Dynamic Mechanisms with Martingale Utilities

Santiago R. Balseiro[1], Vahab S. Mirrokni[2], and Renato Paes Leme[2]

[1]Graduate School of Business, Columbia University
srb2155@columbia.edu

[2]Google Research
mirrokni@google.com, renatoppl@google.com

Previous version: February 2016, August 2016 and February 2017
This version: November 2017

Abstract

We study the dynamic mechanism design problem of a seller who repeatedly sells independent items to a buyer with private values. In this setting, the seller could potentially extract the entire buyer surplus by running efficient auctions and charging an upfront participation fee at the beginning of the horizon. In some markets, such as internet advertising, participation fees are not practical since buyers expect to inspect items before purchasing them. This motivates us to study the design of dynamic mechanisms under successively more stringent requirements that capture the implicit business constraints of these markets. We first consider a *periodic individual rationality constraint*, which limits the mechanism to charge at most the buyer's value in each period. While this prevents large upfront participation fees, the seller can still design mechanisms that spread a participation fee across multiple initial auctions. These mechanisms have the unappealing feature that they provide close-to-zero buyer utility in earlier auctions in exchange for higher utility in later auctions. To address this problem, we introduce a *martingale utility constraint*, which imposes the requirement that from the perspective of the buyer, the next item's expected utility is equal to the present one's. Our main result is providing a dynamic auction satisfying martingale utility and periodic individual rationality whose loss in profit with respect to first-best (full extraction of buyer surplus) is optimal up to polylogarithmic factors. The proposed mechanism is a dynamic two-tier auction with a hard floor and a soft floor that allocates the item whenever the buyer's bid is above the hard floor and charges the minimum of the bid and the soft floor.

Keywords: dynamic mechanism design, martingales, approximations, dynamic auctions, internet advertising, revenue management.

1 Introduction

In this paper, we study the problem of designing mechanisms for markets in which independent items are sold to buyers via repeated auctions. This problem is particularly relevant for display advertising markets, where advertisers (the buyers) increasingly acquire impressions (the items) from online platforms called ad exchanges (Muthukrishnan, 2009). In these platforms, impressions arrive one-by-one in an online fashion, and advertisers submit bids in real-time after inspecting the context associated with each impression. As of today, advertisers spend over 10 billion dollars per year in real-time bidding markets (eMarketer, 2017).

While a seller could repeatedly run the optimal single-item auction for each item, a repeated setting allows for dynamic mechanisms that link incentives across auctions and improve on the optimal single-item auction. The intuition is simple: in a single-item setting with private values, the seller needs to pay an information rent for inducing a buyer to reveal her private information, thus limiting the seller's ability to extract surplus (Myerson, 1981). In a repeated setting with independent values, a buyer's information at any given period is still private, however, from today's perspective there is no information asymmetry for future time periods. The seller can thus design mechanisms that charge in advance for the right to get allocated efficiently in the future. By linking auctions together, the seller can bundle items and extract a larger portion of the buyer surplus.[1]

Such mechanisms allow the seller to extract almost the entire buyer surplus, although with – what we argue are – impractical rules that violate the implicit business requirements of real-time bidding markets. Our main contribution is to identify constraints inspired by business requirements in real-time markets, design mechanisms satisfying such constraints, and analyze their performance relative to first-best (full extraction of buyer surplus). While we focus mostly on internet advertising markets, our results can be applied to other settings such as supply chain contracting and repeated principal-agent relationships.

The most basic constraint is that mechanisms should charge payments at the end of each period and should provide incentives for buyers to participate. This *dynamic individual rationality* (DIR) constraint prevents the seller from extracting the expected surplus for the first item, leaving the seller with no alternative other than running the optimal single-item auction for the first item. At the end

[1] This is similar to the practice of *reservation contracts* in internet advertising. In reservation contracts, the seller agrees to deliver a fixed number of impressions satisfying a coarse targeting criteria over a particular time horizon. In this setup, the buyer cannot inspect individual impressions and bid in real-time based on specific viewer information as in ad exchanges. A major business goal in internet advertising is to capture the revenue advantages of reservation contracts in real-time markets.

of the first period, however, both parties are informationally symmetric, and the seller can charge an additional upfront payment for all subsequent items (Besanko, 1985).

The mechanism just described, however, goes against the essence of ad exchanges, where buyers are granted the possibility to bid in real-time based on specific viewer information. When bids are submitted it is often expected (or required) for the ad exchange to honor the semantic of 'bid': the mechanism should respect the buyer's desire not to pay more than a given amount for a given impression. This leads to the second constraint we impose: buyers should derive a positive utility per auction. This *periodic individual rationality* (PIR) constraint limits mechanisms to charge each buyer at most her value for an item, and prevents the seller from charging upfront fees.[2] However, the seller is still able to push this constraint to its limits by spreading an "upfront payment" across multiple initial auctions. For example, by promising buyers that items will be efficiently allocated in the future, the seller can incentivize buyers to initially pay their values. As observed by Biais et al. (2007) and Krishna et al. (2013), this constraint allows for mechanisms in which buyers' utilities are backloaded, that is, buyers are forced to forfeit their utility in earlier auctions, only to be rewarded in later auctions. This is undesirable as buyers typically prefer that the utility accrued from the mechanism is smooth over time.

A third requirement is that the flow utility (expected utility per auction) of a buyer is "stationary" throughout time. If we impose, however, that the flow utility of the buyer is the same for all items, then the best possible mechanism involves offering the optimal single-item auction for each item. To capture the desire for stationary flow utilities while still allowing for interesting dynamics, we introduce the concept of *martingale utilities* (MU), which imposes the requirement that from the perspective of a buyer, the next item's expected utility is equal to the present one's. This constraint is motivated by the common practice of *smooth delivery or pacing* of ads in internet advertising markets. Typically, advertisers have a strong preference for receiving items and spending money uniformly throughout time (Bhalgat et al., 2012). Allocating impressions and spending money smoothly throughout time is a proxy for a more fundamental measure we seek to be stationary over time: the utility that the buyer accrues from the mechanism. Figure 1 provides a pictorial representation of a buyer's expected flow utility under optimal mechanisms satisfying these successively more stringent requirements.

[2]The periodic individual rationality constraint can be motivated by the presence of consumer withdrawal rights (Krähmer and Strausz, 2015). For example, the European Union recently introduced legislation on "distance sales contracts" governing internet and mail order contracts. Section 37 of Directive 2011/83/EU of the European Parliament and of the Council grants a consumer the right "to test and inspect the goods he has bought to the extent necessary" and the right of withdrawal after inspecting the goods. In the case of withdrawal, the seller is required to reimburse the buyer.

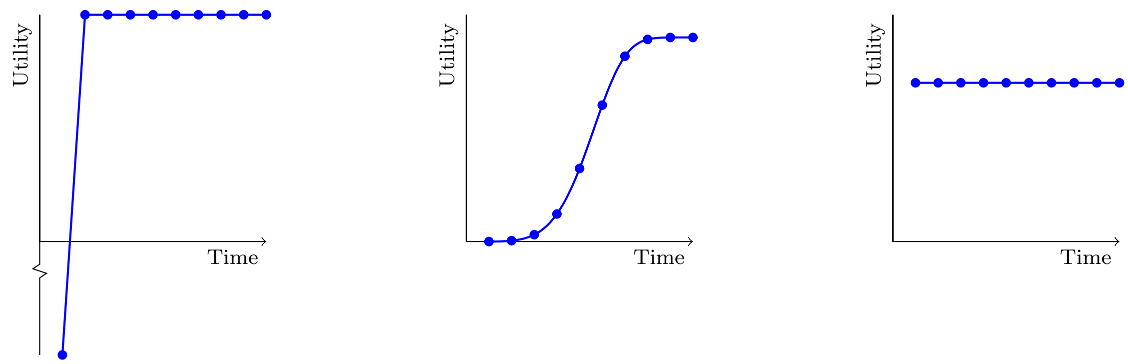

(a) Dynamic Individual Rationality (DIR)

(b) Periodic Individual Rationality (PIR)

(c) Martingale utilities (MU)

Figure 1: Comparison of a buyer's expected flow utility (the utility per item) under optimal mechanisms satisfying successively more stringent requirements.

The main question we investigate in this paper is whether it is possible to achieve close to full surplus extraction in a dynamic setting with limited payments and martingale utilities. We formulate this question within the framework of dynamic mechanism design and provide simple dynamic mechanisms satisfying these requirements. Throughout this paper, we compare the seller's profit under different mechanisms against first-best and show that the average profit of the proposed mechanisms converge to first-best as agents become more patient. We provide sharp characterizations of convergence rates to first-best as a function of the discount rate and prove that the proposed mechanisms are optimal up to polylogarithmic factors. In doing so, we quantify the loss in profit the seller has to take, relative to first-best, to impose these requirements. While the seller has to take a loss to impose periodic individual rationality, an important conclusion from our analysis is that martingale utilities can be imposed at no additional loss (up to polylogarithmic factors).

There are barriers that need to be overcome to apply those mechanisms in practice. The main one is that mechanisms are designed for single-buyer settings. While many advertising sub-markets are thin and can be considered mostly single-buyer, it is important to generalize our ideas to multiple-buyer setting. We believe the mechanisms described in this paper can be extended to accommodate multiple buyers. A second barrier is the assumption of a common prior shared by the seller and buyer. In settings without priors, the seller should take into account the implications on incentives of attempting to "learn" the distribution of buyer's values. We refer the reader to Kanoria and Nazerzadeh (2017) for promising steps in this direction. Nevertheless, we hope that structural insights provided here can be useful to guide the design of mechanisms for real-time markets.

1.1 Main Contributions

We consider the problem of a seller auctioning items that arrive sequentially over an infinite time horizon. The seller has an opportunity cost for each item he sells. To simplify the exposition and highlight the dynamic incentives at play, we assume that there is one buyer with private values that are independent across items. Both parties are risk neutral and discount future payoffs with a common rate.

We employ the framework of promised utility to formulate the dynamic mechanism design problem. This framework, originally introduced by Green (1987), Spear and Srivastava (1987) and Thomas and Worrall (1990), hinges upon the remarkable observation that, when values are independently distributed, an optimal contract can be designed by considering the expected discounted utility of the buyer as the state variable. This forward-looking variable summarizes the stream of all future utilities and allows to formulate the problem recursively via dynamic programming. While this framework can easily accommodate various constraints, the resulting dynamic program is challenging to solve because the action space is complex (the space of all single-item mechanisms). In particular, the classical Myersonian approach cannot be applied to characterize an optimal mechanism because the seller's objective incorporates continuation values, which are typically nonlinear functions. Instead of characterizing optimal dynamic mechanisms, we provide simple dynamic approximation mechanisms[3] that are incentive compatible and whose expected discounted profits converge to first-best as the discount rate converges to one. We are particularly interested in the regime where the discount rate is close to one since advertisers typically participate in a large number of auctions per day. We prove these results by studying the dynamics induced by these stateful mechanisms using martingale concentration arguments.

Our first contribution is the introduction of the *Deferred Utility Mechanism*, a simple incentive compatible approximation mechanism satisfying PIR. The proposed mechanism has three phases: an initial *savings phase* in which the buyer pays her full surplus to the seller in each period, and a subsequent *income phase* in which the buyer is allowed to purchase the item from the seller at cost. The mechanism transitions to the income phase when the present value of all buyer payments exceeds a predefined threshold. The buyer's utilities under this mechanism are backloaded, that is, she experiences zero flow utility in earlier auctions and considerable flow utility in later auctions (see Figure 1b). If the present value of all buyer payments falls below a predefined threshold, the mechanism enters a

[3] An approximation mechanism is a sub-optimal mechanism with theoretical guarantees on performance relative to some benchmark.

temporary *no-trade phase* in which items are not allocated. This phase acts as a threat to guarantee incentive compatibility.

We show that the average discounted profit of the Deferred Utility Mechanism approaches first-best at a rate of $(1-\beta)^{1/2}$, where $\beta \in (0,1)$ denotes the discount rate (we ignore polylogarithmic factors throughout this paper). Note that in the case of dynamic individual rationality inefficiencies only arise in the first period, and the convergence rate to first-best under DIR is $(1-\beta)$. As expected, the convergence rate of the Deferred Utility Mechanism to first-best is slower since the mechanism is constrained to be periodic individually rational and cannot charge large upfront payments. Additionally, we show, using a perfect information relaxation, that the convergence rate of the Deferred Utility Mechanism is tight, in the sense that no other mechanism subject to periodic individual rationality can achieve a better convergence rate. This quantifies the loss in profit the seller has to take to satisfy PIR.

Our second contribution is the introduction of the *Martingale Utility Mechanism*, a simple incentive compatible approximation mechanism satisfying both PIR and MU. The proposed mechanism is a dynamic two-tier auction with a hard floor and a soft floor. The item is allocated whenever the bid is higher than the hard floor. If the bid is above the soft floor, then the auction works like a second-price auction and the buyer pays the soft floor. If the bid lies between the hard floor and the soft floor, then the auction works like a first-price auction and the buyer pays her bid. While this auction is not truthful in a one-shot setting, the mechanism attains truthful reporting via dynamic incentives: the mechanism gives incentives to the buyer to report her value by dynamically adjusting the floors of future auctions based on her present bid. Interestingly, similar two-tier auctions have been used in internet advertising markets, albeit with static floors, which do not guarantee truthful reporting.[4] We show that the average discounted profit of this mechanism approaches first-best at a rate of $(1-\beta)^{1/2}$. Convergence to first-best follows because, as the discount rate increases, the soft floor can be chosen higher and in the limit the seller can run a sequence of truthful "almost" first-price auctions.

1.2 Further Applications

Our analysis can be extended to supply chain contracting with private price information and repeated principal-agent models with hidden cost. The first application involves a manufacturer (the principal) that repeatedly sells a perishable good to a retailer (the agent) facing a newsvendor problem. The retail price in each period is privately observed by the retailer before making the ordering decision. Because the good is perishable, inventory is not carried over and, in the event of a stock out, unmet demand

[4] See, for example, `http://www.mopub.com/2013/04/12/introducing-soft-price-floors/`.

is lost. The manufacturer offers the retailer a dynamic contract that maximizes his expected total discounted profits. This model is a natural extension of Lariviere and Porteus (2001) to a repeated setting with private retail prices. In the second application, a principal contracts with an agent to repeatedly produce output on her behalf. The agent's marginal production cost in each period is privately observed by the agent before making her production decision. This model captures many real-word examples such as retail franchising, labor contracts, and procurement contracts (Laffont and Martimort, 2001).

In both settings the Deferred Utility Mechanism has three phases as before. In the savings phase the principal allocates efficiently and captures the entire surplus from efficient trade. Truthful reporting during the savings phase is achieved exclusively via future promises. In the income phase the principal also allocates efficiently and the agent captures the entire surplus. The mechanism transitions to the income phase when the present value of all transfers exceeds a predefined threshold. In the Martingale Utility Mechanism the principal allocates efficiently and transfers are determined using a dynamic threshold. The principal captures the entire surplus from efficient trade when the agent type is inferior to this threshold, and the agent is provided the additional surplus generated when her type is superior to this threshold. Truthful reporting is achieved via dynamic incentives: if the agent reports a type superior (inferior) to the current threshold, the threshold is adjusted to provide larger (smaller) surplus in the next period. Appendix D gives an in depth exposition of these applications.

1.3 Related Work

In this section we discuss the connection of our work to several streams of literature. First, our paper naturally relates to problems of dynamic mechanism design. Pavan et al. (2014) and Kakade et al. (2013) provide necessary and sufficient conditions for optimal mechanisms in large classes of environments in which agents' information changes over time. This approach involves relaxing most incentive compatibility constraints, solving the relaxed problem, and then showing that the candidate mechanism satisfies all incentive compatibility constraints. To the best of our knowledge, this approach cannot be extended to accommodate constraints such as the ones we consider in this paper. Balseiro et al. (2016) study a dynamic mechanism design when buyers face a cumulative budget constraint and the seller has limited commitment power, that is, he cannot commit to uphold the rules of future auctions. Without budget constraints the optimal mechanism with limited commitment reduces to repeatedly implementing Myerson's optimal auction, which we show to be far from optimal in our setting in which the seller can commit to a dynamic mechanism. Papadimitriou et al. (2014) study the computational complexity of designing optimal dynamic mechanisms and show a strong separation between adaptive, non-adaptive and randomized dynamic mechanisms.

A related stream of work including Vulcano et al. (2002), Gallien (2006), Board and Skrzypacz (2015) and Gershkov and Moldovanu (2014) studies dynamic pricing and revenue management problems using dynamic mechanism design. In these models the designer sells multiple non-perishable items over a finite horizon, introducing option value of waiting for better future opportunities for the seller. This option value is not present in our setting because items are perishable. Chen and Farias (2015) study the design of posted pricing policies in a revenue management setting with customers who arrive dynamically throughout the horizon and can strategically time their purchases. They consider robust pricing policies with a sub-martingale constraint on prices, which incentivizes buyers to purchase items on arrival instead of waiting. In this paper, we impose a martingale restriction on flow utilities, instead of prices, in order to capture the desire for stationary mechanisms. Akan et al. (2015) employ a mechanism design approach to characterize a firm's optimal screening strategy when consumers learn their valuations for future consumption over time. Deb and Said (2015) study a model in which consumers arrive over two time periods and the firm cannot commit in advance to the contractual terms it offers in the second period. Our model differs in that goods sold are independent and perishable, instead of durable; and buyers are not unit-demand. Krähmer and Strausz (2015) argue that sellers no longer benefit from sequential screening with ex-post participation constraints. Moreover, the optimal selling contract with ex-post participation constraints is static and coincides with the optimal posted price contract in the static screening model. In our model, the seller can still improve upon the static allocation with ex-post participation constraints because there are multiple independent items for sale allowing incentives to be linked across auctions. Finally, the result that all inefficiency arises in the first period under risk-neutrality, independent private information and dynamic individual rationally was originally established by Besanko (1985).

The papers closer to ours are Biais et al. (2007), Krishna et al. (2013) and Belloni et al. (2015), where they characterize the optimal dynamic mechanism using the promised utility framework when the agent is liquidity constrained or has limited liability, i.e., under periodic individual rationality constraints. Biais et al. (2007) consider an entrepreneur with limited liability that needs to finance an investment project and explore the optimal financial contract when the entrepreneur can divert operating cash flows. Krishna et al. (2013) study a setting in which a principal contracts with an agent to operate a firm over an infinite time horizon when the agent is liquidity constrained and privately observes the sequence of cost realizations. In both settings, agents are initially incentivized with future promises and agents utilities are backloaded. Belloni et al. (2015) provide a sharp characterization of the optimal dynamic mechanism in a procurement setting and show that the optimal mechanism may require randomization (ironing) even when the virtual value is non-decreasing. While these papers characterize the structure of the optimal dynamic mechanism, the principal still needs to compute the value function to implement the optimal mechanism, which might be challenging in some settings. Additionally, it is not possible to fully characterize the mechanism in closed-form. Instead, here we introduce dynamic approximation mechanisms that are easy to implement and asymptotically achieve first-best. In recent and simultaneous work, Mirrokni et al. (2016) and Ashlagi et al. (2016) show how to compute a $(1 + \epsilon)-$approximation of an optimal dynamic mechanism in polynomial time in a finite horizon model with periodic individual rationality constraints. While those papers compare the profit of their mechanisms against the profit of an optimal dynamic mechanism (that can be obtained via dynamic programming), the objective of our paper is to compare dynamic mechanism satisfying successively more stringent requirement against first-best (full surplus extraction). To the best of our knowledge, ours is the first paper to consider the notion of martingale utilities, show that first-best can be asymptotically achieved via simple repeated auctions in the presence of periodic individual rationality and martingale constraints, and provide sharp bounds on the rate of convergence to first-best as a function of the discount rate.

Our work is related to the literature on linking incentives constraints in mechanism design. Jackson and Sonnenschein (2007) show that linking decisions across identical and independent repeated problems allows the principal to improve on the isolated problem. In their model, agents learn all the values for each problem at time zero and participation decisions are made at the beginning. Instead, our model is motivated by online markets in which values are sequentially revealed as items become available for sale, agents are uncertain about their value for future items, and participation decisions are made for every item. When values are known in advance, the seller can improve his revenue by bundling items together: by selling different goods together, the seller is able to reduce the relative variance of buyers' valuations, therefore reducing the information asymmetry and making the revenue-optimal auction more effective. There is a stream of literature going back to Manelli and Vincent (2006) that studies bundling in non-dynamic mechanisms. We refer the reader to Hart and Nisan (2012); Babaioff et al. (2014); Yao (2015) for a modern discussion on the power of bundling in non-dynamic settings.

