 697-717.
- [15] Deb, Rahul (2011, 2014), ‘Intertemporal Price Discrimination with Stochastic Values,’ mimeo University of Toronto.
- [16] Deb, Rahul and Mallesh Pai (2013), ‘Ironing in Dynamic Revenue Management: Posted Prices and Biased Auctions,’ *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*.
- [17] Deb, Rahul and Maher Said (2015), ‘Dynamic screening with limited commitment,’ *Journal of Economic Theory*, 159, 891-928.
- [18] Dilme, Francesco and Fei Li (2016), ‘Revenue Management without Commitment: Dynamic Pricing and Periodic Fire Sales,’ mimeo University of Bonn and University of North Carolina, Chapel Hill.
- [19] Ely, Jeff, Daniel Garrett and Toomas Hinnosaar (forthcoming), ‘Overbooking,’ *Journal of the European Economic Association*.

- [20] Eso, Peter and Balazs Szentes (2007), ‘Optimal Information Disclosure in Auctions and the Handicap Auction,’ *Review of Economic Studies*, 74, 705-731.
- [21] Garrett, Daniel (2011), ‘Durable Goods Sales with Dynamic Arrivals and Changing Values,’ working paper version, mimeo Northwestern University.
- [22] Garrett, Daniel (2016), ‘Intertemporal price discrimination: dynamic arrivals and changing values,’ *American Economic Review*, 106, 3275-3299.
- [23] Garrett, Daniel and Alessandro Pavan (2012), ‘Managerial Turnover in a Changing World,’ *Journal of Political Economy*, 120, 879-925.
- [24] Gershkov, Alex and Benny Moldovanu (2009), ‘Dynamic Revenue Maximization with Heterogeneous Objects: A Mechanism Design Approach,’ *American Economic Journal: Microeconomics*, 1, 168-198.
- [25] Jullien, Bruno (2000), ‘Participation Constrains in Adverse Selection Models,’ *Journal of Economic Theory*, 93, 1-47.
- [26] Mierendorff, Konrad (2016), ‘Optimal dynamic mechanism design with deadlines,’ *Journal of Economic Theory*, 161, 190-222.
- [27] Mussa, Michael and Sherwin Rosen (1978), ‘Monopoly and Product Quality,’ *Journal of Economic Theory*, 18, 301-317.
- [28] Nocke, Volker, Martin Peitz and Frank Rosar (2011), ‘Advance-purchase discounts as a price discrimination device,’ *Journal of Economic Theory*, 146, 141-162.
- [29] Pai, Mallesh M. and Rakesh Vohra (2013), ‘Optimal Dynamic Auctions and Simple Index Rules,’ *Mathematics of Operations Research*, 38, 682-697.
- [30] Pavan, Alessandro, Ilya Segal and Juuso Toikka (2014), ‘Dynamic Mechanism Design: A Myersonian Approach,’ *Econometrica*, 82, 601-653.
- [31] Said, Maher (2012), ‘Auctions with dynamic populations: Efficiency and revenue maximization,’ *Journal of Economic Theory*, 147, 2419-2438.
- [32] Skrzypacz, Andrzej and Juuso Toikka (2015), ‘Mechanisms for Repeated Trade,’ *American Economic Journal: Microeconomics* 7(4), 252-293

A Appendix: Proofs of results for two-value case (main text)

Proof of Lemma 1. Consider a buyer who at date t has reported $\hat{\theta}_{\tau,\tau}^{t-1}$ from date τ up to date $t-1$. That the buyer must be willing to report truthfully a date- t value θ_H implies

$$\begin{aligned} V_{\tau,t}^{\Omega}(\theta_H; \hat{\theta}_{\tau,\tau}^{t-1}) - V_{\tau,t}^{\Omega}(\theta_L; \hat{\theta}_{\tau,\tau}^{t-1}) &\geq (\theta_H - \theta_L) q_{\tau,t}(\hat{\theta}_{\tau,\tau}^{t-1}, \theta_L) \\ &\quad + \delta(\alpha_H - \alpha_L) \left(V_{\tau,t+1}^{\Omega}(\theta_H; \hat{\theta}_{\tau,\tau}^{t-1}, \theta_L) - V_{\tau,t+1}^{\Omega}(\theta_L; \hat{\theta}_{\tau,\tau}^{t-1}, \theta_L) \right). \end{aligned} \quad (\text{A.1})$$

Suppose that

$$\begin{aligned} &V_{\tau,t}^{\Omega}(\theta_H; \hat{\theta}_{\tau,\tau}^{t-1}) - V_{\tau,t}^{\Omega}(\theta_L; \hat{\theta}_{\tau,\tau}^{t-1}) \\ &\geq (\theta_H - \theta_L) \sum_{s=t}^{t'} \delta^{s-t} (\alpha_H - \alpha_L)^{s-t} q_{\tau,s}(\hat{\theta}_{\tau,\tau}^{t-1}, \theta_L^{s-t+1}) \\ &\quad + \delta^{t'-t+1} (\alpha_H - \alpha_L)^{t'-t+1} \left(V_{\tau,t'+1}^{\Omega}(\theta_H; \hat{\theta}_{\tau,\tau}^{t-1}, \theta_L^{t'-t+1}) - V_{\tau,t'+1}^{\Omega}(\theta_L; \hat{\theta}_{\tau,\tau}^{t-1}, \theta_L^{t'-t+1}) \right) \end{aligned}$$

holds for some $t' > t$. Using (A.1) to substitute for the final term then yields

$$\begin{aligned} &V_{\tau,t}^{\Omega}(\theta_H; \hat{\theta}_{\tau,\tau}^{t-1}) - V_{\tau,t}^{\Omega}(\theta_L; \hat{\theta}_{\tau,\tau}^{t-1}) \\ &\geq (\theta_H - \theta_L) \sum_{s=t}^{t''} \delta^{s-t} (\alpha_H - \alpha_L)^{s-t} q_{\tau,s}(\hat{\theta}_{\tau,\tau}^{t-1}, \theta_L^{s-t+1}) \\ &\quad + \delta^{t''-t+1} (\alpha_H - \alpha_L)^{t''-t+1} \left(V_{\tau,t''+1}^{\Omega}(\theta_H; \hat{\theta}_{\tau,\tau}^{t-1}, \theta_L^{t''-t+1}) - V_{\tau,t''+1}^{\Omega}(\theta_L; \hat{\theta}_{\tau,\tau}^{t-1}, \theta_L^{t''-t+1}) \right) \end{aligned}$$

for $t'' = t' + 1$. The result then follows by induction and (for the case of $T = +\infty$) the observation that, in an incentive-compatible mechanism, $V_{\tau,s+1}^{\Omega}(\theta_H; \hat{\theta}_{\tau,\tau}^{t-1}, \theta_L^{s-t+1}) - V_{\tau,s+1}^{\Omega}(\theta_L; \hat{\theta}_{\tau,\tau}^{t-1}, \theta_L^{s-t+1})$ must be uniformly bounded for all $s \geq t$. ■