2 Model and Problem Formulation

We study a discrete-time infinite horizon setting in which items are sequentially available for sale. We index the sequence of items by $t \in \mathbb{N}$. Future payoffs are discounted according to the discount rate $\beta \in (0,1)$. There is one buyer (she) whose value for the t^{th} item, which we denote by v_t, is independently and identically distributed with cumulative distribution function $F(\cdot)$ and density $f(\cdot)$. We denote $\bar{F}(\cdot) = 1 - F(\cdot)$. Values are supported in the bounded set $\mathcal{V} \triangleq [0, \bar{v}]$ and the density is strictly positive on its domain. The distribution of values is common-knowledge. The value of the t^{th} item is privately observed by the buyer when the t^{th} item arrives. The buyer has a quasilinear utility function given by the difference between the discounted sum of the valuations generated by the items won minus the payments over all auctions she participates. The seller (he) incurs a cost of $c \in [0, \bar{v})$ for selling each item, which is assumed to be common knowledge. The objective of the seller is to maximize expected discounted profit as given by the difference of the payments collected over all auctions and the cost of the goods sold.

2.1 Dynamic Mechanisms

The timing of events is as follows. Initially, the seller announces a dynamic mechanism. A dynamic mechanism is a contingency plan, which the seller commits to honor, specifying an allocation and payment function for every possible state of nature. Then, the following steps are sequentially repeated as a new item becomes available for sale. Firstly, the buyer learns her valuation for the current item and submits a report to the seller. Secondly, the seller publicly announces the outcome of the mechanism for that item, and the payment the buyer should make to the seller. By the Revelation Principle we can focus without loss on direct mechanisms in which the buyer reports her value truthfully to the seller.

We refer to the game associated with the sale of an individual item as the stage game and to the corresponding single-item mechanism as the *stage mechanism*. More formally, a stage mechanism is a pair of functions $(q, z) \in \mathcal{M}$ where $q : \mathcal{V} \to [0,1]$ is an *allocation function* and $z : \mathcal{V} \to \mathbb{R}_+$ is a *payment function*. That is, when the buyer reports v for the item, $q(v)$ determines the probability that the item is allocated and $z(v)$ determines the payment to be charged. A *dynamic mechanism* π is a non-anticipative, adaptive policy that determines a stage mechanism $(q_t^\pi, z_t^\pi) \in \mathcal{M}$ for each time period t (i.e., it makes decisions based on the history of past buyer's reports and seller's actions).

Given a direct dynamic mechanism $\pi \in \mathbb{M}$ the discounted profit of the seller is given by

$$\Pi^\pi = \mathbb{E}^\pi \left[\sum_{t=1}^{\infty} \beta^{t-1} \big(z_t^\pi(v_t) - c q_t^\pi(v_t) \big) \right] ,$$

where the expectation is taken with respect to the history induced by π. Similarly, the discounted utility of the buyer is given by

$$U^\pi = \mathbb{E} \left[\sum_{t=1}^{\infty} \beta^{t-1} \big(v_t q_t^\pi(v_t) - z_t^\pi(v_t) \big) \right] .$$

Finally, in order to compare the utility of the buyer and the profit of the seller for different discount rates β it is useful to define the notion of average discounted utility and average discounted profit per item as $\bar{\Pi}^\pi = (1-\beta)\Pi^\pi$ and $\bar{U}^\pi = (1-\beta)U^\pi$, respectively.[5]

Recursive formulation A challenge with the previous formulation is that the stage mechanism offered at time t potentially depends on the history observed up to time t. We employ the framework of *promised utility* to characterize an optimal dynamic mechanism. Because the seller has commitment power and values are independent, Lemma 1 in Thomas and Worrall (1990) implies that it suffices to consider the expected discounted utility-to-go of the buyer $w_t = \mathbb{E}\left[\sum_{s=t}^{\infty} \beta^{s-t} \big(v_s q_s^\pi(v_s) - z_s^\pi(v_s) \big)\right]$ as the state variable to guarantee dynamic incentive compatibility. We can formulate the mechanism design problem recursively by having the seller arrive at time t, before v_t is realized, with a previously made promise w_t. The seller then delivers the promised utility w_t by choosing a stage mechanism and the continuation value w_{t+1} for the next time period.

Because the problem has an infinite horizon with discounting and values are identically and independently distributed, without loss we can restrict attention to time-independent dynamic mechanisms in which the state variable is the *discounted promised utility* $w \in \mathcal{U} \triangleq [0, \bar{u}]$. Let $y^+ = \max(y, 0)$ be the positive part of a number $y \in \mathbb{R}$. Here $\bar{u} = \mathbb{E}_v\left[(v-c)^+\right]/(1-\beta)$ is an upper bound on the promised utility, which corresponds to the expected social welfare under the first-best allocation. We extend the space of stage mechanism to triples $(q, z, u) \in \mathcal{M}'$ where q and z are allocation and payment functions, respectively, as before, and $u : \mathcal{V} \to \mathcal{U}$ is the *promise function*. That is, when the buyer reports a value of v for the item, $u(v)$ determines the discounted utility-to-go in the next time period.

[5] The infinite horizon discounted model can be interpreted as a random time horizon model with geometric length, i.e., after each time period the problem stops with probability $(1-\beta)$. Under this alternative interpretation, the expected number of items sold is $(1-\beta)^{-1}$.

In this recursive formulation a dynamic mechanism is a contingency plan $(q^\pi(v;w), z^\pi(v;w), u^\pi(v;w))$ that determines the stage mechanism offered at each state $w \in \mathcal{U}$, together with an initial state $w^\pi \in \mathcal{U}$. The dynamic mechanism induces a stochastic process $(w_t^\pi)_{t\geq 1}$ which captures the evolution of the promised utility under the mechanism. This process evolves according to $w_{t+1}^\pi = u^\pi(v_t; w_t^\pi)$ with the initial condition $w_1 = w^\pi$. We emphasize that the promised utility is a state introduced by the mechanism that allows linking decisions across time periods with the objective of improving profits and efficiency. Even though buyer's valuations are independent, the buyer's decisions across time are linked by the state of the mechanism.

Feasible mechanisms A feasible mechanism should satisfy the following constraints. To simplify the notation, we drop the superscript π when referring to a general mechanism. First, the allocation should be feasible, that is,

$$0 \leq q(v;w) \leq 1\,, \tag{1}$$

for every state $w \in \mathcal{U}$ and value v. Additionally, the promised utility should be attainable, that is,

$$0 \leq u(v;w) \leq \bar{u}\,, \tag{2}$$

for every state $w \in \mathcal{U}$ and value v. The mechanism should satisfy the following *promise keeping constraint* for each state $w \in \mathcal{U}$

$$w = \mathbb{E}_v\left[vq(v;w) - z(v;w) + \beta u(v;w)\right]\,, \tag{PK}$$

which imposes that the actual utility delivery by the mechanism coincides with the one promised. The mechanism should be *dynamic incentive compatible* for each state $w \in \mathcal{U}$ of the mechanism, that is, when the buyer's value is v she should be better off reporting her value truthfully over misreporting some value $\tilde{v}$

$$vq(v;w) - z(v;w) + \beta u(v;w) \geq vq(\tilde{v};w) - z(\tilde{v};w) + \beta u(\tilde{v};w)\,. \tag{DIC}$$

This implies, by the one-shot deviation principle, that reporting the value truthfully is weakly dominant for the buyer at every point in time (and regardless of the history).

2.2 Dynamic Revenue Equivalence

In a static setting the celebrated Revenue Equivalence Theorem pins down the payment of a direct mechanism given the allocation. The following result extends this result to our dynamic setting and provides a tractable characterization of a feasible dynamic mechanism. This result essentially adapts Lemma 1 in Besanko (1985) to the promised utility framework. All proofs are deferred to the appendix.

Lemma 2.1. *A dynamic mechanism* $(q(v;w), z(v;w), u(v;w))$ *is feasible if and only if it satisfies constraints* (1), (2)*, the allocation* $q(v;w)$ *is non-decreasing in* v*, and for every state* $w \in \mathcal{U}$

$$z(v;w) - \beta u(v;w) = -U(0;w) + vq(v;w) - \int_0^v q(x;w)\,dx\,, \tag{3}$$

where $U(0;w) = w - \int_0^\infty \bar{F}(x)q(x;w)\,dx$ *is the discounted utility-to-go of the lowest type.*

Using the envelope theorem we establish that given a feasible allocation and a feasible promise function, the payment is pinned down by the dynamic incentive compatibility constraint and the promise keeping constraint. This result readily implies a recursive version of the Revenue Equivalence Theorem: any two feasible dynamic mechanisms with the same allocation and promise function achieve the same expected profit for the seller. Additionally, this lemma provides a straightforward method to construct dynamic incentive compatible mechanisms: for any feasible allocation rule and an attainable promised function, equation (3) pins down a payment rule that makes the mechanism dynamic incentive compatible. While in this paper we study the case of a single buyer, the promise utility framework and the previous result can be easily extended to accommodate multiple buyers.

Lemma 2.1 shows that in a dynamic setting the seller can incentivize the agent to report truthfully by making an instantaneous payment or making future promises. For example, the seller can potentially charge a large payment today and in return promise to provide rents in the future (for example, by allocating efficiently). A priori, it is not clear that such intertemporal substitution of income is beneficial for the seller. The key observation is that while in the current period the buyer's information is private, from today's perspective there is no information asymmetry for future time periods. Thus the seller can charge in advance for the right to get allocated efficiently in the future. Because both parties are informationally symmetric for future time periods, the seller does not need to pay information rents for future periods and he can extract a higher surplus.

2.3 Benchmarks

Dynamic mechanisms are compared against two benchmarks. As a first benchmark we consider the *first-best outcome* (FB), which is the discounted profit an omniscient seller – who observes the realization of values – could extract from the buyer. In this case the seller allocates the item whenever the values is above his marginal cost, and charges the buyer her value. Here the seller achieves full surplus extraction, and his discounted profit is given by

$$\Pi^{\mathrm{FB}} = \frac{1}{1-\beta}\mathbb{E}_v\left[(v-c)^+\right] . \tag{4}$$

The corresponding mechanism is static and given by $q^{\mathrm{FB}}(v) = \mathbf{1}\{v \geq c\}$ and $z^{\mathrm{FB}}(v) = v\mathbf{1}\{v \geq c\}$. This mechanism, while not incentive compatible, provides an upper bound to the objective value of any feasible mechanism.

As a second benchmark we consider the *optimal static* incentive compatible mechanism (S), which amounts to repeatedly allocating items according to Myerson's optimal auction. Let $\phi(v) = v - (1 - F(v))/f(v)$ denote the buyer's virtual value, which is assumed to be increasing.[6] In this case a posted-price scheme is optimal. The optimal price $\phi^{-1}(c)$ solves the problem $\max_{r\geq c}(r-c)\mathbb{P}\{v \geq r\}$ and the seller's discounted profit is given by

$$\Pi^{\mathrm{S}} = \frac{1}{1-\beta}\mathbb{E}_v\left[(\phi(v)-c)^+\right] . \tag{5}$$

The corresponding mechanism is $q^{\mathrm{S}}(v) = \mathbf{1}\{\phi(v) \geq c\}$ and $z^{\mathrm{S}}(v) = \phi^{-1}(c)\mathbf{1}\{\phi(v) \geq c\}$. This mechanism is clearly feasible and provides a lower bound to the objective value of an optimal dynamic mechanism.[7]

It is convenient to note that the ratio $\bar{\Pi}^{\mathrm{FB}}/\bar{\Pi}^{\mathrm{S}}$ between these two benchmarks can be arbitrarily large. A folklore example involves the truncated equal revenue distribution, given by the cumulative distribution function $F(v) = 1 - 1/v$ for $v \in [1, \bar{v})$ and $F(\bar{v}) = 1$. For $c = 0$, we have $\bar{\Pi}^{\mathrm{FB}} = \mathbb{E}[v] = 1 + \log \bar{v}$ while $\bar{\Pi}^{\mathrm{S}} = \max_{r\in[1,\bar{v}]} r\mathbb{P}\{v \geq r\} = 1$. Therefore, $\bar{\Pi}^{\mathrm{FB}}/\bar{\Pi}^{\mathrm{S}} = 1 + \log \bar{v} \to \infty$ when $\bar{v} \to \infty$.

[6] This assumption is typical in the mechanism design literature. Examples of distributions satisfying this condition are uniform, exponential, truncated normal, and Weibull, among others. Interestingly, we do not require this assumption to characterize and analyze our approximation mechanisms.

[7] To simplify the exposition we describe the optimal static mechanism in term of histories. The same mechanism in the promised utility formulation can be given by: $q(v; w) = \mathbf{1}\{v \geq r(w)\}$ where $r(w)$ is such $\mathbb{E}_v[(v - r(w))^+] = (1-\beta)w$, $z(v; w) = r(w)\mathbf{1}\{v \geq r(w)\}$ and $u(v; w) = w$. The initial state is $w^{\mathrm{S}} = \mathbb{E}_v[(v - \phi^{-1}(c))^+]/(1-\beta)$. Because the mechanism is static, the initial state w^{S} is absorbing and all other states are irrelevant.

3 Design Principles

We next formally define the desirable properties a practical dynamic mechanism should satisfy in the markets we consider.

Dynamic Individual Rationality While the seller is committed to honor his promises, the buyer is free to walk away at any time. The first desirable property is that the mechanism should provide incentives for the buyer to participate; such a mechanism is said to be self-enforcing from the buyer's stand point. More formally, a mechanism is *dynamic individually rational* for each state $w \in \mathcal{U}$ if the discounted utility-to-go of the buyer is non-negative for all values v

$$vq(v;w) - z(v;w) + \beta u(v;w) \geq 0\,. \tag{DIR}$$

An optimal dynamic mechanism under DIR allocates the item according to Myerson's optimal auction in the first period and allocates the item efficiently onwards. The seller incorporates a large participation fee in the first auction that allows him to extract the buyer's entire surplus from the second auction onwards. Truthful reporting in the first period is guaranteed by offsetting the first instantaneous payment with a promise of large future reward. Here the seller exploits that at the point of contracting the only private information is the buyer's value for the first auction and that there is no information asymmetry for future time periods. Thus, the seller only pays the buyer an information rent for the first time period and achieves close to full surplus extraction. This result is well-known in the mechanism design literature and was originally established by Besanko (1985). Appendix C provides a closed-form characterization of an optimal dynamic mechanism subject to DIR and shows that the average discounted profit of this mechanism converges at a rate $1 - \beta$ to first-best.

Periodic Individual Rationality An optimal dynamic mechanism satisfying DIR can induce the buyer to make a large upfront payment in the first auction in exchange of repeatedly allocating the item efficiently in the future. In some markets, however, large participation fees are not plausible and buyers expect to inspect items before paying. We exclude such mechanisms by imposing the stronger *periodic individually rationality* constraint (PIR), which imposes that the buyer should derive a positive utility from each auction:

$$vq(v;w) - z(v;w) \geq 0\,. \tag{PIR}$$

This constraint limits the instantaneous payment of the buyer to be less than or equal to her value for the current item and restricts the use of large upfront fees.

Optimal mechanisms satisfying PIR attempt to spread the "upfront payment" across multiple auctions. In earlier auctions the seller captures the entire surplus from trade and truthful reporting is achieved by future promises: in return, the item is afterwards allocated efficiently and the buyer pays the seller's marginal cost for each item. Section 4 introduces the Deferred Utility Mechanism, a simple approximation mechanism satisfying PIR that exhibits this behaviour. We show that the average discounted profit of this mechanism converges to first-best at a rate $(1-\beta)^{1/2}$, and that no other mechanism can converge faster to first-best.

Martingale Utilities While PIR limits payments to at most the buyer's value, it allows for non-stationary mechanisms in which the buyer is initially incentivized with promised future payments and in return derives little-to-no flow utility in earlier auctions. This is undesirable as buyers typically prefer that the utility accrued from the mechanism is smooth over time, instead of back-loaded. Thus motivated we study optimal mechanisms under the constraint that the flow utility $\ell(w) = \mathbb{E}_v[vq(v;w) - z(v;w)]$ is a martingale, that is, $\mathbb{E}_v\left[\ell\big(u(v;w)\big)\right] = \ell(w)$ for all states $w \in \mathcal{U}$. This imposes that the expected utility of the buyer for the next auction, conditioning on the current state, is equal to expected utility of the current auction. The previous constraint is equivalent to requiring that the promised utility is a martingale, that is, the mechanism satisfies

$$\mathbb{E}_v[u(v;w)] = w\,, \tag{MU}$$

for all states $w \in \mathcal{U}$ (see Lemma B.2 in the appendix for a proof).[8] This last constraint is more tractable from an analytical perspective and in the rest of paper we work exclusively with it.

The MU constraint implies that promised future payments are now limited as the mechanism needs to consistently provide rents throughout the horizon. Section 5 introduces the Martingale Utility Mechanism, a simple approximation mechanism satisfying PIR and MU that achieves truthful reporting via a combination of instantaneous payments and promised future payments. Additionally, the average discounted profit of this mechanism converges to first-best at a rate $(1-\beta)^{1/2}$.

[8]This equivalence only holds in the time-discounted model. In a finite horizon model with T time periods, the corresponding constraint on the promised utilities is $\mathbb{E}[w_{t+1} \mid w_t] = (1 - 1/(T-t+1))\, w_t$ for all $t = 1, \dots, T$, which implies that the promised utility is a supermartingale.

4 Dynamic Mechanisms with Periodic Individual Rationality

An optimal dynamic mechanism under PIR can be recursively characterized by the Principle of Optimality. Let $\Pi^{\mathrm{PIR}}(w)$ be the optimal expected profit-to-go of the seller when the promised utility is $w \in \mathcal{U}$. We have:

$$\Pi^{\mathrm{PIR}}(w) = \max_{(q,z,u)\in\mathcal{M}} \mathbb{E}_v[z(v) - cq(v) + \beta \Pi^{\mathrm{PIR}}(u(v))] \tag{6}$$
$$\text{s.t } (1), (2), (\mathrm{PK}), (\mathrm{DIC}), (\mathrm{PIR})\,.$$

Given a value function $\Pi^{\mathrm{PIR}}(\cdot)$ solving the previous Bellman Equation, the optimal profit of the seller is given by $\Pi^{\mathrm{PIR}} = \max_{w\in\mathcal{U}} \Pi^{\mathrm{PIR}}(w)$, and the initial state w^{PIR} is a maximizer of the previous problem.[9]

In this setting an optimal dynamic mechanism cannot be characterized in closed-form. Using Lemma 2.1 it is still possible to eliminate payments from the inner optimization problem in (6), and obtain a simpler problem constrained only by (1), (2) and (PIR). In this case, however, the allocation and utility functions are coupled by (PIR) and the objective is not separable. This leads to an infinite dimensional optimization problem over the space of feasible allocations, and because the objective is non-linear it is not possible to optimize the objective in a pointwise fashion, as done in the single-item setting (Myerson, 1981). Belloni et al. (2015) provide sufficient conditions under which the optimal stage mechanism can be characterized given the value function. In order to implement the optimal dynamic mechanism, however, the seller needs to determine the value function, and it is not possible to fully characterize the mechanism in closed-form. Additionally, they show that the optimal mechanism may require randomization (ironing) even when the virtual value is non-decreasing. Thus motivated, we consider a dynamic approximation mechanism that is easy to implement and asymptotically optimal. We first provide the intuition behind the derivation of the approximation mechanism.

Let $S^{\mathrm{PIR}}(w) = w + \Pi^{\mathrm{PIR}}(w)$ be the social welfare generated by an optimal mechanism. After eliminating payments using Lemma 2.1, we obtain the simpler equivalent recursion

$$S^{\mathrm{PIR}}(w) = \max_{(q,u)} \mathbb{E}_v[(v - c)q(v) + \beta S^{\mathrm{PIR}}(u(v))] \tag{7}$$

[9]The optimal value function $\Pi^{\mathrm{DIR}}(\cdot)$ is a fixed-point of the Bellman operator $(T\Pi)(w) = \max_{(q,z,u)\in\mathcal{M}':(1),(2),(\mathrm{PK}),(\mathrm{DIC}),(\mathrm{DIR})} \mathbb{E}_v[z(v) - cq(v) + \beta\Pi(u(v))]$ for all $w \in \mathcal{U}$. A unique optimal solution can be shown to exist because the Bellman operator T is a contraction mapping under the sup-norm when $\beta < 1$.

subject to (1), (2) and

$$\beta u(v) \le w + \int_0^v q(x)\mathrm{d}x - \int_0^{\bar v} \bar F(x) q(x)\mathrm{d}x\,. \tag{PIR'}$$

Note that $S^{\text{PIR}}(w)$ is non-decreasing in w because the constraint set of the inner optimization problem expands as w increases. Suppose the state is not close to $\bar u$. For any fixed allocation q the seller would like to set the promise function as large as possible and constraint (PIR') should be binding. This pins down the promise function. The change in the promised utility from one auction to the next is at most $2\bar v/\beta$, which is small relatively to $\bar u$ when agents are patient. Thus, continuation values are approximately linear in the region of achievable states from the current state. Performing a first-order expansion of the value function around w we obtain $S^{\text{PIR}}(u(v)) \approx S^{\text{PIR}}(w) + \frac{d}{dw} S^{\text{PIR}}(w)(u(v) - w)$, which implies after taking expectations that $\beta\mathbb{E}[S^{\text{PIR}}(u(v))] \approx \beta S^{\text{PIR}}(w) + \frac{d}{dw} S^{\text{PIR}}(w)(1-\beta)w$ because $\beta\mathbb{E}[u(v)] = w$ from (PIR'). Therefore, the allocation has a negligible impact on the second term of (7) and the seller optimizes the objective by maximizing flow profit. Thus, allocating efficiently according to $q(v) = \mathbf{1}\{v \ge c\}$ is nearly optimal. A similar argument shows that allocating efficiently is nearly optimal when $u(v)$ is close to $\bar u$.