Proof of Lemma 2. For an incentive-compatible mechanism Ω and any date τ , the buyer's expected payoff must satisfy

$$\begin{aligned} V_{\tau,\tau}^{\Omega}(\theta_L; \emptyset) &\geq \delta \left((1 - \alpha_L) V_{\tau+1,\tau+1}^{\Omega}(\theta_L; \emptyset) + \alpha_L V_{\tau+1,\tau+1}^{\Omega}(\theta_H; \emptyset) \right) \\ &= \delta \left(V_{\tau+1,\tau+1}^{\Omega}(\theta_L; \emptyset) + \alpha_L \left(V_{\tau+1,\tau+1}^{\Omega}(\theta_H; \emptyset) - V_{\tau+1,\tau+1}^{\Omega}(\theta_L; \emptyset) \right) \right) \\ &\geq \delta \left(V_{\tau+1,\tau+1}^{\Omega}(\theta_L; \emptyset) + \alpha_L (\theta_H - \theta_L) \sum_{s=\tau+1}^T \delta^{s-\tau-1} (\alpha_H - \alpha_L)^{s-\tau-1} q_{\tau+1,s}(\theta_L^{s-\tau}) \right), \end{aligned}$$

where the final equality follows from Lemma 1. The same inequality holds also for $V_{\tau+1,\tau+1}^{\Omega}(\theta_L; \emptyset)$. Hence,

$$\begin{aligned} V_{\tau,\tau}^{\Omega}(\theta_L; \emptyset) &\geq \delta \alpha_L (\theta_H - \theta_L) \sum_{s=\tau+1}^T \delta^{s-\tau-1} (\alpha_H - \alpha_L)^{s-\tau-1} q_{\tau+1,s}(\theta_L^{s-\tau}) \\ &\quad + \delta^2 \alpha_L (\theta_H - \theta_L) \sum_{s=\tau+2}^T \delta^{s-\tau-2} (\alpha_H - \alpha_L)^{s-\tau-2} q_{\tau+2,s}(\theta_L^{s-\tau-1}) \\ &\quad + \delta^2 V_{\tau+2,\tau+2}^{\Omega}(\theta_L; \emptyset). \end{aligned}$$

The result then follows from induction, and the fact that $V_{\tau',\tau'}^\Omega(\theta_L; \emptyset)$ remains bounded uniformly over $\tau' > \tau$. ■

Proof of Lemma 3. First note that the buyer's expected rent is given by

$$\begin{aligned} & \sum_{\tau=1}^T \delta^{\tau-1} \rho_\tau (\mu V_{\tau,\tau}^\Omega(\theta_H; \emptyset) + (1-\mu) V_{\tau,\tau}^\Omega(\theta_L; \emptyset)) \\ = & \sum_{\tau=1}^T \delta^{\tau-1} \rho_\tau V_{\tau,\tau}^\Omega(\theta_L; \emptyset) + \mu \sum_{\tau=1}^T \delta^{\tau-1} \rho_\tau (V_{\tau,\tau}^\Omega(\theta_H; \emptyset) - V_{\tau,\tau}^\Omega(\theta_L; \emptyset)). \end{aligned}$$

The first term reflects the rent that the buyer expects to earn even if his value is low at the arrival date, while the second term reflects the additional rent he expects if his value is instead high. A lower bound for the first term is available from Lemma 2:

$$\sum_{\tau=1}^T \delta^{\tau-1} \rho_\tau V_{\tau,\tau}^\Omega(\theta_L; \emptyset) \geq \alpha_L (\theta_H - \theta_L) \sum_{\tau=2}^T \sum_{s=\tau}^T \delta^{s-1} \beta_\tau (\alpha_H - \alpha_L)^{s-\tau} q_{\tau,s} (\theta_L^{s-\tau+1}).$$

A lower bound for the second term is available from simply substituting the expression in Lemma 1:

$$\mu \sum_{\tau=1}^T \delta^{\tau-1} \rho_\tau (V_{\tau,\tau}^\Omega(\theta_H; \emptyset) - V_{\tau,\tau}^\Omega(\theta_L; \emptyset)) \geq \mu (\theta_H - \theta_L) \sum_{\tau=1}^T \rho_\tau \left(\sum_{s=\tau}^T \delta^{s-1} (\alpha_H - \alpha_L)^{s-\tau} q_{\tau,s} (\theta_L^{s-\tau+1}) \right).$$

Therefore, the rents that a buyer is expected to earn must be at least

$$\sum_{\tau=1}^T \sum_{s=\tau}^T \delta^{s-1} (\alpha_L \beta_\tau + \mu \rho_\tau) (\theta_H - \theta_L) (\alpha_H - \alpha_L)^{s-\tau} q_{\tau,s} (\theta_L^{s-\tau+1}).$$

The expression for profits in the lemma is then simply the expected surplus less the lower bound on buyer expected rents. ■

Proof of Proposition 1. The allocations $\mathbf{q}^* = (q_{\tau,t}^*)_{1 \leq \tau \leq t}$ are chosen to maximize (4). (A unique optimum exists by convexity of the cost function $c(\cdot)$.) It remains to verify the existence of a system of transfers \mathbf{p} which implements \mathbf{q}^* as part of an incentive-compatible mechanism. To this end, we begin by specifying the payoff that the buyer expects from truthful reporting at each date t following any history of reports $\hat{\theta}_{\tau,\tau}^{t-1}$ from date τ . We choose these payoffs so that the inequalities (1) and (3) hold with equality, which in turn implies that the buyer's expected rents are as small as possible in an incentive-compatible and individually-rational mechanism implementing \mathbf{q}^* . This means that expected profits are equal to the expression in (4).

There is still much freedom in how payoffs are spread across time. One possible specification is as follows: At each date τ of first reporting

$$V_{\tau,\tau}^\Omega(\theta_L; \emptyset) = \alpha_L (\theta_H - \theta_L) \sum_{i=1}^{T-\tau} \delta^i \left(\sum_{s=\tau+i}^T \delta^{s-\tau-i} (\alpha_H - \alpha_L)^{s-\tau-i} q_{\tau+i,s}^* (\theta_L^{s-\tau-i+1}) \right).$$

For each $t > \tau$, and each history of reports $\hat{\theta}_{\tau,\tau}^{t-1}$, $V_{\tau,t}^\Omega(\theta_L; \hat{\theta}_{\tau,\tau}^{t-1}) = 0$. For each τ , each $t \geq \tau$, and each $\hat{\theta}_{\tau,\tau}^{t-1}$,

$$V_{\tau,t}^\Omega(\theta_H; \hat{\theta}_{\tau,\tau}^{t-1}) = V_{\tau,t}^\Omega(\theta_L; \hat{\theta}_{\tau,\tau}^{t-1}) + (\theta_H - \theta_L) \sum_{s=t}^T \delta^{s-t} (\alpha_H - \alpha_L)^{s-t} q_{\tau,s}^* \left(\hat{\theta}_{\tau,\tau}^{t-1}, \theta_L^{s-t+1} \right).$$

Next, one can choose prices to ensure that these payoffs are realized if the buyer reports truthfully. This is achieved if, for each τ , each $t \geq \tau$ and each $\hat{\theta}_{\tau,\tau}^{t-1}$, we let

$$p_{\tau,t}^* \left(\hat{\theta}_{\tau,\tau}^{t-1}, \theta_{\tau,t} \right) = \theta_{\tau,t} q_{\tau,t}^* \left(\hat{\theta}_{\tau,\tau}^{t-1}, \theta_{\tau,t} \right) - V_{\tau,t}^\Omega \left(\theta_{\tau,t}; \hat{\theta}_{\tau,\tau}^{t-1} \right) + \delta \mathbb{E} \left[V_{\tau,t+1}^\Omega \left(\tilde{\theta}_{\tau,t+1}; \hat{\theta}_{\tau,\tau}^{t-1}, \theta_{\tau,t} \right) \mid \theta_{\tau,t} \right].$$

Now, we wish to check that the mechanism $\langle \mathbf{q}^*, \mathbf{p}^* \rangle$, with $\mathbf{p}^* = (p_{\tau,t}^*)_{1 \leq \tau \leq t}$ is incentive compatible. Two kinds of incentive constraints must be checked. First, conditional on the buyer having reported to the mechanism, he must be willing to report his values truthfully (checking these constraints requires arguments akin to those in Battaglini, 2005; see footnote 9 for further discussion). Second, he must be willing to participate in the mechanism and report to it immediately on the date of his arrival.