This argument readily suggests an approximation mechanism for the seller: allocate efficiently whenever possible, set the promise function so that (PIR') is binding and determine payments using Lemma 2.1. However, we must deal with two issues: (i) when the promised utility is low the mechanism needs to be modified so that the promised utility never drop below zero and (ii) the initial state needs to be determined to maximize expected profits. We address these issues in the next section.

4.1 The Deferred Utility Mechanism

In this section we propose the *Deferred Utility Mechanism* (DUM), a dynamic mechanism satisfying dynamic incentive compatibility and periodic individual rationality that asymptotically achieves first-best as the discount rate converges to one. We study the problem when the discount rate satisfies the following assumption.

Assumption 4.1. *The discount rate satisfies* $\beta > (\bar v - c)/(\mathbb{E}_v\left[(v-c)^+\right] + \bar v - c) \ge 1/2$.

For example, when values are $U[0,1]$ and the seller cost is zero, this assumption requires that $\beta \ge 2/3$, or at least 3 items are sold on average if we interpret discounting as geometric lifetimes.

We denote the approximation mechanism by $(q^{\text{DUM}}, z^{\text{DUM}}, u^{\text{DUM}})$. Let $\underline{w} = \mathbb{E}_v\left[(v-c)^+\right]$ and $\bar w =$

$\mathbb{E}_v\left[(v-c)^+\right]/(1-\beta)-(\bar{v}-c)$. The allocation function is given by

$$q^{\text{DUM}}(v;w)=\begin{cases}0\,, & \text{if } w\in[0,\underline{w})\,,\\ \mathbf{1}\{v\geq c\}\,, & \text{if } w\in[\underline{w},\bar{u}]\,,\end{cases}$$

the payment function is given by

$$z^{\text{DUM}}(v;w)=\begin{cases}0\,, & \text{if } w\in[0,\underline{w})\,,\\ v\mathbf{1}\{v\geq c\}\,, & \text{if } w\in[\underline{w},\bar{w}]\,,\\ c\mathbf{1}\{v\geq c\}\,, & \text{if } w\in(\bar{w},\bar{u}]\,,\end{cases}$$

and the promise function is given by

$$u^{\text{DUM}}(v;w)=\begin{cases}\frac{1}{\beta}w\,, & \text{if } w\in[0,\underline{w})\,,\\ \frac{1}{\beta}\left(w+\Delta^{\text{DUM}}(v)\right)\,, & \text{if } w\in[\underline{w},\bar{w}]\,,\\ \frac{1}{\beta}\left(w+\Delta^{\text{DUM}}(c)\right)\,, & \text{if } w\in(\bar{w},\bar{u}]\,,\end{cases}$$

where we denote by $\Delta^{\text{DUM}}(x)=(x-c)^+-\mathbb{E}_v\left[(v-c)^+\right]$ the mean adjusted surplus. The initial promised utility is chosen so that $w^{\text{DUM}}\in[\underline{w},\bar{w}]$. The next result uses Lemma 2.1 to show that the proposed dynamic mechanism is feasible.

Proposition 4.2. *Suppose Assumption 4.1 holds. The Deferred Utility Mechanism* $(q^{\text{DUM}},z^{\text{DUM}},u^{\text{DUM}})$ *is periodic individually rational and dynamic incentive compatible for all initial states* $w\in[0,\bar{u}]$*.*

The proposed mechanism has three phases depending on the value of the promised utility. The *savings phase* is characterized by promised utilities in the intermediate range $[\underline{w},\bar{w}]$. In this phase the item is allocated according to a first-price auction (FPA) with reserve price c, i.e., the buyer obtains the item when her value is above the cost c and pays her value. In a single-item setting the first-price auction is known to be non-truthful as bidders have an incentive to shade their bids. In our dynamic setting, the mechanism attains truthful reporting via dynamic incentives: the buyer is promised a relatively higher (lower) discounted surplus in future auctions when her report is high (low). Thus the seller provides an incentive for the buyer to pay her value via future promises. The *income phase* is characterized by promised utilities in the highest range $(\bar{w},\bar{u}]$. In this phase the item is allocated efficiently according to a second-price auction (SPA) with reserve price c, i.e., the buyer obtains the item when her value is above the cost c and pays this cost. The *no-trade phase* is characterized by promised utilities in the lowest range $[0,\underline{w})$. In this phase the item is not allocated and no payments are collected, and in return the promised utility is monotonically increasing to prevent the promised utility from being negative. The no-trade phase plays a critical role in guaranteeing DIC. For example, if the buyer decides to consistently report c in the savings phase, the state would eventually reach the no-trade phase and items would not be allocated until the state leaves that phase.[10]

Dynamics The initial promised utility w^{DUM} is chosen to lie in the savings phase. In the savings phase the seller's expected flow profit is $\mathbb{E}_v\left[(v-c)^+\right]$ and thus he extracts the whole buyer surplus in each auction. The mechanism substitutes present surplus with future surplus, i.e., it provides no surplus to the buyer in the current auction in return for positive surplus in the future. In this phase the promised function satisfies $\mathbb{E}_v[u^{\text{DUM}}(v;w)] = w/\beta$, and thus the promised utility drifts upwards towards the income phase. Because the seller derives his profit during the savings phase, the initial promised utility w^{DUM} is chosen to carefully balance two effects. If the initial promised utility is high, the time spent in the savings phase is low as the promised utility would quickly drift towards the unprofitable income phase. Alternatively, if the initial promised utility is low, the time spent in the savings phase is low again as the promised utility could likely fall into the no-trade phase. Figure 2 illustrates the dynamics of the approximation mechanism.

Compared to an optimal mechanism with dynamic individual rationality, the seller can no longer extract the buyer's surplus via large upfront payments. Because payments are limited to be smaller than values, the seller instead spreads the "upfront payment" across the multiple auctions of the savings phase and employs a first-price auction as an instrument to collect this payment.

The proposed mechanism resembles a deferred annuity contract between a principal and an agent. Interpreting the discount rate as the buyer having a geometric lifetime, we can interpret the mechanism as an insurance contract: the buyer invests money into the account during the savings phase until a given level is reached, and then the buyer gets a "lifetime" of payments in the form of being provided the entire surplus from efficient allocation.[11]

[10] While the buyer might be indifferent between reporting a lower value or truthful bidding, the seller strictly prefers the latter. The seller can modify the mechanism so that the buyer is strictly better off reporting truthfully, at the expense of a small loss in profit (see, e.g., Fiat et al. (2013)).

[11] Because the drift in the income phase is non-positive, the promised utility might alternate between the income and savings phase. However, the length of these incursions into the savings phase are small.

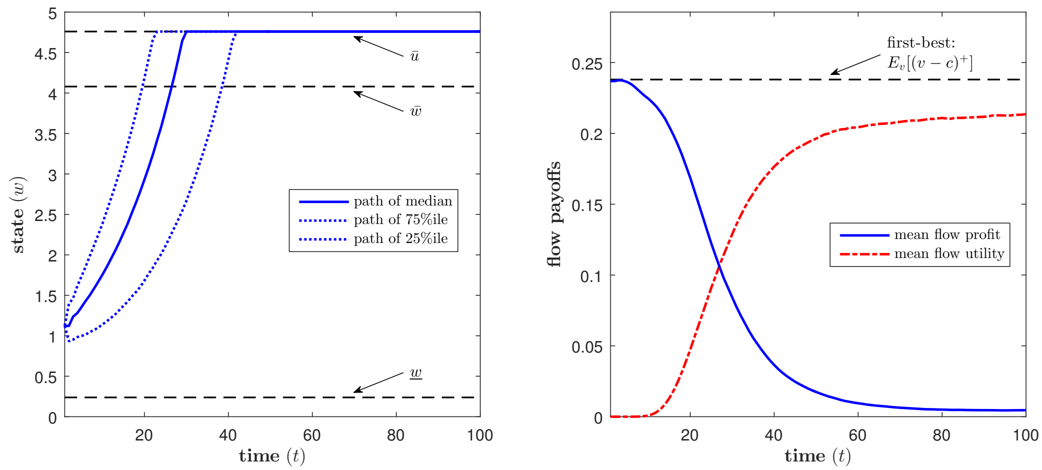

Figure 2: Dynamics of approximation mechanism for periodic individual rationality. The left plot shows the evolution of the state, and the right plot shows the evolution of the mean flow profit of the seller and mean flow utility of the buyer. Values are $U[0,1]$, the seller's cost is $c = 0.3$, and the discount rate is $\beta = 0.95$. Results are obtained using Monte Carlo simulation of 50,000 sample paths.

4.2 Performance Analysis

The next result compares the performance of the approximation mechanism to that of the first-best mechanism.[12]

Theorem 4.3. *Suppose Assumption 4.1 holds. Let* $w^{\text{DUM}} = \underline{w} + \frac{\sqrt{2}\bar{v}}{(1-\beta)^{1/2}} \log^{1/2}\left(\frac{1}{1-\beta}\right)$. *We have*

$$\Pi^{\text{DUM}} \leq \Pi^{\text{PIR}} \leq \Pi^{\text{FB}} \leq \Pi^{\text{DUM}} + \tilde{O}\left(\frac{1}{(1-\beta)^{1/2}}\right).$$

A straightforward corollary of the previous result is that the average profit of the mechanism $(q^{\text{DUM}}, z^{\text{DUM}}, u^{\text{DUM}})$ and the optimal PIR mechanism achieve first-best as the discount rate converges to one:

$$\lim_{\beta \to 1} \bar{\Pi}^{\text{DUM}} = \lim_{\beta \to 1} (1-\beta) \Pi^{\text{DUM}} = \mathbb{E}_v\left[(v-c)^+\right] = \bar{\Pi}^{\text{FB}}.$$

The rate of convergence to first-best is $(1-\beta)^{1/2}$, which is slower than the rate of convergence under an optimal DIR mechanism, since in this case the mechanism is constrained to be periodic individually rational and cannot charge large upfront payments. Also the initial promised utility is higher than the one provided by an optimal DIR mechanism (see Appendix C). We next provide some intuition for

[12] We say $f(\beta)$ is big O of $g(\beta)$ or $f(\beta) = O(g(\beta))$ if and only if there exists $C > 0$ and $\beta_0 \in (0,1)$ such that $|f(\beta)| \leq C|g(\beta)|$ for all $\beta_0 \leq \beta < 1$. We say $f(\beta)$ is soft O of $g(\beta)$ or $f(\beta) = \tilde{O}(g(\beta))$ if and only if $f(\beta) = O\big(g(\beta) \log^k g(\beta)\big)$ for some $k > 0$. We say $f(\beta)$ is big Omega of $g(\beta)$ or $f(\beta) = \Omega(g(\beta))$ if and only if there exists $C > 0$ and $\beta_0 \in (0,1)$ such that $|f(\beta)| \geq C|g(\beta)|$ for all $\beta_0 \leq \beta < 1$.

this result.

Let $S^{\text{DUM}} = \Pi^{\text{DUM}} + w^{\text{DUM}}$ be the social welfare generated by the mechanism. Ignore momentarily the possibility of the state going below $\underline{w}$ and entering the no-trade phase. Because the mechanism allocates efficiently whenever possible and payments are an internal transfer of wealth, we obtain that $S^{\text{DUM}} \approx \mathbb{E}_v\left[(v-c)^+\right]/(1-\beta) = \Pi^{\text{FB}}$. This implies $\Pi^{\text{DUM}} \approx \Pi^{\text{FB}} - w^{\text{DUM}}$ and thus the seller would like to pick the initial promised utility as small as possible to maximize the number of time periods spent in the savings phase. Recall that in order to enforce DIC, the promised utility has to decrease with positive probability, which could eventually lead the mechanism to the undesirable no-trade phase. Hence the previous analysis is incomplete as it does not take into account the possibility of entering the no-trade phase. The loss in profit with respect to the first-best given in Theorem 4.3 arises from the fact that the initial promised utility should be chosen large enough so that the probability that the state falls to the no-trade phase is small. We next briefly discuss convergence to first-best.

As the discount rate increases the promised utility concentrates around its expected path, which has an upward drift during the savings phase (since $\mathbb{E}_v[u^{\text{DUM}}(v;w)] = w/\beta$ because $\mathbb{E}_v[\Delta^{\text{DUM}}(v)] = 0$). Hence, the initial promised utility w^{DUM} can be chosen to be lower relative to Π^{FB} while guaranteeing that the probability that the promised utility hits the no-trade phase is small, which allows the seller to increase the number of time periods spent in the savings phase and, in the limit, achieve first-best.

4.3 The Cost of Periodic Individual Rationality

We next study the cost of imposing PIR on the seller's profit by measuring the loss in profit the seller has to take in order to guarantee the PIR constraint. We denote this loss by $\Pi^{\text{FB}} - \Pi^{\text{PIR}}$. Theorem 4.3 shows that the cost of imposing PIR is bounded from above as $\Pi^{\text{FB}} - \Pi^{\text{PIR}} \le \tilde{O}\left((1-\beta)^{-1/2}\right)$. The next result shows that no other mechanism satisfying PIR can achieve a loss smaller than $\Omega\left((1-\beta)^{-1/2}\right)$. This implies that the convergence rate of the Deferred Utility Mechanism given in Theorem 4.3 is tight, in the sense that no other mechanism satisfying PIR can convergence faster to first-best (up to polylogarithmic factors). We prove the result by considering a model in which buyer's values are discrete.

Theorem 4.4. *Suppose that the seller's cost is $c = 0$ and the buyer's value is a non-degenerate and positive two-point distribution. Then $\Pi^{\text{FB}} - \Pi^{\text{PIR}} \ge \Omega\left((1-\beta)^{-1/2}\right)$.*

The proof proceeds by using structural properties of the optimal value function to simplify the optimal mechanism design problem. As described above, it is possible to eliminate the payment and promise functions from the mechanism design problem. Because there are no distortions at the top, the problem reduces to optimally determining the probability of selling the item when the buyer's value is low. Let q_t be the probability of selling the item when the type is low at time t. When the promised utility is close to $\bar{u}$, the seller cannot profit since promises are delivered by providing all surplus to the buyer. Hence, the seller profits only when the promised utility is not close to $\bar{u}$. In this case, we obtain that $\mathbb{E}[w_{t+1}|w_t] = w_t/\beta$ regardless of q_t because (PIR') is binding. Therefore, the seller would like to pick the initial promised utility as small as possible. When the state is close to zero, however, the seller must reduce the volatility of the promised utility process, at a loss, by setting $q_t \approx 0$ to avoid being absorbed at $w = 0$. Because of stochastic fluctuations in the promised utility process, every mechanism needs to incur a minimal loss of $\Omega\left((1-\beta)^{-1/2}\right)$ to avoid being absorbed by the lower boundary.

Bounding the loss in profit with respect to first-best of an optimal non-anticipative policy is challenging. Instead, we provide an upper bound to this problem using a perfect information relaxation in which the seller knows in advance all realizations of the buyer's value. We then characterize the optimal objective value of the perfect information relaxation in terms of the expected deviations from the origin of a reflected random walk, and we conclude by bounding these deviations using concentration inequalities.

5 Dynamic Mechanisms with Martingale Utilities

As before, an optimal dynamic mechanism can be recursively characterized by the Principle of Optimality. Let $\Pi^{\mathrm{MU}}(w)$ be the optimal expected profit-to-go of the seller when the promised utility is $w \in \mathcal{U}$. We have:

$$\Pi^{\mathrm{MU}}(w) = \max_{(q,z,u)\in\mathcal{M}'} \mathbb{E}_v[z(v) - cq(v) + \beta\Pi^{\mathrm{MU}}(u(v))] \quad (8)$$
$$\text{s.t } (1), (2), (\mathrm{PK}), (\mathrm{DIC}), (\mathrm{PIR}), (\mathrm{MU})\,.$$

Given a value function $\Pi^{\mathrm{MU}}(\cdot)$ solving the previous Bellman Equation, the optimal profit of the seller is given by $\Pi^{\mathrm{MU}} = \max_{w\in\mathcal{U}} \Pi^{\mathrm{MU}}(w)$, and the initial state w^{MU} is a maximizer of the previous problem. Because an optimal dynamic mechanism cannot be characterized in closed-form, we consider a dynamic approximation mechanism that is easy to implement and asymptotically optimal.

We next provide the intuition behind the derivation of the approximation mechanism. Let $S^{\mathrm{MU}}(w) = w + \Pi^{\mathrm{MU}}(w)$ be the social welfare of an optimal mechanism. Using Lemma 2.1 to eliminate payments as before, we obtain the simpler equivalent recursion

$$S^{\mathrm{MU}}(w) = \max_{(q,u)} \mathbb{E}_v[(v-c)q(v) + \beta S^{\mathrm{MU}}(u(v))] \tag{9}$$

subject to (1), (2), (MU) and (PIR'). It is not hard to see that $S^{\mathrm{MU}}(w)$ is nondecreasing and concave in w, where the latter follows because the constraint set of the inner optimization problem is linear and the objective is concave. Fix an allocation q and let $\beta \bar{u}_q(v)$ be the right-hand side of (PIR'), that is, $u(v) \le \bar{u}_q(v)$. The seller determines the promise function by solving a knapsack problem with concave and non-decreasing rewards: the knapsack has a capacity of w given by (MU), "objects" are indexed by v, and $u(v)$ denotes the amount of v put into the knapsack. Object v has a reward $f(v)S^{\mathrm{MU}}(u(v))$, volume $f(v)u(v)$, and upper bounds $u(v) \le \bar{u}_q(v)$. Because the ratio of reward to volume is the same for all objects, the seller would like to simultaneously increase $u(v)$ across all objects as much as possible until the knapsack is full. Since the upper bounds $\bar{u}_q(v)$ are non-decreasing in v, it follows that there exists a threshold s such that $u(v) = \bar{u}_q(v)$ for $v \le s$ and $u(v) = \bar{u}_q(s)$ for $v > s$. That is, (PIR') is binding for types lower than s and the threshold should satisfy $\mathbb{E}_v[\bar{u}_q(\min(v,s))] = w$. This pins down the promise function. Performing a first-order expansion of the value function around w, as before, and using that the flow utility is a martingale, we obtain that the allocation has a negligible impact on the second term of (9) and the seller optimizes flow profit by allocating efficiently according to $q(v) = \mathbf{1}\{v \ge c\}$. This argument readily suggests an approximation mechanism for the seller.

5.1 The Martingale Utility Mechanism

In this section we propose the *Martingale Utility Mechanism* (MUM), a dynamic mechanism satisfying dynamic incentive compatibility, periodic individual rationality, and martingale utilities that asymptotically achieves first-best as the discount rate converges to one. We study the problem when the discount rate satisfies the following assumption.

Assumption 5.1. *The discount rate satisfies* $\beta > 1 - \mathbb{E}_v[(v-c)^+]/(\bar{v}-c)$ *and* $\beta \ge 1/2$.

For example, when values are $U[0,1]$ and the seller cost is zero, this assumption requires that $\beta \ge 1/2$, or at least 2 items are sold on average if we interpret discounting as geometric lifetimes.

We denote the Martingale Utility Mechanism by $(q^{\mathrm{MUM}}, z^{\mathrm{MUM}}, u^{\mathrm{MUM}})$. Let $\underline{w} = \mathbb{E}_v[(v-c)^+]$. The allocation function is given by

$$q^{\mathrm{MUM}}(v;w) = \begin{cases} (w/\underline{u})\,\mathbf{1}\{v \geq c\}\,, & \text{if } w \in [0,\underline{w})\,,\\ \mathbf{1}\{v \geq c\}\,, & \text{if } w \in [\underline{w},\bar{u}]\,,\end{cases}$$

the payment function is given by

$$z^{\mathrm{MUM}}(v;w) = \begin{cases} (w/\underline{u})c\,\mathbf{1}\{v \geq c\}\,, & \text{if } w \in [0,\underline{w})\,,\\ \min(v, s^{\mathrm{MUM}}(w))\mathbf{1}\{v \geq c\}\,, & \text{if } w \in [\underline{w},\bar{u}]\,,\end{cases}$$

where we let *soft floor* $s^{\mathrm{MUM}}(w)$ be a solution of the equation

$$(1-\beta)w = \mathbb{E}_v\left[(v - s^{\mathrm{MUM}}(w))^+\right]\,, \tag{10}$$

and the promise function is given by

$$u^{\mathrm{MUM}}(v;w) = \begin{cases} w\,, & \text{if } w \in [0,\underline{w})\,,\\ \frac{1}{\beta}\Big(w + \big(\min(v, s^{\mathrm{MUM}}(w)) - c\big)^+ - \mathbb{E}_v[(v-c)^+]\Big)\,, & \text{if } w \in [\underline{w},\bar{u}]\,.\end{cases}$$

The initial promised utility will be chosen so that $w^{\mathrm{MUM}} \in [\underline{w},\bar{u}]$. The next result uses Lemma 2.1 to show that the proposed dynamic mechanism is feasible.

Proposition 5.2. *Suppose Assumption 5.1 holds. The Martingale Utility Mechanism* $(q^{\mathrm{MUM}}, z^{\mathrm{MUM}}, u^{\mathrm{MUM}})$ *is dynamic incentive compatible, periodic individually rational and satisfies the martingale utilities constraint for all states* $w \in [0,\bar{u}]$*. Additionally, the soft floor satisfies* $c \leq s^{\mathrm{MUM}}(w) \leq \bar{v}$*.*

The proposed mechanism has an *allocation phase* characterized by promised utilities in the range $[\underline{w},\bar{u}]$. In this phase items are allocated according to a two-tier auction with a hard floor c and soft floor $s^{\mathrm{MUM}}(w) \geq c$. The item is allocated whenever the value is higher than c. If the value is above the soft floor $s^{\mathrm{MUM}}(w)$, then the auction works like a SPA (second-price auction) and the buyer pays the soft floor $s^{\mathrm{MUM}}(w)$. If the value lies between the hard floor c and the soft floor $s^{\mathrm{MUM}}(w)$, then the auction works like a FPA (first-price auction) and the buyer pays her value. The soft floor is dynamic and determined using (10) as a function of the promised utility. In particular, the soft floor $s^{\mathrm{MUM}}(w)$ is decreasing in the promised utility, and equals c when the promised utility is $\bar{u}$ and equals $\bar{v}$ when the promised utility is zero. While this auction is not truthful in a single-item setting, the mechanism attains truthful reporting via dynamic incentives: if the buyer reports a high (low) value, the soft floor in the next auction goes down (up) leading to higher (lower) future surplus. The mechanism also has a *throttled phase* characterized by promised utilities in the lowest range $[0, \underline{w})$. In this phase the item is allocated with small probability whenever the value is above the cost. The probability is set low enough to guarantee that the promised utility remains constant. The throttled phase plays a critical role in guaranteeing DIC. For example, if the buyer decides to consistently report c, the state would eventually reach the throttled phase and the item would start being awarded only with small probability.

Dynamics The initial promised utility w^{MUM} is chosen to lie in the allocation phase. Recall that in this phase, the lower the promised utility, the higher the soft floor. Thus if the initial promise is low, the seller's expected flow profit is higher. However, because the promised utility is a martingale, the lower the initial promised utility, the larger the likelihood that the state could drift downwards into the unprofitable throttled phase. The initial promised utility is chosen to carefully balance these two effects. Figure 3 illustrates the dynamics of the approximation mechanism.

The DIR and PIR mechanism given in Section C and Section 4 have two consecutive phases: first the seller profits either by charging a large upfront fee or by running a sequence of first-price auctions, and then the buyer derives the bulk of her surplus by getting allocated efficiently. Because the mechanism is now constrained to satisfy (PIR) and (MU), the seller can no longer incentivize the buyer exclusively using future promises. Instead the mechanism here described incentivizes the buyer using a combination of instantaneous payments and future promises, and implements a soft floor/hard floor auction in which both parties simultaneously derive surplus. In particular, the seller captures the entire surplus when the value is below the soft floor and the buyer is provided the additional surplus generated when the value is above the soft floor.

5.2 Performance Analysis

The next result compares the performance of the approximation mechanism to that of the first-best mechanism.

Theorem 5.3. *Suppose Assumption 5.1 holds. Let* $w^{\mathrm{MUM}} = \underline{w} + \frac{\bar{v}}{\beta} \frac{1}{(1-\beta)^{1/2}} \log\left(\frac{1}{1-\beta}\right)$. *We have*

$$\Pi^{\mathrm{MUM}} \le \Pi^{\mathrm{MU}} \le \Pi^{\mathrm{FB}} \le \Pi^{\mathrm{MUM}} + \tilde{O}\left(\frac{1}{(1-\beta)^{1/2}}\right).$$

A straightforward corollary of the previous result is that both the proposed mechanism $(q^{\mathrm{MUM}}, z^{\mathrm{MUM}}, u^{\mathrm{MUM}})$

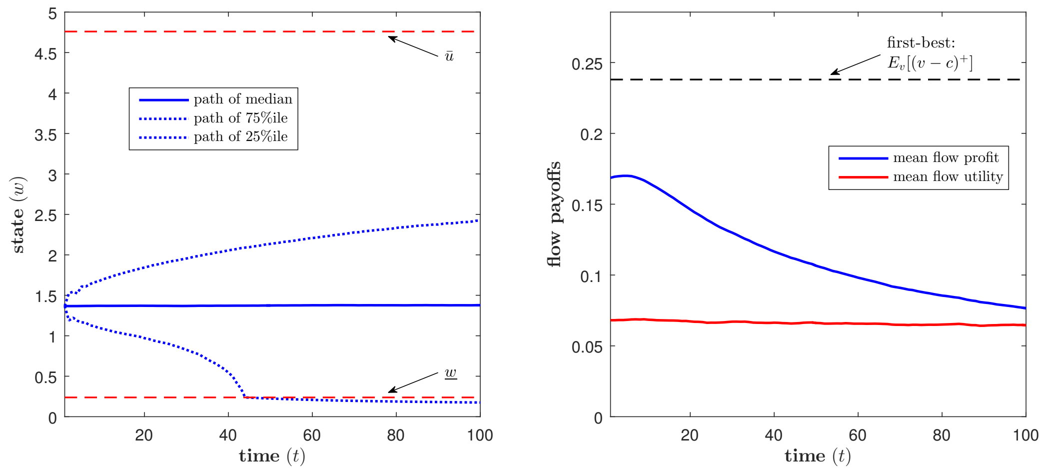

Figure 3: Dynamics of approximation mechanism for martingale utilities. The left plot shows the evolution of the state, and the right plot shows the evolution of the mean flow profit of the seller and mean flow utility of the buyer. Values are $U[0,1]$, the seller's cost is $c = 0.3$, and the discount rate is $\beta = 0.95$. Results are obtained using Monte Carlo simulation of 50,000 sample paths.

and the optimal MU mechanism achieve first-best as the discount rate converges to one:

$$\lim_{\beta \to 1} \bar{\Pi}^{\mathrm{MUM}} = \lim_{\beta \to 1} (1-\beta)\Pi^{\mathrm{MUM}} = \mathbb{E}_v\left[(v-c)^+\right] = \bar{\Pi}^{\mathrm{FB}}.$$

The rate of convergence to first-best is $(1-\beta)^{1/2}$, which is slower than the rate of convergence under an optimal DIR mechanism, yet similar to the rate of convergence of an optimal PIR mechanism. We next provide some intuition for this result.

Ignoring momentarily the possibility of the state going below $\underline{w}$ and entering the throttled phase, we obtain, as in Section 4.2, that $\Pi^{\mathrm{MUM}} \approx \Pi^{\mathrm{FB}} - w^{\mathrm{MUM}}$ because the proposed mechanism allocates efficiently whenever possible. Hence the seller is better off when the promised utility is smaller, which is expected because the soft floor employed in the two-tier auction increases as the promised utility decreases. The initial promised utility is chosen to balance the loss in profit derived from eventually entering the unprofitable throttled phase (which increases as w^{MUM} decreases) and the profit collected by the two-tier auction (which increases as w^{MUM} decreases). We next briefly discuss convergence to first-best.

As the discount rate increases, the promised utility concentrates around the initial promised utility w^{MUM} because the process is a martingale. Hence, the initial promised utility w^{MUM} can be chosen to be relatively closer to $\underline{w}$ while guaranteeing that the loss in profit derived from eventually entering the throttled phase is small, which allows the seller to increase the soft floor employed in the two-tier auction. In the limit the soft floor converges to $\bar{v}$ and the two-tier auction reduces to a first-price auction with reserve price c. Thus the dynamic mechanism asymptotically achieves first-best since the seller is running a sequence of truthful "almost" first-price auctions.

While in Theorem 5.3 the initial state is chosen to maximize seller's profit, the buyer can be provided any utility in the range $[0, \bar{u}]$ by setting the initial state appropriately. In particular, any initial state w satisfying $(1-\beta)w^2 \to \infty$ as $\beta \to 1$ would guarantee that the average social welfare achieved by the Martingale Utility Mechanism converges to first-best as agents become more patient. For example, by setting the initial state to the utility of the buyer under the optimal static mechanism $w^{\mathrm{S}} = \mathbb{E}_v[(v - \phi^{-1}(c))^+]/(1-\beta)$, the seller can guarantee the same level of buyer utility as in the optimal static mechanism and, at the same time, improve efficiency.[13]

In light of Theorem 4.4, the convergence rate of the Martingale Utility Mechanism given in Theorem 5.3 is tight, in the sense that no other mechanism satisfying PIR and MU can convergence faster to first-best (up to polylogarithmic factors). The loss in profit is mainly driven by the PIR constraint: it is possible to design mechanisms satisfying DIR and MU whose the convergence to first-best is of order $(1-\beta)$ (for example, by considering a mechanism that allocates efficiently and charges in each period the expected utility of the next period).

6 Numerical Experiments

In this section we numerically study the dynamic mechanism design problem for different values of the discount rate β, and in each setting compare the proposed approximation mechanism to an optimal dynamic mechanism. In order to compute an optimal dynamic mechanism, we discretize values by setting a uniform grid with 50 points and discretize the state space (the promised utility) by setting a logarithmic grid with 500 points. Dynamic program (6) and (8) are solved via policy iteration with the stopping condition that the n^{th} iterate $\Pi^{(n)}$ satisfies $\|\Pi^{(n)} - \Pi^{(n-1)}\|_\infty \le \epsilon$ with $\epsilon = 0.0001$ (see Bertsekas 2012, Chapter 2 for an overview of policy iteration). Because Π is concave (since the action space is convex and the flow profit is linear), we can write it as the minimum of linear envelopes and thus the inner problems in each value iteration can be solved via linear programming. We estimate the performance of the approximation mechanisms in Section 4 and Section 5 using Monte Carlo simulation of 1,000 sample paths; the resulting mean standard errors are small in all examples. Throughout these experiments we assume that values are $U[0,1]$ and the seller's cost is $c = 0.3$. Results for different distribution of values and seller's costs are similar, and not reported.

[13]We thank an anonymous referee for suggesting this analysis.

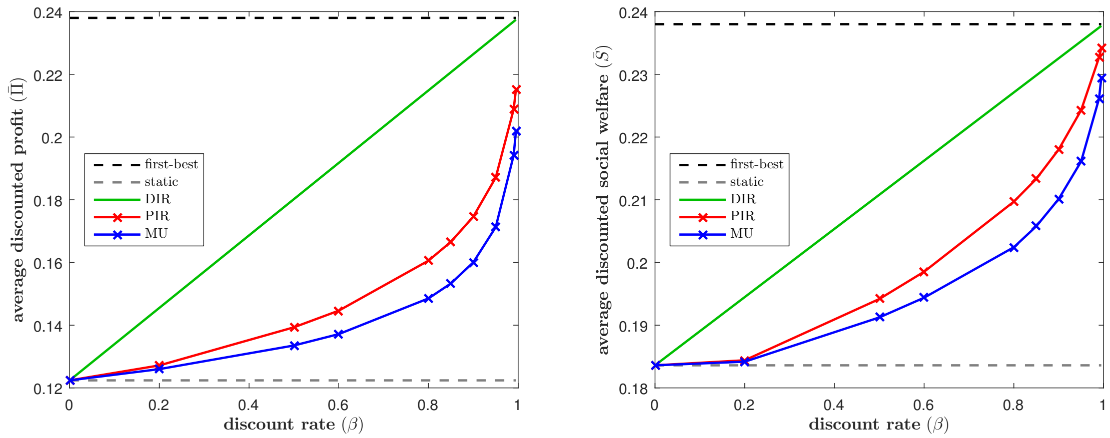

(a) Profit

(b) Social Welfare

Figure 4: Plots of seller's average profit and average social welfare of optimal mechanisms as a function of the discount rate. Values are $U[0,1]$ and the seller's cost is $c = 0.3$.

Figure 4a plots the average discounted seller's profit $\bar{\Pi}^{\pi}$ for different mechanisms as a function of the discount rate (we report average discounted performance metrics so that we can compare results across different discount rates). Consistent with our theoretical results, we observe that the seller's optimal profit under dynamic individual rationality $\bar{\Pi}^{\mathrm{DIR}}$, periodic individual rationality $\bar{\Pi}^{\mathrm{PIR}}$, and martingale utilities $\bar{\Pi}^{\mathrm{MU}}$ converge to first-best. The theoretical analyses of Theorem 4.3 and Theorem 5.3 show that rate of convergence to first-best of $\bar{\Pi}^{\mathrm{PIR}}$ and $\bar{\Pi}^{\mathrm{MU}}$ is $(1-\beta)^{1/2}$. Furthermore, Theorem 4.4 implies that this convergence rate is tight up to polylogarithmic factors. From our numerical experiments we observe that the loss in performance of our approximation mechanisms $\bar{\Pi}^{\mathrm{DUM}}$ and $\bar{\Pi}^{\mathrm{MUM}}$ relative to the respective optimal values $\bar{\Pi}^{\mathrm{PIR}}$ and $\bar{\Pi}^{\mathrm{MU}}$ are not substantial when the discount rate is large (see Figure 5). Even though the convergence rates to first-best of the Deferred Utility Mechanism and the Martingale Utility Mechanism are similar up to polylogarithmic factors, we observe that in practice the expected profit of the Martingale Utility Mechanism is lower, as expected, because of the additional martingale utility constraint. Finally, the profit improvement relative to the optimal static auction is substantial in this particular instance: the relative difference is $(\Pi^{\mathrm{FB}} - \Pi^{\mathrm{S}})/\Pi^{\mathrm{S}} \approx 94\%$ and the seller can benefit considerably by adopting a dynamic mechanism when agents are patient.

Figure 4b plots the average discounted social welfare for different mechanisms as a function of the discount rate, where the average discounted social welfare of dynamic mechanism $\pi \in \mathbb{M}$ is given by $\bar{S}^{\pi} = \bar{\Pi}^{\pi} + \bar{U}^{\pi}$. We observe that the social welfare of an optimal mechanism under dynamic individual rationality $\bar{S}^{\mathrm{DIR}}$, periodic individual rationality $\bar{S}^{\mathrm{PIR}}$, and martingale utilities $\bar{S}^{\mathrm{MU}}$ converge to first-best. This is a straightforward consequence of the fact that under these mechanisms the seller's

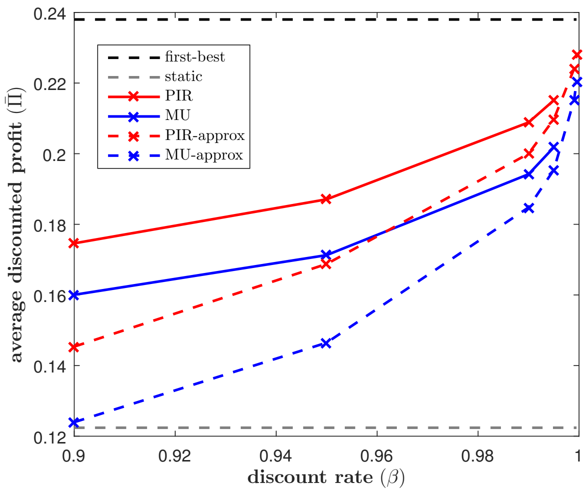

Figure 5: Comparison of the seller's average profit from an optimal mechanism (obtained from dynamic programming) and the approximation mechanism for different models, as a function of the discount rate. Values are $U[0, 1]$ and the seller's cost is $c = 0.3$. The approximation mechanisms' performance is estimated using Monte Carlo simulation for discount rates $\beta \in \{0.9, 0.95, 0.99, 0.995, 0.999, 0.9995\}$, while the dynamic programs are solved via value iteration for discount rates $\beta \in \{0.9, 0.95, 0.99, 0.995\}$. Dynamic programs for larger discount rates were challenging to solve numerically.

profit converges to first-best and the participation constraint guarantees that the buyer's utility is non-negative. Thus, social efficiency is achieved at the expense of the buyer; in the limit the entire social surplus is captured by the seller.

Figure 6 compares the drift $\delta(w) = \mathbb{E}_v [u(v; w)] - w$, flow profit $\rho(w) = \mathbb{E}_v [z(v; w) - cq(v; w)]$, flow utility $\ell(w) = \mathbb{E}_v [vq(v; w) - z(v; w)]$, and social welfare $\mathbb{E}_v [(v - c)q(v; w)]$ of an optimal mechanism for different models when the discount rate is $\beta = 0.95$. We first discuss an optimal mechanism under PIR, which is obtained by numerically solving dynamic program (6). The plots of the flow profit and flow utility suggest three clear phases. When the state is low no trade is materialized (akin to the no-trade phase of the approximation mechanism), when the state is high the item is allocated efficiently and the buyer captures the entire surplus (akin to the income phase), and in intermediate states the item is allocated efficiently and the seller captures the entire surplus (akin to the savings phase). As in our approximation mechanism, the initial state is chosen to be low and the state drifts upwards to the absorbing state $\bar{u}$. Thus, under an optimal mechanism the buyer pays her full surplus to the seller during earlier auctions and then the buyer is allowed to purchase the item from the seller at cost. We note that, in contrast to our approximation mechanism, the transition between phases in an optimal mechanism is smooth.

We now turn to an optimal mechanism under MU, which is obtained by numerically solving dynamic

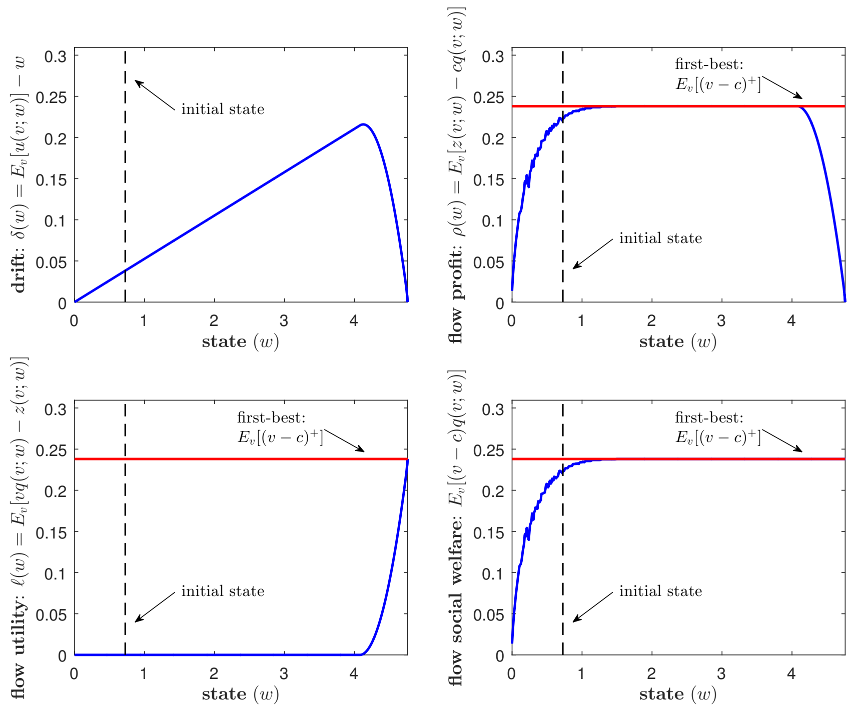

(a) Optimal PIR mechanism

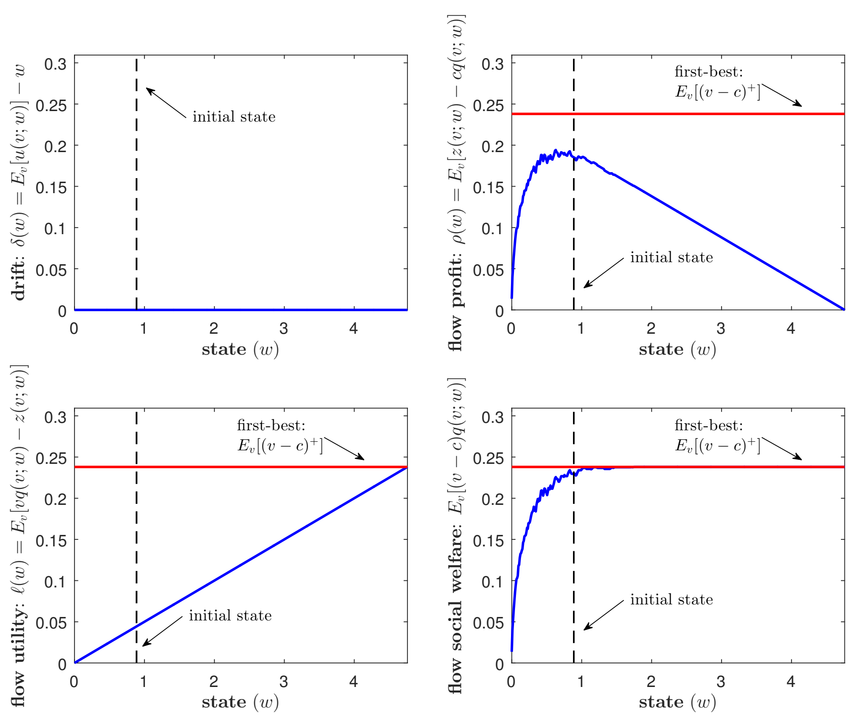

(b) Optimal MU mechanism

Figure 6: Plots of drift, flow profit, flow utility and flow social welfare as a function of the state for an optimal mechanisms under PIR and MU constraints, respectively. Values are $U[0,1]$, the seller's cost is $c = 0.3$, the discount rate is $\beta = 0.95$, $\mathbb{E}_v\left[(v-c)^+\right] \approx 0.238$ and $\underline{u} \approx 4.76$.

program (8). Since the mechanism satisfies the martingale utility constraint (MU), the drift is zero and the flow utility increases linearly with the state (since the promise keeping constraint implies that $\ell(w) = (1-\beta)w - \beta\delta(w)$ and here $\delta(w) = 0$). The plot of the flow profit suggests two phases. When the state is low, trade is materialized with low probability (akin to the throttled phase of the approximation mechanism). Otherwise, the item is allocated according to an auction that yields lower flow profits as the state increases (akin to the allocation phase). Though not shown, the auction in this last phase resembles the soft floor/hard floor auction of our approximation mechanism (with a soft floor that is decreasing with the state). Note that the initial state does not necessarily maximize the initial flow profit. As in our approximation mechanism, the initial state is chosen to balance the potential loss of eventually hitting the throttled phase and the seller's flow profit; which both increase as the state decreases.