Truthful reporting of values. By the ‘‘one-shot deviation principle’’ of Blackwell (1965), it is enough to check that one-shot deviations from truth-telling are never optimal, for any history of past reports. Because the process is first-order Markov, the payoffs available to the buyer at any date t depend only on his date- t value $\theta_{\tau,t}$, and the past reports $\hat{\theta}_{\tau,\tau}^{t-1}$, and not on any previous values. Verifying that the buyer does not profit from a one-shot deviation when his value is high amounts to verifying (A.1), which holds by construction. Verifying that the buyer does not profit from a one-shot deviation when his value is low amounts to checking

$$\begin{aligned} V_{\tau,t}^\Omega(\theta_L; \hat{\theta}_{\tau,\tau}^{t-1}) &\geq V_{\tau,t}^\Omega(\theta_H; \hat{\theta}_{\tau,\tau}^{t-1}) - (\theta_H - \theta_L) q_{\tau,t}^* \left(\hat{\theta}_{\tau,\tau}^{t-1}, \theta_H \right) \\ &\quad - \delta (\alpha_H - \alpha_L) \left(V_{\tau,t+1}^\Omega \left(\theta_H; \hat{\theta}_{\tau,\tau}^{t-1}, \theta_H \right) - V_{\tau,t+1}^\Omega \left(\theta_L; \hat{\theta}_{\tau,\tau}^{t-1}, \theta_H \right) \right). \end{aligned}$$

That (1) holds with equality at all histories implies that this is equivalent to

$$\begin{aligned} &(\theta_H - \theta_L) \sum_{s=t}^T \delta^{s-t} (\alpha_H - \alpha_L)^{s-t} q_{\tau,s}^* \left(\hat{\theta}_{\tau,\tau}^{t-1}, \theta_L^{s-t+1} \right) \\ &\leq (\theta_H - \theta_L) q_{\tau,t}^* \left(\hat{\theta}_{\tau,\tau}^{t-1}, \theta_H \right) + \delta (\alpha_H - \alpha_L) \left(V_{\tau,t+1}^\Omega \left(\theta_H; \hat{\theta}_{\tau,\tau}^{t-1}, \theta_H \right) - V_{\tau,t+1}^\Omega \left(\theta_L; \hat{\theta}_{\tau,\tau}^{t-1}, \theta_H \right) \right) \\ &= (\theta_H - \theta_L) \sum_{s=t}^T \delta^{s-t} (\alpha_H - \alpha_L)^{s-t} q_{\tau,s}^* \left(\hat{\theta}_{\tau,\tau}^{t-1}, \theta_H, \theta_L^{s-t} \right). \end{aligned}$$

This is satisfied because, for all τ , t and s , with $\tau \leq t \leq s$, and all $\hat{\theta}_{\tau,\tau}^{t-1}$, $q_{\tau,s}^* \left(\hat{\theta}_{\tau,\tau}^{t-1}, \theta_H, \theta_L^{s-t} \right) \geq q_{\tau,s}^* \left(\hat{\theta}_{\tau,\tau}^{t-1}, \theta_L^{s-t+1} \right)$.

Timely participation. The above implies that, if the buyer participates at date τ , he then reports his values truthfully from then on, and therefore expects to earn the payoffs specified above. Since transition probabilities do not depend on the buyer's arrival date to the market, the buyer's problem of whether to participate is identical irrespective of his true arrival date. Specifying that the buyer participates at every opportunity, we can then check that the buyer does not gain from one-shot deviations, i.e. from delaying participation. This follows when the buyer's value is low by (2), which is satisfied by construction.

For a high value, we need to check

$$V_{\tau,\tau}^{\Omega}(\theta_H; \emptyset) \geq \delta \left((1 - \alpha_H) V_{\tau+1,\tau+1}^{\Omega}(\theta_L; \emptyset) + \alpha_H V_{\tau+1,\tau+1}^{\Omega}(\theta_H; \emptyset) \right).$$

This is equivalent to

$$(\theta_H - \theta_L) q_{\tau,\tau}^*(\theta_L) + (\theta_H - \theta_L) \sum_{s=\tau+1}^T \delta^{s-\tau} (\alpha_H - \alpha_L)^{s-\tau} (q_{\tau,s}^*(\theta_L^{s-\tau+1}) - q_{\tau+1,s}^*(\theta_L^{s-\tau})) \geq 0. \quad (\text{A.2})$$

It is readily checked that, for all $\tau \leq T - 1$, $m_{\tau+1}^{\tau+1}(\theta_L) \leq m_{\tau}^{\tau+1}(\theta_L^2)$ implies $m_{\tau+1}^s(\theta_L^{s-\tau}) \leq m_{\tau}^s(\theta_L^{s-\tau+1})$ for all $s \geq \tau + 1$. Therefore, $q_{\tau+1,s}^*(\theta_L^{s-\tau+1}) \leq q_{\tau,s}^*(\theta_L^{s-\tau+1})$ for all $s \geq \tau + 1$, so that (A.2) is indeed satisfied. ■

Proof of Corollary 1. The first part follows directly from the qualities specified in Proposition 1. The second part follows using these optimal qualities and the fact that the buyer payoffs at the participation/arrival date $V_{\tau,\tau}(\theta_{\tau,\tau}; \emptyset)$ satisfy the inequalities (1) and (3) with equality. ■

Proof of Corollary 2. The first part follows directly from the qualities specified in Proposition 1. The implication for buyer rents follows using the optimal qualities and the fact that the buyer payoffs at the participation/arrival date $V_{\tau,\tau}(\theta_{\tau,\tau}; \emptyset)$ satisfy the inequalities (1) and (3) with equality. ■

Proof of Corollary 3. This follows immediately from the qualities specified in Proposition 1. ■

Proof of Proposition 2. The lower bound (3) in Lemma 2 becomes

$$V_{\tau,\tau}^{\Omega}(\theta_L; \emptyset) \geq \alpha_L^W (\theta_H - \theta_L) \sum_{i=1}^{T-\tau} \sum_{s=\tau+i}^T \delta^{s-\tau} (\alpha_H - \alpha_L)^{s-\tau-i} q_{\tau+i,s}(\theta_L^{s-\tau-i+1}).$$

Together with the inequalities in Lemma 1 one obtains a lower bound on buyer expected rents. This allows us to derive the upper bound on expected profits

$$\mathbb{E} \left[\sum_{s=\tilde{\tau}}^T \delta^{s-1} \left(m_{\tilde{\tau}}^{W,s}(\tilde{\theta}_{\tilde{\tau},\tilde{\tau}}^s) q_{\tilde{\tau},s}(\tilde{\theta}_{\tilde{\tau},\tilde{\tau}}^s) - c \left(q_{\tilde{\tau},s}(\tilde{\theta}_{\tilde{\tau},\tilde{\tau}}^s) \right) \right) \right],$$

with virtual values $m_\tau^{W,s}$ given in the present result. One then chooses qualities to maximize this expression and then proposes an appropriate implementation as in the proof of Proposition 1. In particular, one should specify payments such that all incentive (and individual rationality) constraints bind, and can then specify that, for each $t > \tau$, and each history of reports $\hat{\theta}_{\tau,\tau}^{t-1}$, $V_{\tau,t}^\Omega(\theta_L; \hat{\theta}_{\tau,\tau}^{t-1}) = 0$.