7 Conclusions

In this paper, we show that the seller can asymptotically achieve first-best via simple dynamic auctions satisfying appealing business constraints. We have purposely designed our model to be the simplest possible extension of a classical setting to crisply highlight the effects at play. In particular, we assume that there is a single buyer with independent and identically distributed values in an infinite horizon model with discounting. While our analysis is exclusively focused on the case of a single buyer with discounting, we conjecture that our approximation mechanisms and performance results extend to the case of multiple buyers. Extending our results to a finite horizon should be possible. The stationarity of values is not burdensome and can be relaxed at the expense of more complicated mechanisms. The independence of values is critical for our results and is predicated on the fact that, in internet advertising markets, user visiting websites arrive essentially at random, so intertemporal correlation in values (which is driven by the context associated with each impression) is typically weak.

A critical assumption in our model is that the buyer's value distribution is known by the seller. An interesting research direction, which is currently outside the scope of this paper, is to explore the design of practical mechanisms in a prior-free or less prior dependent setting. Another avenue of research is to study the robustness of the proposed mechanisms, in terms of profit and incentive properties, to statistical errors introduced by the estimation of values. We conjecture that the mechanisms herein presented can be suitably modified to remain robust to statistical errors and model misspecification (by imposing, for example, that the buyer is strictly better off reporting truthfully).

Acknowledgements

The authors thank Dirk Bergemann, Omar Besbes, David Brown, Jim Giles, Nitish Korula, Max Lin, Ilan Lobel, David Parkes, Gabriel Weintraub, and Song Zuo for providing helpful comments, as well as seminar participants at the Google Market Algorithms Workshop, Utah Winter Operations Conference and Marketplace Innovation Workshop for their feedback. They also thank the associate editor and three anonymous referees for valuable input that significantly improved the paper. The first author's research was partially supported by a Google Faculty Research Award.

References

Akan, M., Ata, B. and Dana, Jr, J. D. (2015), 'Revenue management by sequential screening', *Journal of Economic Theory* **159, Part B**, 728 – 774.

Ashlagi, I., Daskalakis, C. and Haghpanah, N. (2016), Sequential mechanisms with ex-post participation guarantees, *in* 'Proceedings of the 2016 ACM Conference on Economics and Computation, EC '16, Maastricht, The Netherlands, July 24-28, 2016', pp. 213–214.

Asmussen, S. (2008), *Applied probability and queues*, Vol. 51, Springer Science & Business Media.

Babaioff, M., Immorlica, N., Lucier, B. and Weinberg, S. M. (2014), A simple and approximately optimal mechanism for an additive buyer, *in* '55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014', pp. 21–30.

Balseiro, S., Besbes, O. and Weintraub, G. (2016), Dynamic mechanism design with budget constrained buyers under limited commitment. Working Paper.

Belloni, A., Chen, B. and Sun, P. (2015), Structures of optimal contracts in dynamic mechanism design with one agent. Working Paper.

Bertsekas, D. P. (2012), *Dynamic Programming and Optimal Control*, Vol. 2, 4th edn, Athena Scientific.

Besanko, D. (1985), 'Multi-period contracts between principal and agent with adverse selection', *Economics Letters* **17**(1), 33 – 37.

Bhalgat, A., Feldman, J. and Mirrokni, V. (2012), Online allocation of display ads with smooth delivery, *in* 'Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining', KDD '12, ACM, New York, NY, USA, pp. 1213–1221.

Biais, B., Mariotti, T., Plantin, G. and Rochet, J.-C. (2007), 'Dynamic security design: Convergence to continuous time and asset pricing implications', *The Review of Economic Studies* **74**(2), 345–390.

Board, S. and Skrzypacz, A. (2015), Revenue management with forward-looking buyers. forthcoming, *Journal of Political Economy*.

Chen, Y. and Farias, V. F. (2015), Robust dynamic pricing with strategic customers. Working paper.

Deb, R. and Said, M. (2015), 'Dynamic screening with limited commitment', *Journal of Economic Theory* **159, Part B**, 891 – 928.

eMarketer (2017), 'eMarketer releases new programmatic advertising estimates'. Retrieved from http://www.emarketer.com.

Ëso, P. and Szentes, B. (2007), 'Optimal information disclosure in auctions and the handicap auction', *The Review of Economic Studies* **74**(3), 705–731.

Fiat, A., Karlin, A., Koutsoupias, E. and Vidali, A. (2013), Approaching utopia: Strong truthfulness and externality-resistant mechanisms, *in* 'Proceedings of the 4th Conference on Innovations in Theoretical Computer Science', ITCS '13, ACM, New York, NY, USA, pp. 221–230.

Gallien, J. (2006), 'Dynamic mechanism design for online commerce', *Operations Research* **54**(2), 291–310.

Gershkov, A. and Moldovanu, B. (2014), *Dynamic Allocation and Pricing: A Mechanism Design Approach*, The MIT Press.

Green, E. J. (1987), 'Lending and the smoothing of uninsurable income', *Contractual arrangements for intertemporal trade* **1**, 3–25.

Hart, S. and Nisan, N. (2012), Approximate revenue maximization with multiple items, *in* 'ACM Conference on Electronic Commerce, EC '12, Valencia, Spain, June 4-8, 2012', p. 656.

Jackson, M. O. and Sonnenschein, H. F. (2007), 'Overcoming incentive constraints by linking decisions1', *Econometrica* **75**(1), 241–257.

Kakade, S., Lobel, I. and Nazerzadeh, H. (2013), 'Optimal dynamic mechanism design and the virtual pivot mechanism', *Operations Research* **61**(3), 837–854.

Kanoria, Y. and Nazerzadeh, H. (2017), Dynamic reserve prices for repeated auctions: Learning from bids. Working paper. https://ssrn.com/abstract=2444495.

Krähmer, D. and Strausz, R. (2015), 'Optimal sales contracts with withdrawal rights', *The Review of Economic Studies* **82**(2), 762.

Krishna, R. V., Lopomo, G. and Taylor, C. R. (2013), 'Stairway to heaven or highway to hell: Liquidity, sweat equity, and the uncertain path to ownership', *The RAND Journal of Economics* **44**(1), 104–127.

Laffont, J.-J. and Martimort, D. (2001), *The Theory of Incentives: The Principal-Agent Model*, Princeton University Press.

Laforgia, A. (1991), 'Bounds for modified bessel functions', *Journal of Computational and Applied Mathematics* **34**(3), 263–267.

Lariviere, M. A. and Porteus, E. L. (2001), 'Selling to the newsvendor: An analysis of price-only contracts', *Manufacturing & Service Operations Management* **3**(4), 293–305.

Lobel, I. and Xiao, W. (2016), Optimal long-term supply contracts with asymmetric demand information. Working Paper.

Luenberger, D. G. (1969), *Optimization by Vector Space Methods*, 1st edn, John Wiley & Sons, Inc., New York, NY, USA.

Manelli, A. M. and Vincent, D. R. (2006), 'Bundling as an optimal selling mechanism for a multiple-good monopolist', *Journal of Economic Theory* **127**(1), 1–35.

Marcinkiewicz, J. and Zygmund, A. (1937), 'Quelques théoremes sur les fonctions indépendantes', *Fund. Math* **29**, 60–90.

McDiarmid, C. (1998), *Concentration*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 195–248.

Milgrom, P. (2004), *Putting Auction Theory to Work*, Cambridge University Press.

Milgrom, P. and Segal, I. (2002), 'Envelope theorems for arbitrary choice sets', *Econometrica* **70**(2), 583–601.

Mirrokni, V. S., Leme, R. P., Tang, P. and Zuo, S. (2016), 'Optimal dynamic mechanisms with ex-post IR via bank accounts', *CoRR* **abs/1605.08840**.
URL: *http://arxiv.org/abs/1605.08840*

Muthukrishnan, S. (2009), Ad exchanges: Research issues., *in* 'Internet and Network Economics, Lecture Notes in Computer Science', pp. 1–12.

Myerson, R. (1981), 'Optimal auction design', *Mathematics of Operations Research* **6**(1), 58–73.

Papadimitriou, C. H., Pierrakos, G., Psomas, C. and Rubinstein, A. (2014), 'The intractability of dynamic mechanism design', *CoRR* **abs/1407.5373**.

Pavan, A., Segal, I. and Toikka, J. (2014), 'Dynamic mechanism design: A myersonian approach', *Econometrica* **82**(2), 601–653.

Spear, S. E. and Srivastava, S. (1987), 'On repeated moral hazard with discounting', *The Review of Economic Studies* **54**(4), 599–617.

Thomas, J. and Worrall, T. (1990), 'Income fluctuation and asymmetric information: An example of a repeated principal-agent problem', *Journal of Economic Theory* **51**(2), 367 – 390.

Vulcano, G., van Ryzin, G. and Maglaras, C. (2002), 'Optimal dynamic auctions for revenue management', *Management Science* **48**(11), 1388–1407.

Yao, A. C. (2015), An n-to-1 bidder reduction for multi-item auctions and its applications, *in* 'Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015', pp. 92–109.

A Proof of selected results

A.1 Proof of Theorem 4.4

Suppose that the seller's cost is $c = 0$ and the buyer's value is v_1 with probability f_1 and v_2 with probability $f_2 = 1 - f_1$. Because values are non-degenerate and positive we have that $0 < v_1 < v_2$ and $f_1, f_2 \in (0, 1)$.

We prove the result in four steps. First, we use structural properties of the optimal value function to simplify the optimal mechanism design problem. Second, we provide an upper bound to the simplified problem using a perfect information relaxation in which the seller knows in advance all realizations of the buyer's value. In this step we characterize the optimal objective value of the perfect information relaxation in terms of the expected deviations from the origin of a reflected random walk. Third, we use the perfect information bound to control the cost of imposing PIR. Fourth, we conclude by bounding the latter stochastic deviations using concentration inequalities.

Step 1 The optimal value function under PIR satisfies the Bellman equation

$$
\begin{aligned}
\Pi(w) = \max_{(q_i, z_i, u_i)_i} \ & \sum\nolimits_i f_i \left(z_i + \beta \Pi(u_i) \right) \\
\text{s.t } & w = \sum\nolimits_i f_i (v_i q_i - z_i + \beta u_i)\,, && \text{(PK)} \\
& v_i q_i - z_i \geq 0\,, \quad \forall i\,, && \text{(PIR)} \\
& v_i q_i - z_i + \beta u_i \geq v_i q_j - z_j + \beta u_j\,, \quad \forall i, j \neq i\,, && \text{(DIC)} \\
& 0 \leq q_i \leq 1\,,\ 0 \leq u_i \leq \bar{u}\,, \quad \forall i\,,
\end{aligned}
$$

with boundary conditions $\Pi(0) = 0$ and $\bar{u} = f_1 v_1 + f_2 v_2$. The optimal objective value is given by $\Pi^{\text{PIR}} = \max_{w \in \mathcal{U}} \Pi(w)$.

The following result is a discrete analogue of Lemma 2.1. We state the result without proof as this follows directly from Lemma 1 in Krishna et al. (2013).

Lemma A.1. *Let $U_i = v_i q_i - z_i + \beta u_i$ be the utility-to-go of type $i = 1, 2$. Without loss of optimality we can replace the dynamic incentive compatibility constraint by $q_1 \le q_2$ and $U_2 = U_1 + (v_2 - v_1)q_1$.*

Combining the previous lemma with the promise keeping constraint $w = f_1 U_1 + f_2 U_2$, we obtain that $U_1 = w - f_2(v_2 - v_1)q_1$ and $U_2 = w + f_1(v_2 - v_1)q_1$ because $f_1 + f_2 = 1$. Note that (PIR) can be written as $u_i \le U_i/\beta$. Let $S(w) = w + \Pi(w)$ be the social welfare generated by an optimal mechanism. Using the promised keeping constraint to eliminate payments from the objective we obtain the new Bellman equation

$$\begin{aligned} S(w) = \max_{(q_i, u_i)_i} & \sum\nolimits_i f_i \left(v_i q_i + \beta S(u_i) \right) \\ \text{s.t } & U_1 = w - f_2(v_2 - v_1)q_1 \,, \\ & U_2 = w + f_1(v_2 - v_1)q_1 \,, \\ & 0 \le u_i \le U_i/\beta \,, \quad \forall i \,, \\ & 0 \le q_1 \le q_2 \le 1 \,, u_i \le \bar{u} \,, \quad \forall i \,. \end{aligned} \tag{PIR'}$$

Because the objective is increasing in the allocation of the highest type, we obtain that $q_2 = 1$, i.e., there is no distortion at the top. Note that $S(w)$ is non-decreasing in w because the constraint set of the inner optimization problem expands as w increases. Thus, for any fixed allocation q_1 the seller would like to set the promise function as large as possible and (PIR') should be binding whenever $u_i \le \bar{u}$. This implies that $u_i = \min(U_i/\beta, \bar{u})$.

Consider the two-point random variable

$$\xi_t = \begin{cases} -f_2(v_2 - v_1) \,, & \text{w.p. } f_1 \,, \\ f_1(v_2 - v_1) \,, & \text{w.p. } f_2 \,. \end{cases}$$

The previous discussion implies that the state evolves according to $w_{t+1} = \min\left((w_t + \xi_t q_t)/\beta, \bar{u}\right)$, where with some abuse we denote by $q_t \triangleq q_{1,t}$ the allocation of the lowest type at time t (since we know that the highest type is always allocated the item). The seller's problem can be cast as the following stochastic control problem

$$\begin{aligned} S(w) = \frac{f_2 v_2}{1 - \beta} + \max_{\mathbf{q} \in \mathcal{Q}} & \, \mathbb{E}\left[\sum_{t=1}^{\infty} f_1 v_1 \beta^{t-1} q_t \right] \\ \text{s.t } & w_{t+1} = \min\left((w_t + \xi_t q_t)/\beta, \bar{u} \right) \,, \\ & w_1 = w \,, \\ & w_t \ge 0 \,,\; 0 \le q_t \le 1 \,, \end{aligned}$$

where $\mathcal{Q}$ denotes the set of all adaptive, non-anticipative policies $\mathbf{q} = (q_t)_{t=1}^{\infty}$ that map a history at time t to an allocation q_t of the lowest type. Note that the state $\bar{u}$ is absorbing: since $\frac{1}{\beta}\left(\bar{u} - f_2(v_2 - v_1)\right) > \bar{u}$ we have that if $w_t = \bar{u}$, then $w_{t+1} = \bar{u}$ regardless of q_t and ξ_t. As a consequence, if the state hits the upper bound $\bar{u}$, it is optimal to set $q_t = 1$ thereafter. This observation implies that we can replace the dynamics by $w_{t+1} = (w_t + \xi_t q_t)/\beta$ without loss of optimality (because $w_{t+1} \ge \bar{u}$ if $w_t \ge \bar{u}$).

Step 2 Consider a perfect information relaxation in which the decision maker has access to all realizations of the random variables $\boldsymbol{\xi} = (\xi_t)_{t=1}^{\infty}$. Given a sample path $\boldsymbol{\xi} \in \mathbb{R}^{\infty}$ we can calculate the optimal value for the sample path in "hindsight" by solving a deterministic linear program. The expected value with perfect information provides an upper bound on $S(w)$. More formally, we denote by $S^{\mathrm{H}}(\boldsymbol{\xi}; w)$ the optimal (deterministic) vale of the perfect information problem for sample path $\boldsymbol{\xi}$, where H stands for hindsight. We then have $S(w) \leq \mathbb{E}\left[S^{\mathrm{H}}(\boldsymbol{\xi}; w)\right]$, where the perfect information problem is given by

$$
\begin{aligned}
S^{\mathrm{H}}(\boldsymbol{\xi}; w) = \frac{f_2 v_2}{1-\beta} + \max_{(w_t, q_t)_{t=1}^{\infty}} & \sum_{t=1}^{\infty} f_1 v_1 \beta^{t-1} q_t \\
\text{s.t } & w_{t+1} = (w_t + \xi_t q_t)/\beta \,, \\
& w_1 = w \,, \\
& w_t \geq 0 \,, \; 0 \leq q_t \leq 1 \,.
\end{aligned}
$$

It is not hard to see that the perfect information problem admits a simple optimal solution: greedily set q_t as large as possible. This implies that $q_t = 1$ whenever $\xi_t > 0$ and $q_t = \min\{1, w_t/(f_2(v_2 - v_1))\}$ when $\xi_t < 0$. Consider the scaled stochastic process $\hat{w}_t = \beta^{t-1} w_t$. Because $\hat{w}_{t+1} = \hat{w}_t + \beta^{t-1}\xi_t q_t$, this process evolves according to the reflected geometric random walk

$$
\hat{w}_{t+1} = \max\{\hat{w}_t + \beta^{t-1}\xi_t, 0\} \,.
$$

We next provide a closed-form expression for the optimal objective value of the perfect information problem using Skorokhod's map for reflected random walks. Using that $q_t = 1$ whenever $\xi_t > 0$, we obtain that the optimal objective value is given by

$$
S^{\mathrm{H}}(\boldsymbol{\xi}; w) = \frac{f_2 v_2}{1-\beta} + f_1 v_1 \sum_{t=1}^{\infty} \beta^{t-1} \mathbf{1}\{\xi_t > 0\} + f_1 v_1 \sum_{t=1}^{\infty} \beta^{t-1} q_t \mathbf{1}\{\xi_t < 0\} \,.
$$

We can eliminate the third term by considering the state dynamics. These are given by

$$
\hat{w}_{s+1} - w = \sum_{t=1}^{s} \beta^{t-1} q_t \xi_t = f_1(v_2 - v_1) \sum_{t=1}^{s} \beta^{t-1} \mathbf{1}\{\xi_t > 0\} - f_2(v_2 - v_1) \sum_{t=1}^{s} \beta^{t-1} q_t \mathbf{1}\{\xi_t < 0\} \,.
$$

Because $\beta \in (0,1)$, the series on the right-hand side are absolutely convergent and $\hat{w}_\infty \triangleq \lim_{s\to\infty} \hat{w}_s$ is finite. Letting $s \to \infty$, multiplying the second equation by $\rho \triangleq f_1 v_1/(f_2(v_2 - v_1)) > 0$, and adding these last two equations together we obtain that

$$
S^{\mathrm{H}}(\boldsymbol{\xi}; w) = \frac{f_2 v_2}{1-\beta} + \frac{f_1 v_1}{f_2} \sum_{t=1}^{\infty} \beta^{t-1} \mathbf{1}\{\xi_t > 0\} - \rho\left(\hat{w}_\infty - w\right) \,.
$$

Let $X_t = -\sum_{s=1}^{t} \beta^{s-1} \xi_s$ be the state of the geometric random walk by time t. Skorokhod's map (see, e.g., Asmussen, 2008, ch. IX.2) implies that the reflected random walk satisfies $\hat{w}_{t+1} = w - X_t + \max_{1\leq s\leq t}(X_s - w)^+$. Taking expectations we obtain that

$$
\mathbb{E}\left[S^{\mathrm{H}}(\boldsymbol{\xi}; w)\right] = \Pi^{\mathrm{FB}} - \rho \mathbb{E}\left[\max_{s\geq 1}(X_s - w)^+\right] \,, \tag{11}
$$

where we used that $\mathbb{P}\{\xi_t > 0\} = f_2$; $\Pi^{\mathrm{FB}} = (f_1 v_1 + f_2 v_2)/(1-\beta)$; and $\mathbb{E}\left[\lim_{t\to\infty} X_t\right] = 0$ from Dominated Convergence Theorem because $\beta \in (0,1)$ and $\mathbb{E}[\xi_t] = 0$.