Verifying that the buyer, having chosen to participate in the mechanism, is willing to report all values truthfully follows the same steps as in the proof of Proposition 1. It then remains to verify the buyer's willingness to participate at his arrival date. By construction, he is indifferent to doing so when his value is low. When his value is high, it is enough to verify that

$$\begin{aligned} & (\theta_H - \theta_L) \sum_{s=\tau}^T \delta^{s-\tau} (\alpha_H - \alpha_L)^{s-\tau} q_{\tau,s}^W (\theta_L^{s-\tau+1}) \\ & \geq (\alpha_H^W - \alpha_L^W) (\theta_H - \theta_L) \sum_{s=\tau+1}^T \delta^{s-\tau} (\alpha_H - \alpha_L)^{s-\tau-1} q_{\tau+1,s}^W (\theta_L^{s-\tau}). \end{aligned}$$

That this is satisfied follows because $\alpha_H^W - \alpha_L^W \leq 1$, and because $\frac{\beta_\tau}{\rho_\tau}$ is increasing in τ . The latter guarantees that $q_{\tau,s}^W (\theta_L^{s-\tau+1}) \geq q_{\tau+1,s+1}^W (\theta_L^{s-\tau+1})$ for all $s \leq T-1$. ■

B Appendix: Continuum of values

This appendix provides results for a continuum of values. The process for arrivals is as described in the model set-up. The key difference will be that, in each period t after arrival, the buyer draws a value from $\Theta = [\underline{\theta}, \bar{\theta}]$. The buyer's initial type $\theta_{\tau, \tau}$ at arrival date τ is drawn from a continuously differentiable c.d.f. F_{In} with density f_{In} and full support on Θ .

For each date $t > \tau$, if the buyer's date $t - 1$ value is $\theta_{\tau, t-1} \in \Theta$, then his date- t value is drawn according to a continuously differentiable c.d.f. $F_{Tr}(\cdot | \theta_{\tau, t-1})$ with density $f_{Tr}(\cdot | \theta_{\tau, t-1})$ and with support on $[\underline{\theta}_{Tr}(\theta_{\tau, t-1}), \bar{\theta}_{Tr}(\theta_{\tau, t-1})] \subset \Theta$. The function $F_{Tr}(\cdot | \cdot)$ is also continuously differentiable in its second argument. Following Garrett and Pavan (2012), we specify that, for any $\theta_{\tau, t-1}, \theta_{\tau, t} \in \Theta$,

$$-f_{Tr}(\theta_{\tau, t} | \theta_{\tau, t-1}) \leq \frac{\partial F_{Tr}(\theta_{\tau, t} | \theta_{\tau, t-1})}{\partial \theta_{\tau, t-1}} \leq 0.$$

The second inequality implies that the conditional distributions $F_{Tr}(\theta_{\tau, t} | \theta_{\tau, t-1})$ are ranked in terms of first-order stochastic dominance, while the first inequality ensures that we can apply a ‘‘dynamic revenue equivalence’’ result developed in Pavan, Segal and Toikka (2014).¹¹

Following Pavan, Segal and Toikka (2014), we introduce the notion of ‘‘impulse responses’’. For any τ, t, s , with $\tau \leq t < s$, and any $\theta_{\tau, t}^s$, let

$$J_t^s(\theta_{\tau, t}^s) \equiv \prod_{l=t+1}^s \left(-\frac{\partial F_{Tr}(\theta_{\tau, l} | \theta_{\tau, l-1}) / \partial \theta_{\tau, l-1}}{f_{Tr}(\theta_{\tau, l} | \theta_{\tau, l-1})} \right).$$

If instead $s = t$, then $J_t^s(\theta_{\tau, t}^s) = J_t^t(\theta_{\tau, t}) = 1$. The value of the impulse response $J_t^s(\theta_{\tau, t}^s)$ can be interpreted as capturing the effect of an infinitesimal variation in $\theta_{\tau, t}$ on $\theta_{\tau, s}$. A simple example is any first-order autoregressive process with persistence parameter γ , in which case the impulse responses are independent of valuations and given by $J_t^s = \gamma^{s-t}$. A property that is helpful for understanding the impulse response function is that, given $\tau \leq t < s$,

$$\frac{\partial}{\partial \theta_{\tau, t}} \mathbb{E} \left[\tilde{\theta}_{\tau, s} | \tilde{\theta}_{\tau, t} = \theta_{\tau, t} \right] = \mathbb{E} \left[J_t^s(\tilde{\theta}_{\tau, t}^s) | \tilde{\theta}_{\tau, t} = \theta_{\tau, t} \right].$$

As for the two-type case, the buyer's expected payoff from truthful reporting in a mechanism $\Omega = \langle \mathbf{q}, \mathbf{p} \rangle$ is denoted $V_{\tau, t}^{\Omega}(\theta_{\tau, t}; \hat{\theta}_{\tau, \tau}^{t-1})$. The following analog to Lemma 1 then determines how, if the mechanism is incentive compatible, the buyer's expected payoff must depend on his value at each date.

Lemma B.1 (Pavan, Segal and Toikka, 2014) *Fix an incentive-compatible mechanism Ω , and consider a buyer who first reports at date τ , and then reports a sequence $\hat{\theta}_{\tau, \tau}^{t-1}$ up to date $t - 1$ (or*

¹¹As noted in Garrett and Pavan (2012), the lower bound on $\frac{\partial F_{Tr}(\theta_t | \theta_{t-1})}{\partial \theta_{t-1}}$ is equivalent to the assumption that, for any $\theta_{t-1} \in \Theta$, and any $x \in \mathbb{R}$, $1 - F(\theta_{t-1} + x | \theta_{t-1})$ is nonincreasing in θ_{t-1} .

makes no reports in case $t = \tau$). If the buyer's value at date t is $\theta_{\tau,t}$, then his expected payoff satisfies

$$V_{\tau,t}^{\Omega}(\theta_{\tau,t}; \hat{\theta}_{\tau,\tau}^{t-1}) = V_{\tau,t}^{\Omega}(\underline{\theta}; \hat{\theta}_{\tau,\tau}^{t-1}) + \int_{\underline{\theta}}^{\theta_{\tau,t}} \mathbb{E} \left[\sum_{s=t}^T \delta^{s-t} J_t^s(\tilde{\theta}_{\tau,t}^s) q_{\tau,s}(\hat{\theta}_{\tau,\tau}^{t-1}, \tilde{\theta}_{\tau,t}^s) \mid \tilde{\theta}_{\tau,t} = r \right] dr.$$

Proof of Lemma B.1. This follows immediately from Theorem 1 of Pavan, Segal and Toikka (2014). ■

Our goal is to use this result to derive a lower bound on the buyer's expected payoff when arriving at any date τ , as in Lemma 2 for the two-type case. To do this we consider the incentive of a buyer who arrives at date τ with the lowest possible initial value $\underline{\theta}$ to misreport his arrival date by delaying participation until the following period. For this deviation to be unprofitable requires

$$V_{\tau,\tau}^{\Omega}(\underline{\theta}; \emptyset) \geq \delta \mathbb{E} \left[V_{\tau+1,\tau+1}^{\Omega}(\tilde{\theta}_{\tau,\tau+1}; \emptyset) \mid \tilde{\theta}_{\tau,\tau} = \underline{\theta} \right]. \quad (\text{B.1})$$

Iterating this requirement yields the following result.