Step 3 We now proceed to bound the cost of imposing PIR. Using that $S(w) = w + \Pi(w)$ and (11) we obtain

$$\Pi^{\mathrm{FB}} - \Pi(w) = \Pi^{\mathrm{FB}} + w - S(w) \geq \Pi^{\mathrm{FB}} + w - \mathbb{E}\left[S^{\mathrm{H}}(\boldsymbol{\xi}; w)\right] \geq w + \rho\mathbb{E}\left[\max_{s\geq 1}(X_s - w)^+\right] .$$

The second term can be further bounded from below as follows

$$\mathbb{E}\left[\max_{s\geq 1}(X_s - w)^+\right] \geq \mathbb{E}\left[(X_\infty - w)^+\right] \geq \left(\mathbb{E}\left[X_\infty^+\right] - w\right)^+ ,$$

where the second inequality follows because $(X_\infty - w)^+ = (X_\infty^+ - w)^+$ because $w \geq 0$ together with Jensen's inequality. Using that $\Pi^{\mathrm{PIR}} = \max_{w\geq 0} \Pi(w)$ we obtain the lower bound

$$\Pi^{\mathrm{FB}} - \Pi^{\mathrm{PIR}} \geq \min_{w\geq 0}\left\{w + \rho\left(\mathbb{E}\left[X_\infty^+\right] - w\right)^+\right\} = \min(\rho, 1)\mathbb{E}\left[X_\infty^+\right] ,$$

where the last equality follows because the minimum of the right-hand is achieved at $w = 0$ if $\rho \in (0, 1)$ and at the break-point $w = \mathbb{E}\left[X_\infty^+\right]$ if $\rho \in [1, \infty)$.

Step 4 We next lower bound the expectation $\mathbb{E}\left[X_\infty^+\right]$. Because the random variable X_∞ is mean zero, we obtain that $\mathbb{E}\left[X_\infty^+\right] = \mathbb{E}\left|X_\infty\right|/2$. Let $Q_t = \sum_{s=1}^{t}\left(\beta^{s-1}\xi_s\right)^2$ be the quadratic variation process. Marcinkiewicz-Zygmund inequality implies that there exists a constant $c_1 > 0$ independent of t such that $\mathbb{E}\left|X_t\right| \geq c_1\mathbb{E}\left|Q_t^{1/2}\right|$ for all t (Marcinkiewicz and Zygmund, 1937). Because the random steps are lower bounded by $|\xi_s| \geq (v_2 - v_1)\min(f_1, f_2)$, we have that the quadratic variation process is lower bounded by $Q_t^{1/2} \geq (v_2 - v_1)\min(f_1, f_2)\left((1-\beta^{2t})/(1-\beta^2)\right)^{1/2}$. Since $\beta \in (0, 1)$, we obtain from Dominated Convergence Theorem

$$\mathbb{E}\left|X_\infty\right| = \lim_{t\to\infty}\mathbb{E}\left|X_t\right| \geq c_1 \liminf_{t\to\infty}\mathbb{E}\left|Q_t^{1/2}\right| \geq c_1(v_2 - v_1)\min(f_1, f_2)\left(1-\beta^2\right)^{-1/2} .$$

Therefore, there exists some constant $c_2 > 0$ such that $\mathbb{E}\left[X_\infty^+\right] \geq c_2(1-\beta)^{-1/2}$ since $(1-\beta^2) = (1-\beta)(1+\beta)$ and the result follows.

A.2 Proof of Proposition 5.2

First note that $0 \leq \underline{w} \leq \bar{u}$ because $\beta \in (0, 1)$. Recall that the soft floor $s^{\mathrm{MUM}}(w)$ is a solution of equation (10). Note that $G(s) \triangleq \mathbb{E}_v[(v-s)^+]$ can be written as $G(s) = \int_s^{\bar{v}} \bar{F}(x)\mathrm{d}x$, and thus it is decreasing and continuous (because the density is strictly positive). Additionally, we have that $G(c) = \mathbb{E}_v[(v-c)^+] = (1-\beta)\bar{u}$ and $G(\bar{v}) = 0$. Thus we obtain that for all $w \in [0, \bar{u}]$ there exists a unique $s^{\mathrm{MUM}}(w)$ solving the latter equation, and the soft floor satisfies that $c \leq s^{\mathrm{MUM}}(w) \leq \bar{v}$.

Periodic individual rationality follows because payments are never larger than values for all $w \in [0, \bar{u}]$. The martingale constraint follows by construction when $w \in [0, \underline{w})$. When $w \in [\underline{w}, \bar{u}]$ we have that

$$\begin{aligned}
\mathbb{E}_v\left[u^{\text{MUM}}(v;w)\right] &= \frac{1}{\beta}\left(w + \mathbb{E}_v\left[\left(\min(v, s^{\text{MUM}}(w)) - c\right)^+\right] - \mathbb{E}_v[(v-c)^+]\right) \\
&= \frac{1}{\beta}\left(w - \mathbb{E}_v\left[(v - \min(v, s^{\text{MUM}}(w)))\,\mathbf{1}\{v \geq c\}\right]\right) \\
&= \frac{1}{\beta}\left(w - \mathbb{E}_v\left[(v - s^{\text{MUM}}(w))^+\mathbf{1}\{v \geq c\}\right]\right) \\
&= \frac{1}{\beta}\left(w - \mathbb{E}_v\left[(v - s^{\text{MUM}}(w))^+\right]\right) = w\,,
\end{aligned}$$

where we used that $x - \min(x,y) = (x-y)^+$ together with the fact that $s^{\text{MUM}}(w) \geq c$ and that the soft floor is chosen so to be solution of equation (10).

In order for the mechanism to be dynamic incentive compatible it suffices to show that it satisfies (2) and (3) because the allocation is trivially non-decreasing. Equation (3) follows from the definition of the mechanism together with the fact that $G(c) = \mathbb{E}_v\left[(v-c)^+\right]$. Constraint (2) is trivially satisfied when $w \in [0, \underline{w})$. When $w \in [\underline{w}, \bar{u}]$ we use that the promise function is non-decreasing in v to obtain

$$u^{\text{MUM}}(v;w) \geq u^{\text{MUM}}(0;w) = \frac{1}{\beta}\left(w - \mathbb{E}_v\left[(v-c)^+\right]\right) \geq \frac{1}{\beta}\left(\underline{w} - \mathbb{E}_v\left[(v-c)^+\right]\right) = 0\,,$$

because $w \geq \underline{w}$ and $\underline{w} = \mathbb{E}_v\left[(v-c)^+\right]$. Because $u^{\text{MUM}}(v;w)$ is non-decreasing in v it suffices to show that $u^{\text{MUM}}(\bar{v};w) \leq \bar{u}$. Using the formula for the soft floor $s^{\text{MUM}}(w)$ to eliminate w together with $\bar{u} = \mathbb{E}_v\left[(v-c)^+\right]/(1-\beta)$ and $s^{\text{MUM}}(w) \geq c$ can write the condition that $u^{\text{MUM}}(\bar{v};w) \leq \bar{u}$ as follows

$$\mathbb{E}_v\left[(v - s^{\text{MUM}}(w))^+\right] + (1-\beta)(s^{\text{MUM}}(w) - c) \leq \mathbb{E}_v\left[(v-c)^+\right]\,.$$

The expression on the left-hand side is convex in $s^{\text{MUM}}(w)$ and attains its maximum at one of its end points $s^{\text{MUM}}(w) = c$ or $s^{\text{MUM}}(w) = \bar{v}$. The inequality trivially holds when $s^{\text{MUM}}(w) = c$. Evaluating at $s^{\text{MUM}}(w) = \bar{v}$ we obtain $(1-\beta)(\bar{v} - c) \leq \mathbb{E}_v\left[(v-c)^+\right]$, which is implied by Assumption 5.1.

A.3 Proof of Theorem 5.3

Let $\{w_t^{\text{MUM}}\}_{t=1}^{\infty}$ denote the stochastic process that governs the evolution of the state under the approximation mechanism $(q^{\text{MUM}}, z^{\text{MUM}}, u^{\text{MUM}})$. This process evolves according to $w_t^{\text{MUM}} = u^{\text{MUM}}(v_{t-1}, w_{t-1}^{\text{MUM}})$ with initial condition $w^{\text{MUM}} \in [\underline{w}, \bar{u}]$. Let $\tau = \inf\{t \geq 1 : w_t^{\text{MUM}} \leq \underline{w}\}$ be the first time that the state falls below $\underline{w}$. We have that during time $t = 1, \ldots, \tau - 1$ the dynamic mechanism allocates according to hard floor/soft floor auction.

By construction, the mechanism guarantees that the promised utility stochastic process $\{w_t^{\text{MUM}}\}_{t=1}^{\infty}$ is a martingale, that is, the process satisfies $\mathbb{E}[w_{t+1}^{\text{MUM}} | w_t^{\text{MUM}}] = w_t^{\text{MUM}}$. The next result characterizes some properties of this process.

Lemma A.2. *Under the approximation mechanism $(q^{\text{MUM}}, z^{\text{MUM}}, u^{\text{MUM}})$,*

1. *Let $b = \beta^2(w^{\text{MUM}} - \underline{w})^2/(4\bar{v}^2)$. The c.d.f. of the stopping time τ satisfies*

$$\mathbb{P}(\tau < t) \leq \exp\left(-\frac{b}{t-1}\right)\,.$$

2. Suppose $b \geq 1/(4(1-\beta))$. The stopping time τ satisfies

$$\mathbb{E}[\beta^\tau] \leq \big(1+8b\big)(1-\beta)\exp\left(-2\big((1-\beta)b\big)^{1/2}\right) .$$

Let $L(w) = \mathbb{E}_v\left[(v-c)^+\right] - \mathbb{E}_v\left[(\min(v, s^{\mathrm{MUM}}(w)) - c)\,\mathbf{1}\{v \geq c\}\right]$ be the difference in expected profit between the first-best allocation and the mechanism that allocates according to the hard floor/soft floor auction when the promised utility is w. Because the soft floor is chosen according to (10) we obtain

$$L(w) = \mathbb{E}_v\left[(v - \min(v, s^{\mathrm{MUM}}(w)))\,\mathbf{1}\{v \geq c\}\right] = \mathbb{E}_v\left[(v - s^{\mathrm{MUM}}(w))^+\mathbf{1}\{v \geq c\}\right] = (1-\beta)w\,, \tag{12}$$

where we used that $x - \min(x,y) = (x-y)^+$ together with the fact that $s^{\mathrm{MUM}}(w) \geq c$ because $w \leq \underline{u}$.

We are now in position to prove the main result. We have that $\Pi^{\mathrm{MUM}} \leq \Pi^{\mathrm{FB}}$, because the mechanism in consideration is primal feasible and thus is bounded by first-best. In the remainder of the proof we prove the last inequality, that is

$$\Pi^{\mathrm{MUM}} \geq \Pi^{\mathrm{FB}} - \tilde{O}\left(\frac{1}{(1-\beta)^{1/2}}\right) .$$

Step 1 The expected performance of approximation mechanism can be decomposed as follows:

$$\begin{aligned}
\Pi^{\mathrm{MUM}} &= \sum_{t=1}^{\infty} \mathbb{E}[\beta^{t-1}(z^{\mathrm{MUM}}(v_t;w_t) - cq^{\mathrm{MUM}}(v_t;w_t))] \\
&= \frac{\mathbb{E}_v\left[(v-c)^+\right]}{1-\beta} - \sum_{t=1}^{\infty}\mathbb{E}\Big[\underbrace{\beta^{t-1}\left((v_t-c)^+ - z^{\mathrm{MUM}}(v_t;w_t) + cq^{\mathrm{MUM}}(v_t,w_t)\right)}_{R_t}\Big] \\
&= \Pi^{\mathrm{FB}} - \sum_{t=1}^{\infty}\mathbb{E}[R_t]\,,
\end{aligned}$$

where the second equation follows because values are i.i.d. and using that $\sum_{t=0}^{\infty}\beta^{t-1} = 1/(1-\beta)$. The error terms R_t measure the difference in expected performance between the first-best allocation and the mechanism in consideration. We can decompose the error terms R_t as follows

$$\sum_{t=1}^{\infty}\mathbb{E}[R_t] = \underbrace{\sum_{t=1}^{\infty}\mathbb{E}[R_t\mathbf{1}\{t<\tau\}]}_{E_1} + \underbrace{\sum_{t=1}^{\infty}\mathbb{E}[R_t\mathbf{1}\{t\geq\tau\}]}_{E_2} .$$

In the remainder of the proof we upper bound the error terms E_1 and E_2, respectively.

Step 2 We next upper bound the error term E_1. Because when $t < \tau$ we have that $w_t^{\text{MUM}} \in [\underline{w}, \bar{u}]$, and we obtain

$$
\begin{aligned}
E_1 &= \sum_{t=1}^{\infty} \mathbb{E}[R_t \mathbf{1}\{t < \tau\}] = \sum_{t=1}^{\infty} \mathbb{E}\left[\beta^{t-1}\left((v_t - c)^+ - (\min(v_t, s^{\text{MUM}}(w_t^{\text{MUM}})) - c)\,\mathbf{1}\{v_t \ge c\}\right)\mathbf{1}\{t < \tau\}\right] \\
&= \sum_{t=1}^{\infty} \beta^{t-1}\mathbb{E}\left[L\left(w_t^{\text{MUM}}\right)\mathbf{1}\{t < \tau\}\right] = (1-\beta)\sum_{t=1}^{\infty} \beta^{t-1}\mathbb{E}\left[w_t^{\text{MUM}}\mathbf{1}\{t < \tau\}\right] \\
&\le (1-\beta)\sum_{t=1}^{\infty} \beta^{t-1}\mathbb{E}\left[w_{\tau\wedge t}^{\text{MUM}}\right] = w^{\text{MUM}}(1-\beta)\sum_{t=1}^{\infty}\beta^{t-1} = w^{\text{MUM}}\,,
\end{aligned}
$$

where the second equality follows because the mechanism allocates according to the hard floor/soft floor auction when $t < \tau$, the third equality because values are i.i.d. and τ is a stopping time together with the definition of $L(\cdot)$, the third equality from (12), the inequality because the stopped martingale is non-negative because $u^{\text{MUM}}(v; w) \ge 0$, the fifth equality from the Optional Stopping Theorem because the stopping time $\tau \wedge t$ is bounded, and the last from the geometric series $\sum_{t=0}^{\infty} \beta^t = 1/(1-\beta)$.

Step 3 We next bound the error term E_2. We have

$$
E_2 = \sum_{t=1}^{\infty} \mathbb{E}[R_t \mathbf{1}\{t \ge \tau\}] \le \bar{v}\sum_{t=1}^{\infty}\mathbb{E}[\beta^{t-1}\mathbf{1}\{t \ge \tau\}] = \bar{v}\beta\mathbb{E}\left[\sum_{t=\tau}^{\infty}\beta^t\right] = \frac{\bar{v}\beta}{1-\beta}\mathbb{E}\left[\beta^{\tau}\right]\,,
$$

where the first inequality follows because $|R_t| \le \beta^{t-1}\bar{v}$ since the flow profit is at most $\bar{v}$, and the second equality from Tonelli's Theorem.

Step 4 Putting everything together and using Lemma A.2 we obtain that the error terms are bounded by

$$
E_1 + E_2 \le w^{\text{MUM}} + \bar{v}\big(1 + 8b\big)\beta\exp\left(-2\big((1-\beta)b\big)^{1/2}\right)\,,
$$

with $b = \beta^2(w^{\text{MUM}} - \underline{w})^2/(4\bar{v}^2)$ such that $b \ge 1/(4(1-\beta))$. We balance the error terms by setting the initial promised utility to $w^{\text{MUM}} = \underline{w} + \frac{\bar{v}}{\beta}\frac{1}{(1-\beta)^{1/2}}\log\left(\frac{1}{1-\beta}\right)$. With this choice we have that $b \ge 1/(4(1-\beta))$ for $\beta \ge 1 - 1/e$. This leads to

$$
E_1 + E_2 \le \tilde{O}\left(\frac{1}{(1-\beta)^{1/2}}\right)\,,
$$

and the result follows.

Electronic Companion:

Dynamic Mechanisms with Martingale Utilities

Santiago R. Balseiro, Vahab S. Mirrokni, Renato Paes Leme

November 27, 2017

B Additional Proofs

B.1 Proof of Lemma 2.1

Fix the state $w \in \mathcal{W}$. To simplify the exposition we drop the dependence on the state.

Only if part Let $U(v) = vq(v) - z(v) + \beta u(v)$ be the expected utility-to-go of the buyer when her realized value is v. From the dynamic incentive compatibility constraint (DIC) we have that for any feasible mechanism $U(v) = \max_{\tilde{v}} vq(\tilde{v}) - z(\tilde{v}) + \beta u(\tilde{v})$. Applying the Envelope Theorem we obtain:

$$\frac{\mathrm{d}U(v)}{\mathrm{d}v} = q(v)\,.$$

This implies that the allocation $q(v)$ is non-decreasing in the value v because the $U(v)$ is convex in the value v (since it is the maximum of linear functions). Convexity of $U(v)$ implies absolute continuity and integrating we obtain that:

$$U(v) = U(0) + \int_0^v q(x)\ \mathrm{d}x\,. \tag{B-13}$$

Using the definition of $U(v)$ we get that the envelope condition corresponding to the DIC constraint is given by:

$$vq(v) - z(v) + \beta u(v) = U(0) + \int_0^v q(x)\ \mathrm{d}x\,. \tag{B-14}$$

where $U(0) = -z(0) + \beta u(0)$ is the utility of the lowest type.

The promise keeping constraint (PK) gives that $w = \mathbb{E}[U(v)]$ which combined with equation (B-13) implies that the utility of the lowest type is given by

$$U(0) = w - \mathbb{E}_v\left[\int_0^v q(x)\ \mathrm{d}x\right] = w - \int_0^\infty \bar{F}(x)q(x)\ \mathrm{d}x\,,$$

where the last equation follows from integrating by parts. Solving for payments in (B-14) and using the previous equality we obtain

$$z(v) = \beta u(v) - w + vq(v) - \int_0^v q(x)\ \mathrm{d}x + \int_0^\infty \bar{F}(x)q(x)\ \mathrm{d}x$$

and the result follows.

If part We need to show that constraints (PK) and (DIC) are satisfied. Using (3) in the promise keeping constraint gives

$$\mathbb{E}_v\left[vq(v) - z(v) + \beta u(v;w)\right] = w + \mathbb{E}_v\left[\int_0^v q(x)\mathrm{d}x\right] - \int_0^\infty \bar{F}(x)q(x)\mathrm{d}x = w\,,$$

where the last equation follows from integrating by parts the second term and cancelling terms.

Further, the monotonicity of the allocation together with the envelope condition are necessary and sufficient conditions for incentive compatibility (see, e.g., Theorem 4.3 in Milgrom (2004)). We reproduce the result for completeness. Let $U(v,\tilde{v}) = vq(\tilde{v}) - z(\tilde{v}) + \beta u(\tilde{v})$ be the expected utility-to-go of the buyer when her realized value is v and she reports $\tilde{v}$. For the dynamic incentive compatibility constraint to hold we need to show that $U(v,v) \geq U(v,\tilde{v})$ for all $v,\tilde{v}$. Using (3) we obtain that

$$U(v,\tilde{v}) = w + (v-\tilde{v})q(\tilde{v}) + \int_0^{\tilde{v}} q(x)\mathrm{d}x - \int_0^\infty \bar{F}(x)q(x)\mathrm{d}x\,.$$

When $\tilde{v} \leq v$ we have that

$$U(v,v) - U(v,\tilde{v}) = \int_{\tilde{v}}^v q(x)\mathrm{d}x - (v-\tilde{v})q(\tilde{v}) \geq 0\,,$$

where the equality follows from cancelling constant terms and the inequality because $q(x)$ is non-decreasing. A similar result holds for $\tilde{v} \geq v$ and the result follows.

B.2 Proof of Proposition 4.2

First note that $\underline{w} \leq \bar{w}$ from Assumption 4.1. Periodic individual rationally follows because payments are at most the value.

In order for the mechanism to be dynamic incentive compatible it suffices to show that it satisfies (2) and (3) because the allocation is trivially non-decreasing. Equation (3) follows trivially when $w \in [0,\underline{w})$. When $w \in [\underline{w},\bar{w}]$ we have that (3) can be written as

$$\begin{aligned} z^{\text{DUM}}(v;w) &= \beta u^{\text{DUM}}(v;w) - w + vq^{\text{DUM}}(v;w) - \int_0^v q^{\text{DUM}}(x;w)\ \mathrm{d}x + \int_0^\infty \bar{F}(x)q^{\text{DUM}}(x;w)\ \mathrm{d}x \\ &= \Delta^{\text{DUM}}(v) + v\mathbf{1}\{v \geq c\} - (v-c)^+ + \mathbb{E}_v[(v-c)^+] = v\mathbf{1}\{v \geq c\}\,, \end{aligned}$$

where the second equality follows because $\int_c^\infty \bar{F}(x)\ \mathrm{d}x = \mathbb{E}_v[(v-c)^+]$ and the last because $\Delta^{\text{DUM}}(v) = (v-c)^+ - \mathbb{E}_v[(v-c)^+]$. When $w \in (\bar{w},\bar{u}]$ we obtain that (3) can be written as

$$z^{\text{DUM}}(v;w) = v\mathbf{1}\{v \geq c\} - (v-c)^+ = c\mathbf{1}\{v \geq c\}\,,$$

where the first equality follows from the formulas for the promised utility and allocation functions.