Lemma B.2 Fix an incentive-compatible mechanism $\Omega = \langle \mathbf{q}, \mathbf{p} \rangle$. The buyer's expected payoff for any arrival date τ must satisfy

$$V_{\tau,\tau}^{\Omega}(\underline{\theta}; \emptyset) \geq \sum_{i=1}^{T-\tau} \mathbb{E} \left[\frac{1 - F_{Tr}(\tilde{\theta}_{\tau+i,\tau+i}|\underline{\theta})}{f_{In}(\tilde{\theta}_{\tau+i,\tau+i})} \sum_{t=\tau+i}^T \delta^{t-\tau} J_{\tau+i}^t(\tilde{\theta}_{\tau+i,\tau+i}^t) q_{\tau+i,t}(\tilde{\theta}_{\tau+i,\tau+i}^t) \right]. \quad (\text{B.2})$$

Proof of Lemma B.2. By (B.1), for any τ ,

$$\begin{aligned} V_{\tau,\tau}^{\Omega}(\underline{\theta}; \emptyset) &\geq \delta V_{\tau+1,\tau+1}^{\Omega}(\underline{\theta}; \emptyset) \\ &+ \mathbb{E} \left[\int_{\underline{\theta}}^{\tilde{\theta}_{\tau,\tau+1}} \mathbb{E} \left[\sum_{t=\tau+1}^T \delta^{t-\tau} J_{\tau+1}^t(\tilde{\theta}_{\tau+1,\tau+1}^t) q_{\tau+1,t}(\tilde{\theta}_{\tau+1,\tau+1}^t) \mid \tilde{\theta}_{\tau+1,\tau+1} = r \right] dr \mid \tilde{\theta}_{\tau,\tau} = \underline{\theta} \right] \\ &= \delta V_{\tau+1,\tau+1}^{\Omega}(\underline{\theta}; \emptyset) \\ &+ \mathbb{E} \left[\frac{1 - F_{Tr}(\tilde{\theta}_{\tau,\tau+1}|\underline{\theta})}{f_{Tr}(\tilde{\theta}_{\tau,\tau+1}|\underline{\theta})} \sum_{t=\tau+1}^T \delta^{t-\tau} J_{\tau+1}^t(\tilde{\theta}_{\tau+1,\tau+1}^t) q_{\tau+1,t}(\tilde{\theta}_{\tau+1,\tau+1}^t) \mid \tilde{\theta}_{\tau,\tau} = \underline{\theta} \right] \\ &= \delta V_{\tau+1,\tau+1}^{\Omega}(\underline{\theta}; \emptyset) \\ &+ \mathbb{E} \left[\frac{1 - F_{Tr}(\tilde{\theta}_{\tau+1,\tau+1}|\underline{\theta})}{f_{In}(\tilde{\theta}_{\tau+1,\tau+1})} \sum_{t=\tau+1}^T \delta^{t-\tau} J_{\tau+1}^t(\tilde{\theta}_{\tau+1,\tau+1}^t) q_{\tau+1,t}(\tilde{\theta}_{\tau+1,\tau+1}^t) \right], \end{aligned}$$

where the first equality follows from integration by parts and the second by a simple rearrangement. Iterating then yields (B.2). ■

Lemmas B.1 and B.2 can be used to provide the following analog of Lemma 3.

Lemma B.3 *Suppose that Ω is an incentive-compatible, individually-rational mechanism implementing an allocation \mathbf{q} . Then expected profits are no greater than*

$$\mathbb{E} \left[\sum_{t=\tilde{\tau}}^T \delta^{s-1} \left(m_{\tilde{\tau}}^t \left(\tilde{\theta}_{\tilde{\tau},\tilde{\tau}}^t \right) q_{\tilde{\tau},t} \left(\tilde{\theta}_{\tilde{\tau},\tilde{\tau}}^t \right) - c \left(q_{\tilde{\tau},t} \left(\tilde{\theta}_{\tilde{\tau},\tilde{\tau}}^t \right) \right) \right) \right], \quad (\text{B.3})$$

where, for all τ , all $t \geq \tau$, and all $\theta_{\tau,\tau}^t$,

$$m_{\tau}^t \left(\theta_{\tau,\tau}^t \right) = \theta_{\tau,t} - J_{\tau}^t \left(\theta_{\tau,\tau}^t \right) \left(\frac{\beta_{\tau}}{\rho_{\tau}} \frac{1 - F_{Tr} \left(\theta_{\tau,\tau} | \underline{\theta} \right)}{f_{In} \left(\theta_{\tau,\tau} \right)} + \frac{1 - F_{In} \left(\theta_{\tau,\tau} \right)}{f_{In} \left(\theta_{\tau,\tau} \right)} \right), \quad (\text{B.4})$$

and where expectations are taken over the arrival time $\tilde{\tau}$, as well as the realized values $\tilde{\theta}_{\tilde{\tau},\tilde{\tau}}^T$.

Proof of Lemma B.3. By Lemmas B.1 and B.2, the buyer's expected rent conditional on arriving at date τ is at least

$$\begin{aligned} & \sum_{i=1}^{T-\tau} \mathbb{E} \left[\frac{1 - F_{Tr} \left(\tilde{\theta}_{\tau+i,\tau+i} | \underline{\theta} \right)}{f_{In} \left(\tilde{\theta}_{\tau+i,\tau+i} \right)} \sum_{t=\tau+i}^T \delta^{t-\tau} J_{\tau+i}^t \left(\tilde{\theta}_{\tau+i,\tau+i}^t \right) q_{\tau+i,t} \left(\tilde{\theta}_{\tau+i,\tau+i}^t \right) \right] \\ & + \mathbb{E} \left[\int_{\underline{\theta}}^{\tilde{\theta}_{\tau,\tau}} \mathbb{E} \left[\sum_{s=\tau}^T \delta^{s-\tau} J_{\tau}^s \left(\tilde{\theta}_{\tau,\tau}^s \right) q_{\tau,s} \left(\tilde{\theta}_{\tau,\tau}^s \right) | \tilde{\theta}_{\tau,\tau} = r \right] dr \right]. \end{aligned}$$

For each τ , integrate the second term by parts and subtract the full expression for buyer expected rents from the expected surplus. Taking expectations over the arrival date τ then yields the result.

■

We choose qualities to maximize the expression in (B.3). Under certain conditions, we can then find an incentive-compatible mechanism which implements these allocations, implying the following result.

Proposition B.1 *Suppose that, (i) for all $\tau \leq T-1$, and all pairs $\theta_{\tau,\tau}^{\tau+1}$, $m_{\tau+1}^{\tau+1} \left(\theta_{\tau,\tau+1} \right) \leq m_{\tau}^{\tau+1} \left(\theta_{\tau,\tau}^{\tau+1} \right)$, and (ii) for all τ and all $t \geq \tau$, each $m_{\tau}^t \left(\cdot \right)$ is non-decreasing. Then profit-maximizing qualities are given, for all τ , all $t \geq \tau$, and all $\theta_{\tau,\tau}^t$, by*

$$c' \left(q_{\tau,t}^* \left(\theta_{\tau,\tau}^t \right) \right) = \max \left\{ m_{\tau}^t \left(\theta_{\tau,\tau}^t \right), 0 \right\}.$$

Proof of Proposition B.1. Since the qualities $q_{\tau,t}^*$ are chosen to maximize (B.3), we need only to provide an incentive-compatible mechanism which implements them. As for the proof of Proposition 1, we begin by specifying the buyer's expected payoffs that the mechanism is to deliver the buyer when he reports his values truthfully. For all τ and all $\theta_{\tau,\tau}$, let

$$\begin{aligned} V_{\tau,\tau}^{\Omega} \left(\theta_{\tau,\tau}; \underline{\theta} \right) &= \sum_{i=1}^{T-\tau} \mathbb{E} \left[\frac{1 - F_{Tr} \left(\tilde{\theta}_{\tau+i,\tau+i} | \underline{\theta} \right)}{f_{In} \left(\tilde{\theta}_{\tau+i,\tau+i} \right)} \sum_{t=\tau+i}^T \delta^{t-\tau} J_{\tau+i}^t \left(\tilde{\theta}_{\tau+i,\tau+i}^t \right) q_{\tau+i,t}^* \left(\tilde{\theta}_{\tau+i,\tau+i}^t \right) \right] \\ &+ \int_{\underline{\theta}}^{\theta_{\tau,\tau}} \mathbb{E} \left[\sum_{s=\tau}^T \delta^{s-\tau} J_{\tau}^s \left(\tilde{\theta}_{\tau,\tau}^s \right) q_{\tau,s}^* \left(\tilde{\theta}_{\tau,\tau}^s \right) | \tilde{\theta}_{\tau,\tau} = r \right] dr. \end{aligned}$$