Constraint (2) is satisfied when $w \in [0,\underline{w})$ because $0 \leq u^{\text{DUM}}(v;w) \leq \bar{u}$ since $u^{\text{DUM}}(v;w) = w/\beta \leq \underline{w}/\beta = \mathbb{E}_v[(v-c)^+]/\beta$ and $\beta \geq 1/2$ from Assumption 4.1. When $w \in [\underline{w},\bar{w}]$ we have

$$u^{\text{DUM}}(v;w) \geq u^{\text{DUM}}(0;\underline{w}) = \frac{1}{\beta}\left(\underline{w} - \mathbb{E}_v[(v-c)^+]\right) = 0\,,$$

and

$$u^{\text{DUM}}(v;w) \leq u^{\text{DUM}}(\bar{v};\bar{w}) = \frac{1}{\beta}\left(\bar{w} + (\bar{v}-c)^+ - \mathbb{E}_v[(v-c)^+]\right) = \bar{u}\,,$$

from the definition of $\bar{w}$. When $w \in (\bar{w}, \bar{u}]$ we have that

$$u^{\text{DUM}}(v; w) \leq u^{\text{DUM}}(v; \bar{u}) = \frac{1}{\beta}(\bar{u} - \mathbb{E}_v[(v-c)^+]) = \bar{u}\,,$$

because $\mathbb{E}_v[(v-c)^+] = (1-\beta)\bar{u}$ and

$$u^{\text{DUM}}(v; w) \geq u^{\text{DUM}}(v; \bar{w}) = \frac{1}{\beta}(\bar{w} - \mathbb{E}_v[(v-c)^+]) = \frac{1}{\beta}(\bar{w} - \underline{w}) \geq 0\,,$$

because $\bar{w} \geq \underline{w}$.

B.3 Proof of Theorem 4.3

Let $\{w_t^{\text{DUM}}\}_{t=1}^{\infty}$ denote the stochastic process that governs the evolution of the state under the approximation mechanism $(q^{\text{DUM}}, z^{\text{DUM}}, u^{\text{DUM}})$. This process evolves according to $w_{t+1}^{\text{DUM}} = u^{\text{DUM}}(v_t, w_t^{\text{DUM}})$ with initial condition w^{DUM}. Let $\tau = \inf\{t \geq 1 : w_t^{\text{DUM}} \notin [\underline{w}, \bar{w}]\}$ be the first time that the promised utility falls outside the interval $[\underline{w}, \bar{w}]$. During $t = 1, \ldots, \tau - 1$ the dynamic mechanism allocates according to the first-best auction.

By construction, the mechanism guarantees that the scaled stochastic process $\{\beta^{t-1} w_t^{\text{DUM}}\}_{t=1}^{\infty}$ is a martingale whenever w_t^{DUM} in $[\underline{w}, \bar{w}]$, that is, the process satisfies $\mathbb{E}[\beta^t w_{t+1}^{\text{DUM}} | w_t^{\text{DUM}}] = \beta^{t-1} w_t^{\text{DUM}}$ because $\mathbb{E}[\Delta^{\text{DUM}}(v)] = 0$. The next result leverages this property to characterize some properties of the stopping time τ.

Lemma B.1. *Under the approximation mechanism $(q^{\text{DUM}}, z^{\text{DUM}}, u^{\text{DUM}})$,*

1. *The probability that the promised utility first falls below $\underline{w}$ before exceeding $\underline{w}$ is bounded by*

$$\mathbb{P}\{w_\tau^{\text{DUM}} \leq \underline{w}\} \leq \exp\left(-\frac{(1-\beta)(w^{\text{DUM}} - \underline{w})^2}{2\bar{v}^2}\right)\,.$$

2. *The stopping time τ satisfies*

$$\mathbb{E}[\beta^\tau] \leq \frac{w^{\text{DUM}}}{\bar{w}} + \mathbb{P}\{w_\tau^{\text{DUM}} \leq \underline{w}\}\,.$$

We are now in position to prove the main result. We have that

$$\Pi^{\text{DUM}} \leq \Pi^{\text{FB}}\,,$$

because the mechanism is consideration is primal feasible and thus is bounded by first-best. In the remainder of the proof we prove the last inequality, that is

$$\Pi^{\text{DUM}} \geq \Pi^{\text{FB}} - \tilde{O}\left(\frac{1}{(1-\beta)^{1/2}}\right)\,.$$

Step 1 The expected performance of approximation mechanism can be decomposed as follows:

$$\begin{aligned}
\Pi^{\text{DUM}} &= \sum_{t=1}^{\infty} \mathbb{E}[\beta^{t-1}(z^{\text{DUM}}(v_t; w_t) - cq^{\text{DUM}}(v_t; w_t))] \\
&= \frac{\mathbb{E}_v[(v-c)^+]}{1-\beta} - \sum_{t=1}^{\infty} \mathbb{E}\Big[\underbrace{\beta^{t-1}\left((v_t - c)^+ - z^{\text{DUM}}(v_t; w_t) + cq^{\text{DUM}}(v_t, w_t)\right)}_{R_t} \Big] \\
&= \Pi^{\text{FB}} - \sum_{t=1}^{\infty} \mathbb{E}[R_t] \,,
\end{aligned}$$

where the second equation follows because values are i.i.d. and using that $\sum_{t=0}^{\infty} \beta^{t-1} = 1/(1-\beta)$. The error terms R_t measure the difference in expected performance between the first-best allocation and the mechanism in consideration.

Step 2 Because the mechanism allocates according to first-best up to time τ we obtain

$$\begin{aligned}
\sum_{t=1}^{\infty} \mathbb{E}[R_t] &= \sum_{t=1}^{\infty} \mathbb{E}[R_t \mathbf{1}\{t < \tau\}] + \sum_{t=1}^{\infty} \mathbb{E}[R_t \mathbf{1}\{t \geq \tau\}] = \sum_{t=1}^{\infty} \mathbb{E}[R_t \mathbf{1}\{t \geq \tau\}] \\
&\leq \bar{v} \sum_{t=1}^{\infty} \mathbb{E}[\beta^{t-1} \mathbf{1}\{t \geq \tau\}] = \frac{\bar{v}}{\beta} \mathbb{E}\left[\sum_{t=\tau}^{\infty} \beta^t \right] = \frac{\bar{v}}{\beta(1-\beta)} \mathbb{E}\left[\beta^\tau\right] \\
&\leq \frac{2\bar{v}}{1-\beta} \left(\frac{w^{\text{DUM}}}{\bar{w}} + \exp\left(-\frac{(1-\beta)(w^{\text{DUM}} - \underline{w})^2}{2\bar{v}^2} \right) \right) \,,
\end{aligned}$$

where the first inequality follows because $|R_t| \leq \beta^{t-1}\bar{v}$ since the flow profit is at most $\bar{v}$, the second equality from Tonelli's Theorem, the last equation because the sum is a geometric series, and the last inequality from Lemma B.1 and $\beta \geq 1/2$ from Assumption 4.1.

Step 3 Setting $w^{\text{DUM}} = \underline{w} + \frac{\sqrt{2}\bar{v}}{(1-\beta)^{1/2}} \log^{1/2}\left(\frac{1}{1-\beta}\right)$, we obtain that the error term is bounded by

$$\begin{aligned}
\sum_{t=1}^{\infty} \mathbb{E}[R_t] &\leq \frac{2\bar{v}}{1-\beta} \left(\frac{\underline{w}}{\bar{w}} + \frac{\sqrt{2}\bar{v}}{\bar{w}} \frac{1}{(1-\beta)^{1/2}} \log^{1/2}\left(\frac{1}{1-\beta}\right) + 1 - \beta \right) \\
&= \frac{2\bar{v}\underline{w}}{(1-\beta)\bar{w}} + \frac{2\sqrt{2}\bar{v}^2}{(1-\beta)\bar{w}} \frac{1}{(1-\beta)^{1/2}} \log^{1/2}\left(\frac{1}{1-\beta}\right) + 2\bar{v} \\
&= \tilde{O}\left(\frac{1}{(1-\beta)^{1/2}} \right) \,,
\end{aligned}$$

and the last bound follows because $(1-\beta)\bar{w} \to \mathbb{E}_v\left[(v-c)^+\right]$ as $\beta \to 1$.

B.3.1 Proof of Lemma B.1

We prove each item at a time.

Item 1 Consider an alternate process $(\tilde{w}_t^{\text{DUM}})_t$ coupled with the same realization of values $(v_t)_t$ in which the promised utility evolves according to

$$\tilde{w}_{t+1}^{\text{DUM}} = \frac{1}{\beta}\left(\tilde{w}_t^{\text{DUM}} + \Delta^{\text{DUM}}(v_t)\right) ,$$

regardless of whether the state lies within the interval $[\underline{w}, \bar{w}]$ or not. Because both processes coincide up to time τ, we can upper bound the probability that the promised utility falls below $\underline{w}$ before it goes above $\bar{w}$ by the probability that the process $\beta^{t-1}\tilde{w}_t^{\text{DUM}}$ ever goes below $\underline{w}$, that is,

$$\begin{aligned}\mathbb{P}\{w_\tau^{\text{DUM}} \le \underline{w}\} &= \mathbb{P}\{\tilde{w}_\tau^{\text{DUM}} \le \underline{w}\} \le \mathbb{P}\{\beta^{\tau-1}\tilde{w}_\tau^{\text{DUM}} \le \underline{w}\} \le \mathbb{P}\{\cup_{t=1}^{\infty}\beta^{t-1}\tilde{w}_t^{\text{DUM}} \le \underline{w}\} \\ &= \mathbb{P}\{\min_{t=1,\dots,\infty} \beta^{t-1}\tilde{w}_t^{\text{DUM}} \le \underline{w}\} = \mathbb{P}\{\max_{t=1,\dots,\infty} w^{\text{DUM}} - \beta^{t-1}\tilde{w}_t^{\text{DUM}} \ge w^{\text{DUM}} - \underline{w}\} ,\end{aligned}$$

where the first inequality follows because $\beta \in (0,1)$ and $\tilde{w}_\tau^{\text{DUM}} \ge 0$. Because the scaled process

$$\beta^t \tilde{w}_{t+1}^{\text{DUM}} = \beta^{t-1}\tilde{w}_t^{\text{DUM}} + \beta^{t-1}\Delta^{\text{DUM}}(v_t) ,$$

is a martingale with differences bounded by $|\beta^t\tilde{w}_{t+1}^{\text{DUM}} - \beta^{t-1}\tilde{w}_t^{\text{DUM}}| = \beta^{t-1}|\Delta^{\text{DUM}}(v_t)| \le \beta^{t-1}\bar{v}$, we obtain via Azuma's inequality for maxima (see, e.g., McDiarmid 1998, Section 3.5)

$$\mathbb{P}\{w_\tau^{\text{DUM}} \le \underline{w}\} \le \exp\left(-\frac{(w^{\text{DUM}} - \underline{w})^2}{2\bar{v}^2\sum_{t=0}^{\infty}\beta^{2t}}\right) \le \exp\left(-\frac{(1-\beta)(w^{\text{DUM}} - \underline{w})^2}{2\bar{v}^2}\right) ,$$

where the last inequality follows because $\sum_{t=0}^{\infty}\beta^{2t} \le \sum_{t=0}^{\infty}\beta^t = 1/(1-\beta)$ because $\beta \in (0,1)$. The result follows.

Item 2 Because the stopping time $t \wedge \tau$ is finite and $\beta^{t\wedge\tau-1}w_{t\wedge\tau}^{\text{DUM}}$ is a martingale we obtain by the Optional Stopping Theorem that

$$\begin{aligned}\beta w^{\text{DUM}} &= \mathbb{E}\left[\beta^{t\wedge\tau} w_{t\wedge\tau}^{\text{DUM}}\right] \\ &= \mathbb{E}\left[\beta^\tau w_\tau^{\text{DUM}}\mathbf{1}\{w_\tau^{\text{DUM}} \ge \bar{w}, \tau \le t\}\right] + \mathbb{E}\left[\beta^\tau w_\tau^{\text{DUM}}\mathbf{1}\{w_\tau^{\text{DUM}} \le \underline{w}, \tau \le t\}\right] + \mathbb{E}\left[\beta^t w_t^{\text{DUM}}\mathbf{1}\{\tau > t\}\right] \\ &\ge \bar{w}\mathbb{E}\left[\beta^\tau \mathbf{1}\{w_\tau^{\text{DUM}} \ge \bar{w}, \tau \le t\}\right] ,\end{aligned}$$

where the inequality follows discarding the second and third terms because $w_t^{\text{DUM}} \ge 0$ for all t. This implies that $\mathbb{E}\left[\beta^\tau\mathbf{1}\{w_\tau^{\text{DUM}} \ge \bar{w}, \tau \le t\}\right] \le \beta w^{\text{DUM}}/\bar{w}$. We can write the expectation $\mathbb{E}\left[\beta^{t\wedge\tau}\right]$ as follows

$$\begin{aligned}\mathbb{E}\left[\beta^{t\wedge\tau}\right] &= \mathbb{E}\left[\beta^\tau\mathbf{1}\{w_\tau^{\text{DUM}} \ge \bar{w}, \tau \le t\}\right] + \mathbb{E}\left[\beta^\tau\mathbf{1}\{w_\tau^{\text{DUM}} \le \underline{w}, \tau \le t\}\right] + \mathbb{E}\left[\beta^t\mathbf{1}\{\tau > t\}\right] \\ &\le \frac{\beta w^{\text{DUM}}}{\bar{w}} + \mathbb{P}\{w_\tau^{\text{DUM}} \le \underline{w}\} + \beta^t ,\end{aligned}$$

where the inequality follows from the previous bound, using that $\beta^\tau \le 1$ because $\beta \in (0,1)$, and discarding the events $\{\tau \le t\}$ and $\{\tau > t\}$ in the second and third terms, respectively. Because $\beta^{t\wedge\tau} \le 1$ since $\beta \in (0,1)$ we obtain from Dominated Convergence Theorem that

$$\mathbb{E}\left[\beta^\tau\right] = \lim_{t\to\infty}\mathbb{E}\left[\beta^{t\wedge\tau}\right] \le \frac{\beta w^{\text{DUM}}}{\bar{w}} + \mathbb{P}\{w_\tau^{\text{DUM}} \le \underline{w}\} ,$$

and the result follows because $\beta \le 1$.

B.4 Equivalence of Martingale Constraints

In this section we show that martingale constraints in terms of continuation and flow utilities are equivalent.

Lemma B.2. *Let $(q(v;w), z(v;w), u(v;w))$ be a feasible dynamic mechanism. Then $\mathbb{E}_v\left[\ell\big(u(v;w)\big)\right] = \ell(w)$ for all states $w \in \mathcal{U}$ if and only if $\mathbb{E}_v\left[u(v;w)\right] = w$ for all states $w \in \mathcal{U}$.*

Proof. We prove each direction at a time.

Continuation implies flow By the promise keeping constraint we have that

$$\ell(w) = w - \beta\mathbb{E}_v\left[u(v;w)\right] = (1-\beta)w\,, \tag{B-15}$$

where the last equation follows because the continuation utility is a martingale, i.e, $\mathbb{E}_v\left[u(v;w)\right] = w$. Thus

$$\mathbb{E}_v\left[\ell\big(u(v;w)\big)\right] = (1-\beta)\mathbb{E}_v\left[u(v;w)\right] = (1-\beta)w = \ell(w)\,,$$

where the first equality follows from using (B-15) pointwise for $w = u(v;w)$, the second because the continuation utility is a martingale, and the last equality follows from (B-15).

Flow implies continuation Let $(w_t)_t$ be the stochastic process given by $w_1 = w$ and $w_{t+1} = u(v_t;w_t)$. By the promise keeping constraint we have that the continuation utility can be written as

$$w = \sum_{t=1}^{\infty} \beta^{t-1}\mathbb{E}_{w_t}\left[\ell\big(w_t\big)\right] = \frac{\ell(w)}{1-\beta}\,, \tag{B-16}$$

because $\mathbb{E}_{w_t}\left[\ell\big(w_t\big)\right] = \ell(w)$ since the flow utility is a martingale. By the promise keeping constraint we have that

$$\mathbb{E}_v\left[u(v;w)\right] = \frac{1}{\beta}\left(w - \ell(w)\right) = w\,,$$

because $\ell(w) = (1-\beta)w$ from (B-16). □

B.5 Proof of Lemma A.2

We provide the proof of each item separately.

Item 1 We have that the c.d.f. of the stopping time can be upper bounded by

$$\mathbb{P}(\tau < t) = \mathbb{P}(w_i^{\text{MUM}} \le \underline{w} \text{ for some } i = 1,\dots,t-1) = \mathbb{P}\left(\max_{i=1,\dots,t-1} w^{\text{MUM}} - w_{i\wedge\tau}^{\text{MUM}} \ge w^{\text{MUM}} - \underline{w}\right).$$

When $t < \tau$ we have the promised utility evolves according to the martingale

$$w_{t+1}^{\text{MUM}} = \frac{1}{\beta}\left(w_t^{\text{MUM}} + \Delta_t^{\text{MUM}}\right)\,,$$

where $\Delta_t^{\text{MUM}} = \big(\min(v_t, s^{\text{MUM}}(w_t^{\text{MUM}})) - c\big)^+ - \mathbb{E}_v[v-c]^+$. The martingale differences are bounded by

$$\left|w_{t+1}^{\text{MUM}} - w_t^{\text{MUM}}\right| = \left|\frac{1-\beta}{\beta} w_t^{\text{MUM}} + \frac{1}{\beta}\Delta_t^{\text{MUM}}\right| \le \frac{1}{\beta}\left((1-\beta)\bar{u} + \bar{v}\right) \le \frac{2\bar{v}}{\beta}\,,$$

where the first inequality follows because $w_t^{\text{MUM}} \le \bar{u}$ and $|\Delta_t^{\text{MUM}}| \le \bar{v}$, and the last because $\bar{u} = \mathbb{E}\left[(v-c)^+\right]/(1-\beta) \le \bar{v}/(1-\beta)$. Because the stopped martingale $w_{t\wedge\tau}^{\text{MUM}}$ is a martingale, we obtain from Azuma's inequality for maxima (see, e.g., McDiarmid 1998, Section 3.5) that

$$\mathbb{P}(\tau < t) \le \exp\left(-\frac{\beta^2(w^{\text{MUM}} - \underline{w})^2}{4\bar{v}^2(t-1)}\right).$$

Item 2. The expectation in the statement of the lemma is given as follows.

$$\mathbb{E}[\beta^\tau] = \sum_{t=1}^{\infty}\beta^t\mathbb{P}(\tau = t) = \sum_{t=1}^{\infty}\beta^t\mathbb{P}(\tau < t+1) - \sum_{t=1}^{\infty}\beta^t\mathbb{P}(\tau < t)\,.$$

Note that the first term can equivalently be expressed as follows

$$\sum_{t=1}^{\infty}\beta^t\mathbb{P}(\tau < t+1) = \sum_{t=1}^{\infty}\beta^{t-1}\mathbb{P}(\tau < t) - \mathbb{P}(\tau < 1) = \frac{1}{\beta}\sum_{t=1}^{\infty}\beta^t\mathbb{P}(\tau < t)$$

where we used that $\mathbb{P}(\tau < 1) = 0$. Hence the expectation is bounded as follows:

$$\mathbb{E}[\beta^\tau] = \frac{1-\beta}{\beta}\sum_{t=1}^{\infty}\beta^t\mathbb{P}(\tau < t) \le \frac{1-\beta}{\beta}\sum_{t=1}^{\infty}\beta^t\exp\left(-\frac{\beta^2(w^{\text{MUM}} - \underline{w})^2}{4\bar{v}^2(t-1)}\right) = \frac{1-\beta}{\beta}\sum_{t=0}^{\infty}f(t)\,,$$

where we denote $f(t) = \beta\exp\left(-at - b/t\right)$ with $a = \ln\left(\beta^{-1}\right)$ and $b = \beta^2(w^{\text{MUM}} - \underline{w})^2/(4\bar{v}^2)$. Because the function $f(t)$ is unimodal (initially non-decreasing and then non-increasing), we can bound the summation by

$$\begin{aligned}\sum_{t=0}^{\infty}f(t) &\le \max_{t\ge 0}f(t) + \int_0^{\infty}f(t)\ \mathrm{d}t = \beta\exp(-2(ab)^{1/2}) + 2\beta\left(\frac{b}{a}\right)^{1/2}K_1\left(2(ab)^{1/2}\right)\\ &\le \left(1 + 4K_1(1)eb\right)\beta\exp\left(-2(ab)^{1/2}\right) \le \left(1+8b\right)\beta\exp\left(-2\left((1-\beta)b\right)^{1/2}\right).\end{aligned}$$

where the first equality follows because $\max_{t\ge 0}f(t) = \beta\exp(-2(ab)^{1/2})$ and denoting by $K_1(z) = z/4\int_0^{\infty}\exp(-z^2/4x - 1/x)\ \mathrm{d}x$ the modified Bessel function of the second kind of order 1; the second inequality because $K_1(x)/K_1(y) \ge e^{y-x}x/y$ for $y \ge x$ (Laforgia, 1991) and setting $y = 2(ab)^{1/2}$ and $x = 1$; and the last because $K_1(1)e \le 2$ and $a \ge 1-\beta$. Here we used that $y \ge x$ because $b \ge 1/(4(1-\beta))$. The result follows.