For all τ , all $t > \tau$, and all $(\hat{\theta}_{\tau,\tau}^{t-1}, \theta_{\tau,t})$, let

$$V_{\tau,t}^{\Omega}(\theta_{\tau,t}; \hat{\theta}_{\tau,\tau}^{t-1}) = \int_{\underline{\theta}}^{\theta_{\tau,t}} \mathbb{E} \left[\sum_{s=t}^T \delta^{s-t} J_t^s(\tilde{\theta}_{\tau,t}^s) q_{\tau,s}^*(\hat{\theta}_{\tau,\tau}^{t-1}, \tilde{\theta}_{\tau,t}^s) \mid \tilde{\theta}_{\tau,t} = r \right] dr.$$

We then specify the transfers $p_{\tau,t}^*$ which deliver these payoffs; i.e., we take

$$p_{\tau,t}^*(\hat{\theta}_{\tau,\tau}^{t-1}, \theta_{\tau,t}) = \theta_{\tau,t} q_{\tau,t}^*(\hat{\theta}_{\tau,\tau}^{t-1}, \theta_{\tau,t}) - V_{\tau,t}^{\Omega}(\theta_{\tau,t}; \hat{\theta}_{\tau,\tau}^{t-1}) + \delta \mathbb{E} \left[V_{\tau,t+1}^{\Omega}(\tilde{\theta}_{\tau,t+1}; \hat{\theta}_{\tau,\tau}^{t-1}, \theta_{\tau,t}) \mid \theta_{\tau,t} \right].$$

We now verify that the proposed mechanism is incentive compatible. First note that, since the allocations $q_{\tau,t}^*(\cdot)$ are non-decreasing, Condition (iv) of Corollary 1 in Pavan, Segal and Toikka (2014) is satisfied, so the buyer must be willing to report his values truthfully conditional on participation. This implies that the buyer's expected payoff when participating at any date τ with a value $\theta_{\tau,\tau}$ is equal to

$$V_{\tau,\tau}^{\Omega}(\theta_{\tau,\tau}; \emptyset) = V_{\tau,\tau}^{\Omega}(\underline{\theta}; \emptyset) + \int_{\underline{\theta}}^{\theta_{\tau,\tau}} \mathbb{E} \left[\sum_{s=\tau}^T \delta^{s-\tau} J_{\tau}^s(\tilde{\theta}_{\tau,\tau}^s) q_{\tau,s}^*(\tilde{\theta}_{\tau,\tau}^s) \mid \tilde{\theta}_{\tau,\tau} = r \right] dr.$$

We use this in the remaining step, which is to check the incentive compatibility of immediate participation at the buyer's arrival date.

By the one-shot deviation principle, it suffices to verify that the buyer is willing to participate at an arbitrary date τ . By delaying participation until the following period, the buyer expects a payoff, given $\theta_{\tau,\tau}$, of

$$\mathbb{E} \left[V_{\tau+1,\tau+1}^{\Omega}(\tilde{\theta}_{\tau,\tau+1}; \emptyset) \mid \tilde{\theta}_{\tau,\tau} = \underline{\theta} \right] + \int_{\underline{\theta}}^{\theta_{\tau,\tau}} \mathbb{E} \left[\sum_{s=\tau+1}^T \delta^{s-\tau} J_{\tau}^s(\tilde{\theta}_{\tau,\tau}^s) q_{\tau+1,s}^*(\tilde{\theta}_{\tau,\tau}^s) \mid \tilde{\theta}_{\tau,\tau} = r \right] dr.$$

By construction,

$$V_{\tau,\tau}^{\Omega}(\underline{\theta}; \emptyset) = \delta \mathbb{E} \left[V_{\tau+1,\tau+1}^{\Omega}(\tilde{\theta}_{\tau,\tau+1}; \emptyset) \mid \tilde{\theta}_{\tau,\tau} = \underline{\theta} \right].$$

Therefore, a one-shot deviation at date τ to delaying participation is unprofitable if

$$\begin{aligned} & \int_{\underline{\theta}}^{\theta_{\tau,\tau}} \mathbb{E} \left[\sum_{s=\tau}^T \delta^{s-\tau} J_{\tau}^s(\tilde{\theta}_{\tau,\tau}^s) q_{\tau,s}^*(\tilde{\theta}_{\tau,\tau}^s) \mid \tilde{\theta}_{\tau,\tau} = r \right] dr \\ & \geq \int_{\underline{\theta}}^{\theta_{\tau,\tau}} \mathbb{E} \left[\sum_{s=\tau+1}^T \delta^{s-\tau} J_{\tau}^s(\tilde{\theta}_{\tau,\tau}^s) q_{\tau+1,s}^*(\tilde{\theta}_{\tau,\tau}^s) \mid \tilde{\theta}_{\tau,\tau} = r \right] dr. \end{aligned}$$

To see that this holds, we reason as follows. First consider our assumption that, for all τ and all $\theta_{\tau,\tau}^{\tau+1}$, $m_{\tau+1}^{\tau+1}(\theta_{\tau,\tau+1}) \leq m_{\tau}^{\tau+1}(\theta_{\tau,\tau}^{\tau+1})$. This implies that, for all τ , all $s > \tau$, and all $\theta_{\tau,\tau}^s$, $m_{\tau+1}^s(\theta_{\tau,\tau+1}^s) \leq m_{\tau}^s(\theta_{\tau,\tau}^s)$, and hence $q_{\tau+1,s}^*(\theta_{\tau,\tau+1}^s) \leq q_{\tau,s}^*(\theta_{\tau,\tau}^s)$. That the inequality holds is then immediate from the assumption that J_{τ}^s is non-negative (equivalently, that the distribution of values after date τ are ordered in the sense of first-order stochastic dominance). ■

Conditions (i) and (ii) of this proposition can be understood as follows. First, Condition (ii) guarantees the existence of a mechanism in which, once the buyer has accepted to participate, he truthfully reports his values at all dates. As discussed by Pavan, Segal and Toikka (2014), this condition can be relaxed, although the weaker conditions are often difficult to check. Condition (i) then plays the role of ensuring that the allocations which maximize (B.3) are implementable by an incentive-compatible mechanism in which the constraints in (B.2) are satisfied with equality.

Condition (i), which guarantees timely participation in the mechanism, is new relative to settings where the agent's arrival date is fixed or known (as in Pavan, Segal and Toikka, 2014, for instance). The mechanism we construct ensures that the inequality (B.2) holds as an *equality*, which means that the buyer is indifferent between participating and waiting to participate in the next period when his value is equal to $\underline{\theta}$. Condition (i) then implies that, under the allocations which maximize (B.3), immediate participation at date τ is preferred by all higher types. More precisely, it implies that the benefit of immediate participation is increasing in the buyer's date- τ value. Intuitively, this is because earlier participation gives the buyer access to higher quality levels, for the same evolution of his values (assuming these values are reported truthfully). Like Condition (ii), Condition (i) is a kind of monotonicity condition — it implies monotonicity of the allocations in the participation date. Like Condition (ii), it is somewhat stronger than required, but it is simple to state and a natural analogue to the condition of Proposition 1 for the two-value case considered above.

At least when the conditions of Proposition B.1 hold, we are able to confirm the findings in Corollary 1. If $\frac{\beta_\tau}{\rho_\tau}$ is increasing in τ , then qualities are distorted further below the first-best level at later dates and the buyer expects less rent conditional on his value at arrival. A generalized “principle of vanishing distortions” also applies, provided that the impulse response functions $J_\tau^t(\theta_{\tau,\tau}^t)$ vanish uniformly over time.

We next provide examples of processes for which Proposition B.1 is satisfied.

Example B.1 Let $\bar{\theta} > 0$, let F_{In} be the uniform distribution on $[0, \bar{\theta}]$, and let ϕ be a positive scalar. For each τ and each $t > \tau$, let $\tilde{\theta}_{\tau,t} = \bar{\theta} (1 - \tilde{\varepsilon}_{t-\tau} e^{-\phi\theta_{\tau,t-1}})$, where $\tilde{\varepsilon}_{t-\tau}$ is a random variable distributed uniformly on the unit interval. For each τ , each $t > \tau$, and each $\theta_{\tau,\tau}^t$,

$$m_\tau^t(\theta_{\tau,\tau}^t) = \theta_{\tau,t} - \phi^{t-\tau} \Pi_{s=\tau}^t (\bar{\theta} - \theta_s) \left(\frac{\beta_\tau}{\rho_\tau} + 1 \right). \quad (\text{B.5})$$

The conditions of Proposition B.1 are satisfied provided that, for all τ ,

$$\phi \bar{\theta} \left(\frac{\beta_\tau}{\rho_\tau} + 1 \right) \leq \frac{\beta_{\tau+1}}{\rho_{\tau+1}} + 1. \quad (\text{B.6})$$

Proof of Example B.1 . For any τ, t and $(\theta_{\tau,t-1}, \varepsilon_{t-\tau})$, let $z(\theta_{\tau,t-1}, \varepsilon_{t-\tau}) = \bar{\theta} (1 - \varepsilon_{t-\tau} e^{-\phi\theta_{\tau,t-1}})$. For any sequence of values $\theta_{\tau,\tau}^t$, we may find for each $s \in \{\tau + 1, \dots, t\}$ the shock $\varepsilon_{s-\tau}$ such that

$\theta_{\tau,s} = \bar{\theta} (1 - \varepsilon_{s-\tau} e^{-\phi\theta_{\tau,s-1}})$. Indeed, these are given by

$$\varepsilon_{s-\tau} = \frac{\bar{\theta} - \theta_{\tau,s}}{\theta e^{-\phi\theta_{\tau,s-1}}}.$$

The chain rule yields that

$$J_{\tau}^t(\theta_{\tau,\tau}^t) = \prod_{s=\tau+1}^t \frac{\partial z(\theta_{\tau,s-1}, \varepsilon_{s-\tau})}{\partial \theta_{\tau,s-1}}$$

Hence,

$$\begin{aligned} J_{\tau}^t(\theta_{\tau,\tau}^t) &= \prod_{s=\tau+1}^t \bar{\theta} \phi \varepsilon_{s-\tau} e^{-\phi\theta_{\tau,s-1}} \\ &= \phi^{t-\tau} \prod_{s=\tau+1}^t (\bar{\theta} - \theta_{\tau,s}). \end{aligned}$$

Note also that, for $\theta \in [0, \bar{\theta}]$, $F_{Tr}(\theta|0) = \frac{\bar{\theta}-\theta}{\bar{\theta}}$. Substituting in (B.4) yields (B.5). It is then easy to see that Condition (ii) of Proposition B.1 is satisfied. That Condition (i) is satisfied follows from (B.6). ■

Example B.1 is notable in that it typically includes a wide class of distributions for the arrival date. If $\phi\bar{\theta} \leq 1$, for instance, then it is enough that $\frac{\beta_{\tau}}{\rho_{\tau}}$ is non-decreasing in τ , which is the condition we emphasized for the two-value case. More generally, Condition (i) of Proposition B.1 is more likely to hold if the process is not too persistent. In the above example, the impulse response function is given, for any dates τ , and $t > \tau$, by $J_{\tau}^t(\theta_{\tau,\tau}^t) = \phi^{t-\tau} \prod_{s=\tau+1}^t (\bar{\theta} - \theta_{\tau,s})$; thus ϕ is a parameter which indexes the persistence of the process and the condition holds more easily whenever ϕ is small.

An important class of examples in the literature, beginning with Besanko (1985), concerns autoregressive processes. Suppose the buyer's value evolves according to an autoregressive process, with $\tilde{\theta}_{\tau,t} = \gamma\theta_{\tau,t-1} + \tilde{\varepsilon}_t$ for some $\gamma \in (0, 1]$ and $\tilde{\varepsilon}_t$ an independently distributed ‘‘shock’’. In this case, Condition (ii) is often straightforward to check, while Condition (i) is more difficult, unless restrictive assumptions are made on the distribution of arrivals. Condition (i) is easier to check when there are two possible arrival dates, however, as in the following example.

Example B.2 *Suppose that $T = 2$. Let F_{In} be the uniform distribution on $\Theta = [\underline{\theta}, \bar{\theta}]$, which determines the distribution of $\tilde{\theta}_{1,1}$ and $\tilde{\theta}_{2,2}$. Let $\gamma \in (0, 1]$ and let G be a continuously differentiable c.d.f. on $[\underline{\theta}(1-\gamma), \bar{\theta}(1-\gamma)]$. Suppose that $\tilde{\theta}_{1,2}$ is distributed according to $\gamma\theta_{1,1} + \tilde{\varepsilon}$, where $\tilde{\varepsilon}$ is distributed according to G . Then $m_1^1(\theta_{1,1}) = \theta_{1,1} - \frac{1-F_{In}(\theta_{1,1})}{f_{In}(\theta_{1,1})}$, $m_1^2(\theta_{1,1}^2) = \theta_{1,2} - \gamma \frac{1-F_{In}(\theta_{1,1})}{f_{In}(\theta_{1,1})}$, and $m_2^2(\theta_{2,2}) = \theta_{2,2} - \frac{1-F_{In}(\theta_{2,2})}{f_{In}(\theta_{2,2})} - \frac{\rho_1}{\rho_2} \frac{1-G(\theta_{2,2}-\gamma\underline{\theta})}{f_{In}(\theta_{2,2})}$. Then both conditions of Proposition B.1 are satisfied.*

Proof of Example B.2. Condition (ii) is simple to check. For Condition (i), note that, for each possible θ_1^2 ,

$$\begin{aligned}
\frac{1 - F_{In}(\theta_{1,2})}{f_{In}(\theta_{1,2})} &= \bar{\theta} - \theta_{1,2} \\
&\geq \bar{\theta} - (\gamma\theta_{1,1} + \bar{\theta}(1 - \gamma)) \\
&= \gamma(\bar{\theta} - \theta_{1,1}) \\
&= \gamma \frac{1 - F_{In}(\theta_{1,1})}{f_{In}(\theta_{1,1})}.
\end{aligned}$$

Therefore, $m_2^2(\theta_{1,2}) \leq \theta_{1,2} - \frac{1 - F_{In}(\theta_{1,2})}{f_{In}(\theta_{1,2})} \leq \theta_{1,2} - \gamma \frac{1 - F_{In}(\theta_{1,1})}{f_{In}(\theta_{1,1})} = m_1^2(\theta_{1,1})$, as required. ■

While verifying the conditions in Proposition B.1 can be difficult, certain qualitative properties of the optimal mechanism are quite robust. Indeed, under fairly general conditions (requiring neither of the Conditions (i) or (ii) of Proposition B.1), we can establish a partial analogue of Corollary 2. This describes how the agent fares for all sufficiently late participation dates.

Proposition B.2 *Suppose that $T = +\infty$, with $\rho_\tau > 0$ for all τ . Then the following are true of an optimal mechanism:*

- (i) $V_{\tau,\tau}(\underline{\theta}; \emptyset)$ converges to zero with τ .
- (ii) If, in addition, $F_{Tr}(\cdot|\underline{\theta})$ has full support on Θ , then the buyer's expected rents conditional on participation at date τ , $\mathbb{E}\left[V_{\tau,\tau}(\tilde{\theta}_{\tau,\tau}; \emptyset)\right]$, converge to zero with τ .

Proof of Proposition B.2. We begin with Part (i). Suppose for a contradiction that there exists $\varepsilon > 0$ such that, for all $\bar{\tau}$, there is some $\tau > \bar{\tau}$ with $V_{\tau,\tau}(\underline{\theta}; \emptyset) > \varepsilon$. We consider excluding the buyer after date τ , and then reducing the buyer's rents in case of arrival at date τ or before. We argue that this is possible in such a way that the reduction in (ex-ante) expected buyer rents exceeds the (ex-ante) loss in surplus.

Let $\bar{S} = \max_q\{\bar{\theta}q - c(q)\}$ be the upper bound on the surplus that is generated in each period. The total (discounted life-time) surplus generated by a buyer who participates at some date s is no greater than $\bar{S}\bar{L} = \frac{\bar{S}}{1-\delta}$ in date- s dollars. The contribution to ex-ante expected discounted surplus from arrival after date τ is therefore no greater than

$$\begin{aligned}
\delta^{\tau-1} \sum_{s=\tau+1}^{\infty} \rho_s \delta^{s-\tau} \bar{S}\bar{L} &\leq \delta^\tau \bar{S}\bar{L} \sum_{s=\tau+1}^{\infty} \rho_s \\
&= \frac{\bar{S}\delta^\tau}{1-\delta} \sum_{s=\tau+1}^{\infty} \rho_s.
\end{aligned}$$

Let $R_\tau = V_{\tau,\tau}(\underline{\theta}; \emptyset)$ denote the rent expected conditional on arrival at date τ with value $\underline{\theta}$. Consider excluding participation in the mechanism after date τ and charging an additional participation fee equal to $\delta^{\tau-t}R_\tau$ in case of arrival at each date $t \leq \tau$. The adjusted mechanism remains incentive

compatible and induces immediate participation whenever the buyer arrives at date τ or earlier. The reduction in the ex-ante expected rent left to the buyer is at least $\delta^{\tau-1} R_\tau \sum_{s=1}^{\tau} \rho_s$. The increase in profits is therefore at least

$$\begin{aligned} & \delta^{\tau-1} R_\tau \sum_{s=1}^{\tau} \rho_s - \frac{\bar{S} \delta^\tau}{1-\delta} \sum_{s=\tau+1}^{\infty} \rho_s \\ &= \delta^{\tau-1} \left(R_\tau \sum_{s=1}^{\tau} \rho_s - \frac{\bar{S} \delta}{1-\delta} \sum_{s=\tau+1}^{\infty} \rho_s \right). \end{aligned} \quad (\text{B.7})$$

By assumption, we can pick τ arbitrarily large and such that $R_\tau = V_{\tau,\tau}(\underline{\theta}; \emptyset) > \varepsilon$. Hence, the expression (B.7) can be assured strictly positive for τ chosen sufficiently large. That is, profits are higher under the new mechanism.

Now consider Part (ii). If this result does not hold, then there exists $\varepsilon > 0$ such that we can find a sequence $(\tau_k)_{k=1}^{\infty}$ with the property that $\mathbb{E} \left[V_{\tau_k, \tau_k} \left(\tilde{\theta}_{\tau_k, \tau_k}; \emptyset \right) \right] \geq \varepsilon$ for all $k = 1, 2, \dots$. First note that $V_{\tau_k, \tau_k}(\theta_{\tau_k, \tau_k}; \emptyset)$ is uniformly bounded over k and $\theta_{\tau_k, \tau_k} \in \Theta$. Otherwise, by Lemma B.1 and the assumption that $q \leq \bar{q}$, we must have $V_{\tau_k, \tau_k}(\underline{\theta}; \emptyset)$ is not uniformly bounded, contradicting Part (i) of the proposition. Because $\mathbb{E} \left[V_{\tau_k, \tau_k} \left(\tilde{\theta}_{\tau_k, \tau_k}; \emptyset \right) \right] \geq \varepsilon$ for all k , and because $V_{\tau_k, \tau_k}(\cdot; \emptyset)$ is non-decreasing by Lemma B.1, we can hence find $\nu, \kappa > 0$ such that $V_{\tau_k, \tau_k}(\theta_{\tau_k, \tau_k}; \emptyset) > \nu$ for all $\theta_{\tau_k, \tau_k} > \bar{\theta} - \kappa$. The assumption that $F_{T_\tau}(\cdot | \underline{\theta})$ has full support on Θ then implies the existence of χ such that $\mathbb{E} \left[V_{\tau_k, \tau_k} \left(\tilde{\theta}_{\tau_k-1, \tau_k}; \emptyset \right) | \tilde{\theta}_{\tau_k-1, \tau_k-1} = \underline{\theta} \right] > \chi$ for all k . Since $V_{\tau_k-1, \tau_k-1}(\underline{\theta}; \emptyset) \geq \delta \mathbb{E} \left[V_{\tau_k, \tau_k} \left(\tilde{\theta}_{\tau_k-1, \tau_k}; \emptyset \right) | \tilde{\theta}_{\tau_k-1, \tau_k-1} = \underline{\theta} \right]$ by the incentive constraint (B.1), we have established that $V_{\tau_k-1, \tau_k-1}(\underline{\theta}; \emptyset)$ remains bounded above $\delta \chi > 0$, again contradicting Part (i) of the proposition. ■

For Part (i) of the proposition, the intuition is the familiar one: reducing rents at later dates allows the seller to reduce rents also at all earlier dates, so the seller does well to pick $V_{\tau,\tau}(\underline{\theta}; \emptyset)$ close to zero for large τ . However, this does not necessarily imply that $V_{\tau,\tau}(\theta_{\tau,\tau}; \emptyset)$ should vanish with τ for all $\theta_{\tau,\tau}$. In particular, one can find processes with $\bar{\theta}(\underline{\theta}) < \bar{\theta}$ such that the buyer continues to expect positive rent upon arrival with a value larger than $\bar{\theta}(\underline{\theta})$ under an optimal mechanism. Intuitively, the reason is that permitting the buyer a large rent for high initial values need not create a valuable option for the buyer when he arrives in the previous period. For instance, such high values might only be obtained at the buyer's arrival date. The full-support assumption in Part (ii) of the proposition guarantees that this does not happen.

Proposition B.2 also has implications for optimal qualities, which can be understood by examining Lemma B.1. In particular, note that

$$V_{\tau,\tau}^\Omega(\theta_{\tau,\tau}; \emptyset) = V_{\tau,\tau}^\Omega(\underline{\theta}; \emptyset) + \int_{\underline{\theta}}^{\theta_{\tau,\tau}} \mathbb{E} \left[\sum_{s=\tau}^{\infty} \delta^{s-\tau} J_\tau^s \left(\tilde{\theta}_{\tau,\tau}^s \right) q_{\tau,s}^* \left(\tilde{\theta}_{\tau,\tau}^s \right) | \tilde{\theta}_{\tau,\tau} = r \right] dr$$

under an optimal mechanism with allocation rule $(q_{\tau,t}^*)_{1 \leq \tau \leq t}$. Hence, if $F_{T_\tau}(\cdot | \underline{\theta})$ has full support on Θ , the observation that $\mathbb{E} \left[V_{\tau,\tau} \left(\tilde{\theta}_{\tau,\tau}; \emptyset \right) \right]$ vanishes with τ (Part (ii) of Proposition B.2) implies

that the allocations $q_{\tau,s}^*(\theta_{\tau,\tau}^s)$ cannot be too large at histories $\theta_{\tau,\tau}^s$ such that the impulse responses $J_{\tau}^s(\theta_{\tau,\tau}^s)$ are large, except perhaps with small probability.¹²

¹²Conversely, if the buyer's values are not very persistent, then qualities may not be very distorted after the early periods of the relationship. In the extreme case, where the buyer's values are independently distributed across time, we have $J_{\tau}^s(\theta_{\tau,\tau}^s) = 0$ whenever $s > \tau$. The optimal allocation then coincides with the efficient allocation at every date after the arrival date τ (i.e., $q_{\tau,s}^*(\theta_{\tau,\tau}^s) = \theta_{\tau,s}$ for all $s > \tau$ and all $\theta_{\tau,\tau}^s$).