C Mechanism Design with Dynamic Individual Rationality

An optimal dynamic mechanism can be recursively characterized by the Principle of Optimality. Let $\Pi^{\text{DIR}}(w)$ be the optimal expected profit-to-go of the seller when the promised utility is $w \in \mathcal{U}$. We have that the optimal value function satisfies the following Bellman equation (see, e.g., Bertsekas

(2012, Chapter 1)):

$$\Pi^{\text{DIR}}(w) = \max_{(q,z,u)\in\mathcal{M}'} \mathbb{E}_v[z(v) - cq(v) + \beta\Pi^{\text{DIR}}(u(v))] \qquad \text{(C-17)}$$
$$\text{s.t } (1), (2), (\text{PK}), (\text{DIC}), (\text{DIR})\,.$$

Given a value function $\Pi^{\text{DIR}}(\cdot)$ solving the previous Bellman Equation, the optimal profit of the seller is given by $\Pi^{\text{DIR}} = \max_{w\in\mathcal{U}} \Pi^{\text{DIR}}(w)$, and the initial state w^{DIR} is a maximizer of the previous problem.

In this setting an optimal dynamic mechanism can be characterized in closed-form. The following result shows an optimal dynamic mechanism that allocates the item according to Myerson's optimal auction in the first period and allocates the item efficiently onwards. The seller incorporates a large participation fee in the first auction that allows him to extract the buyer's entire surplus from the second auction onwards. The next result characterizes an optimal dynamic mechanisms in terms of histories, instead of recursively using the promised utility framework.

Proposition C.1. *Suppose the hazard rate $f(v)/(1-F(v))$ is non-decreasing in v. An optimal dynamic mechanism is given by*

- *when $t = 1$ the allocation is $q(v) = \mathbf{1}\{\phi(v) \geq c\}$ and the payment is $z(v) = \phi^{-1}(c)\mathbf{1}\{\phi(v) \geq c\} + \frac{\beta}{1-\beta}\mathbb{E}_v\left[(v-c)^+\right]$, and*
- *when $t > 1$ the allocation $q(v) = \mathbf{1}\{v \geq c\}$ and the payment is $z(v) = c\mathbf{1}\{v \geq c\}$.*

Here the seller exploits that at the point of contracting the only private information is the buyer's value for the first auction and that there is no information asymmetry for future time periods. Thus, the seller only pays the buyer an information rent for the first time period and achieves close to full surplus extraction. This result is well-known in the mechanism design literature and was originally established by Besanko (1985), who showed that with risk-neutrality and independent private information, all inefficiency arises in the first period only. In this case the expected discounted profit of the seller is

$$\Pi^{\text{DIR}} = \mathbb{E}_v\left[(\phi(v)-c)^+\right] + \frac{\beta}{1-\beta}\mathbb{E}_v\left[(v-c)^+\right]\,.$$

This implies that an optimal mechanism achieves first-best as the discount rate converges to one:

$$\lim_{\beta\to 1} \bar{\Pi}^{\text{DIR}} = \lim_{\beta\to 1}(1-\beta)\Pi^{\text{DIR}} = \mathbb{E}_v\left[(v-c)^+\right] = \bar{\Pi}^{\text{FB}}\,,$$

where the rate of convergence to first-best is $O\left(1-\beta\right)$, i.e. $\bar{\Pi}^{\text{FB}} - \bar{\Pi}^{\text{DIR}} \leq O(1-\beta)$.

We prove the result by characterizing in closed-form the optimal value function in (C-17). Using Lemma 2.1 we can eliminate payments from the inner optimization problem in the Bellman Equation, and obtain a simpler problem constrained only by (1), (2) and (DIR). The resulting problem is easy to solve because, after eliminating payments, (DIR) only restricts the allocation. Thus the inner optimization problem is separable in the allocation and the promise functions, which allows us to explicitly characterize an optimal mechanism by optimizing the objective pointwise over the realization of values. We shall see, however, that when we introduce additional constraints the problem is not longer separable and, in most cases, an optimal mechanism may not be obtained in closed-form.

In this case an alternative proof can be provided using a standard relaxation approach from the dynamic mechanism design literature (Ëso and Szentes, 2007; Kakade et al., 2013; Pavan et al., 2014). This approach involves relaxing most incentive compatibility constraints, solving the relaxed problem, and then showing that the candidate mechanism satisfies all incentive compatibility constraints. In the relaxed problem all incentive compatibility constraints, other that the one of the first auction, are relaxed. This corresponds to an environment where the value of the buyer for the first auction is private and all future values are observed by the seller. Thus this problem can be reduced to a single-item mechanism design problem (with proper continuation values) that can be solved in a straightforward manner using standard techniques. In particular, in the case of (DIR) the mechanism can align all incentives for future time periods via a large participation fee, and it is possible to un-relax an optimal mechanism from the relaxed problem. While this approach can handle non-independent valuations, to the best of our knowledge, this approach cannot be extended to accommodate constraints such as the ones we later consider in this paper. Our proof of Proposition C.1, while somewhat longer, better illustrates the use of the promised utility framework, which is the cornerstone of our analysis.

C.1 Proof of Proposition C.1

We prove the result in four steps. First, we use Lemma 2.1 to eliminate payments from the optimization problem and rewrite the seller's optimization problem in terms of social welfare. Second, we characterize the optimal value function. Third, we determine an optimal initial state and optimal seller's profit by optimizing over the initial state. Fourth, we determine an optimal mechanism.

Step 1 (Reformulation) Using Lemma 2.1 we can eliminate payments from the optimization problem. The DIR constraint is now given by

$$w + \int_0^v q(x)\mathrm{d}x - \int_0^\infty \bar{F}(x)q(x)\mathrm{d}x \geq 0\,,$$

for all v. Because the allocation $q(x)$ is non-decreasing, it suffices to impose the DIR constraint at the lowest type:

$$\int_0^\infty \bar{F}(x)q(x)\mathrm{d}x \leq w\,. \tag{C-18}$$

Denoting by $S^{\mathrm{DIR}}(w) \triangleq w + \Pi^{\mathrm{DIR}}(w)$ the social welfare generated by an optimal mechanism, we obtain that the mechanism design problem (C-17) can be written as

$$S^{\mathrm{DIR}}(w) = \max_{q,u} \mathbb{E}_v[(v-c)q(v) + \beta S^{\mathrm{DIR}}(u(v))] \tag{C-19}$$
$$\text{s.t } (1), (2), \text{(C-18)}, q(x) \text{ non-decreasing in } v\,.$$

Step 2 (Optimal value function) Let $\gamma \geq 0$ be the Lagrange multiplier of the individual rationality constraint (C-18), let $h(\gamma) = \mathbb{E}_v\left[\left(v - \gamma\frac{\bar{F}(v)}{f(v)} - c\right)^+\right]$, and let $g(w) = \min_{\gamma\geq 0}\{h(\gamma) + \gamma w\}$. We shall prove that $S(w) = g(w) + \beta\bar{u}$ solves the Bellman equation $S(w) = T(S)(w)$ where $T(\cdot)$ is the Bellman operator in (C-19). We first prove the following properties.

Lemma C.2. *The following hold:*

(i) $h(\gamma)$ *is differentiable and strictly convex. Moreover, the derivative is* $h'(\gamma) = -\mathbb{E}_v\left[(v - v(\gamma))^+\right]$ *where* $v(\gamma)$ *is the unique root of* $v - \gamma\bar{F}(v)/f(v) - c = 0$ *over* $v \in [0, \bar{v}]$.

(ii) $\gamma^*(w) \in \arg\max_{\gamma \geq 0} \left\{ h(\gamma) + \gamma w \right\}$ *is unique, bounded and continuous for* $w > 0$.

(iii) $g(w)$ *is concave and differentiable for* $w > 0$.

(iv) $g(w)$ *is non-decreasing.*

(v) $g(\bar{u}) = (1-\beta)\bar{u}$.

We are in position to prove the result. Because the objective in (C-19) is separable we obtain that

$$T(S)(w) = \underbrace{\begin{array}{l} \max_q \left\{ \mathbb{E}_v[(v-c)q(v)] \right\} \\ \text{s.t } (1), (\text{C-18}), q(x) \text{ non-decreasing in } v \end{array}}_{(I)} \quad + \quad \beta \underbrace{\begin{array}{l} \max_u \left\{ S(u(v)) \right\} \\ \text{s.t } (2) \end{array}}_{(II)}$$

We solve the first problem by introducing a Lagrange multiplier $\gamma \geq 0$ for constraint (C-18) and considering its dual problem. Note that in (I) the objective and constraint (C-18) are linear, and the feasible set $\mathcal{Q} = \{q : 0 \leq q(v) \leq 1, q(x) \text{ non-decreasing in } v\}$ is convex with a feasible interior point. Hence, by the Strong Duality Theorem (p.224 in Luenberger (1969)) we obtain

$$\begin{aligned} (I) &= \min_{\gamma \geq 0} \max_{q \in \mathcal{Q}} \left\{ \mathbb{E}_v \left[\left(v - \gamma \frac{\bar{F}(v)}{f(v)} - c \right) q(v) \right] \right\} + \gamma w \\ &= \min_{\gamma \geq 0} h(\gamma) + \gamma w = g(w) , \end{aligned}$$

where the second equality follows from optimizing pointwise over v to obtain that $q(v) = \mathbf{1}\{v - \gamma\bar{F}(v)/f(v) - c \geq 0\}$ is optimal for all $\gamma \geq 0$ (here we used that $q(v)$ is non-decreasing because the hazard rate $f(v)/\bar{F}(v)$ is non-decreasing), and the last from our definition of $g(w)$. For the second problem we have that

$$(II) = S(\bar{u}) = \bar{u} ,$$

because the optimal $u(v)$ under the value function $S(w)$ is given by $u(v) = \bar{u}$ since $S(w)$ is non-decreasing in w (because $g(w)$ is non-decreasing) and the last equation because $S(\bar{u}) = \bar{u}$ since $g(\bar{u}) = (1-\beta)\bar{u}$. We thus obtain that

$$T(S)(w) = g(w) + \beta\bar{u} = S(w) ,$$

and we conclude that $S(w)$ satisfies the Bellman equation.

Step 3 (Optimal initial state) An optimal initial state is obtained by solving for the optimal seller's profit:

$$\Pi^{\text{DIR}} = \max_{w \in \mathcal{U}} \Pi^{\text{DIR}}(w) = \max_{w \in \mathcal{U}} \left\{ S^{\text{DIR}}(w) - w \right\} = \max_{0 \leq w \leq \bar{u}} \left\{ g(w) + \beta\bar{u} - w \right\} .$$

We claim that $w^{\text{DIR}} = -h'(1) = \mathbb{E}_v \left[(v - \phi^{-1}(c))^+ \right]$ is an optimal solution (the last expression follows from property (i) above). Because the proposed solution is interior, and the objective is convex and differentiable at the proposed solution; it suffices to check that the proposed solution satisfies the first-order condition $0 = g'(w^{\text{DIR}}) - 1$. By the Envelope Theorem, we have that $g'(w) = \gamma^*(w)$ where $\gamma^*(w) \in \arg\max_{\gamma \geq 0} \left\{ h(\gamma) + \gamma w \right\}$. This implies that $\gamma^*(w^{\text{DIR}}) = 1$. In turn, because the solution of the inner optimization problem is interior, and the objective is convex and differentiable;

the first-order condition of the inner optimization problem gives that $h'(\gamma^*(w^{\text{DIR}}))+w^{\text{DIR}}=0$. Using that $\gamma^*(w^{\text{DIR}})=1$ we conclude that $w^{\text{DIR}}=-h'(1)$ as claimed.

Additionally, we obtain that the optimal seller's profit is given by

$$\Pi^{\text{DIR}}=g(w^{\text{DIR}})+\beta\bar{u}-w^{\text{DIR}}=h(1)+\beta\bar{u}=\mathbb{E}_v\left[(\phi(v)-c)^+\right]+\frac{\beta}{1-\beta}\mathbb{E}_v\left[(v-c)^+\right],$$

where the second equality follows because $\gamma^*(w^{\text{DIR}})=1$ is optimal for the inner problem.

Step 4 (Optimal mechanism) In the first auction we have that $w_1=w^{\text{DIR}}$ and $\gamma^*(w_1)=1$. From step 2 we have that the allocation is $q(v)=\mathbf{1}\{\phi(v)\ge c\}$ and the promise function is $u(v)=\bar{u}$. From (3) we obtain that the payment is given by

$$\begin{aligned}z(v)&=\beta u(v)-w^{\text{DIR}}+vq(v)-\int_0^v q(x)\mathrm{d}x+\int_0^\infty \bar{F}(x)q(x)\mathrm{d}x\\&=\frac{\beta}{1-\beta}\mathbb{E}_v\left[(v-c)^+\right]+\phi^{-1}(c)\mathbf{1}\{v\ge\phi^{-1}(c)\},\end{aligned}$$

where the last equation follows because $\int_0^\infty \bar{F}(x)q(x)\mathrm{d}x=\mathbb{E}_v\left[(v-\phi^{-1}(c))^+\right]=w^{\text{DIR}}$ and $\int_0^v q(x)\mathrm{d}x=(v-\phi^{-1}(c))\mathbf{1}\{v\ge\phi^{-1}(c)\}$.

In the second auction we have that $w_2=\bar{u}$ and $\gamma^*(w_2)=0$. From step 2 we have that the allocation is $q(v)=\mathbf{1}\{v\ge c\}$ and the promise function is $u(v)=\bar{u}$. From (3) we obtain that the payment is given by

$$\begin{aligned}z(v)&=\beta u(v)-\bar{u}+vq(v)-\int_0^v q(x)\mathrm{d}x+\int_0^\infty \bar{F}(x)q(x)\mathrm{d}x\\&=c\mathbf{1}\{v\ge c\},\end{aligned}$$

where the last equation follows because $\int_0^\infty \bar{F}(x)q(x)\mathrm{d}x=\mathbb{E}_v\left[(v-c)^+\right]=(1-\beta)\bar{u}$ and $\int_0^v q(x)\mathrm{d}x=(v-c)\mathbf{1}\{v\ge c\}$. Because the promise function is $u(v)=\bar{u}$, then $\bar{u}$ is an absorbing state and the same mechanism is used onwards.

C.2 Proof of Lemma C.2

We prove each item at a time.

Item i Differentiability of $h(\gamma)=\mathbb{E}_v[h(\gamma,v)]$ where $h(\gamma,v)=\left(v-\gamma\bar{F}(v)/f(v)-c\right)^+$ follows from Liebniz's rule because the function $h(\gamma,v)$ is differentiable almost everywhere with respect to γ and its derivative $h_\gamma(\gamma,v)=-\bar{F}(v)/f(v)\mathbf{1}\{v-\gamma\bar{F}(v)/f(v)-c\ge 0\}$ is bounded by an integrable function as follows: $|h_\gamma(\gamma,v)|\le\bar{F}(v)/f(v)$. The derivative is given by $h'(\gamma)=-\int_0^{\bar{v}}\bar{F}(v)\mathbf{1}\{v-\gamma\bar{F}(v)/f(v)-c\ge 0\}\mathrm{d}v$. Because the hazard rate is increasing and $c\in[0,\bar{v})$ we have that $h'(\gamma)=-\int_{v(\gamma)}^{\bar{v}}\bar{F}(v)\mathrm{d}v=-\mathbb{E}_v\left[(v-v(\gamma))^+\right]$ where $v(\gamma)$ is the unique root of $v-\gamma\bar{F}(v)/f(v)-c=0$ over $v\in[0,\bar{v}]$. Because the density $f(\cdot)$ is strictly positive on its domain we conclude that $h(\gamma)$ is strictly convex.

Item ii When $w>0$ we have that the optimal solution is bounded because all solutions such that $\gamma\ge h(0)/w$ are dominated by $\gamma=0$ since the objective is lower bounded by $g(w,\gamma)=h(\gamma)+\gamma w\ge h(0)=g(w,0)$ because $h(\gamma)\ge 0$. The Maximum Theorem implies that $\gamma^*(w)$ is continuous because the objective is jointly continuous, strictly convex in γ and the feasible set is compact.

Item iii Concavity of $g(w)$ follows because $g(w)$ is the minimum of affine functions in w. Differentiability of $g(w)$ follows from Theorem 3 of Milgrom and Segal (2002) because the objective $g(w,\gamma)$ is equidifferentiable in w since $g_w(w,\gamma) = \gamma$, and $\gamma^*(w)$ is bounded and continuous for $w > 0$.

Item iv Non-decreasingness of $g(w)$ follows trivially because $\gamma \geq 0$.

Item v This follows because $\gamma = 0$ is an optimal solution when $w = \bar{u}$. The optimality of $\gamma = 0$ follows because $h(\gamma)$ is convex and the derivative of the objective evaluated at $\gamma = 0$ is non-negative. That is,

$$g_\gamma(\bar{u}, 0) = \bar{u} + h'(0) = \bar{u} - \mathbb{E}_v\left[(v-c)^+\right] = \beta\bar{u} \geq 0\,,$$

where the last equality follows because $\bar{u} = \mathbb{E}_v\left[(v-c)^+\right]/(1-\beta)$.

D Appendix to Section 1.2

In this section we give a more in depth exposition of the models and the proposed mechanisms under the different requirements for the applications discussed in Section 1.2.

D.1 Supply Chain Contracting with Private Price Information

Consider a supply chain contract in which a manufacturer (the principal) repeatedly sells a perishable good to a retailer (the agent) facing a newsvendor problem. The manufacturer has a publicly observed marginal cost $c > 0$.[14] The retailer faces uncertain demand in period t, denoted by d_t, and needs to place an order quantity before demand is realized. The retail price in each period, denoted by v_t, is privately observed by the retailer before making the ordering decision and satisfies $v_t \geq c$. The realized demands $(d_t)_{t\geq 1}$ and retail prices $(v_t)_{t\geq 1}$ are assumed to be independent and identically distributed across time periods with cumulative distribution functions $G(d)$ and $F(v)$. Because the good is perishable, inventory is not carried over and, in the event of a stock out, unmet demand is lost.

The manufacturer offers a non-linear pricing contract to the retailer that specifies a wholesale price as a function of the order quantity. Since the manufacturer has commitment power in our model, we can restrict attention to direct mechanism in which the retailer reports the retail price truthfully to the manufacturer. In this setting a stage mechanism is a pair of functions $(q, z) \in \mathcal{M}$ where $q : \mathbb{R}_+ \to \mathbb{R}_+$ is a supply function and $z : \mathbb{R}_+ \to \mathbb{R}_+$ is a payment function. That is, when the retailer reports a price v, $q(v)$ determines the number of units shipped to the retailer and $z(v)$ determines the payment to be charged. Given a direct dynamic mechanism $\pi \in \mathbb{M}$ the discounted profit of the manufacturer is as before and the discounted utility of the retailer is given by $U^\pi = \mathbb{E}\left[\sum_{t=1}^{\infty} \beta^{t-1}\big(v_t \min(d_t, q_t^\pi(v_t)) - z_t^\pi(v_t)\big)\right]$. We study the manufacturer's problem of designing a contract that maximizes his expected total discounted profits.

This model is a natural extension of Lariviere and Porteus (2001) to a repeated setting with private retail prices. In Lariviere and Porteus (2001) the retail price is publicly observed and distortion is introduced because the manufacturer is restricted to price-only contracts in which the retailer can buy as much as she wants at the posted price. Lobel and Xiao (2016) study the design of supply contracts under DIR in a dynamic setting when the retailer has private demand information and leftover inventory is carried over to the next time period. Kakade et al. (2013)

[14] While for simplicity we assume that the manufacturer's cost is linear, the same analysis can be easily extended to convex costs.

study the design of optimal mechanisms under DIR in a similar model with perishable goods in which the price and demand process are potentially non-stationary.

First-best mechanism In this case the first-best outcome (FB) corresponds to the case of an integrated firm. The average optimal expected profit of the integrated firm is given by the newsvendor problem

$$\bar{\Pi}^{\mathrm{FB}} = \mathbb{E}_v \left[\max_{q \geq 0} \{vD(q) - cq\} \right] ,$$

where we denote by $D(q) = \mathbb{E}_d [\min(d, q)]$ the expected sales when q units are offered. The function $D(q)$ is non-decreasing and concave, with derivative $D'(q) = 1 - G(q)$. The integrated firm optimally determines the quantity q as a function of the retail price v by equating marginal revenue to marginal cost. Because the retail price is never below the marginal manufacturing cost, the optimal ordering quantity is given by the critical fractile formula $q^{\mathrm{FB}}(v) = G^{-1}\big((v - c)/v\big)$. An omniscient manufacturer who observes the realization of retail prices could achieve the profit of the integrated firm by charging the retailer's revenue $z^{\mathrm{FB}}(v) = vD\left(q^{\mathrm{FB}}(v)\right)$.

Optimal static mechanism The optimal static incentive compatible mechanism (S) corresponds to the case of a decentralized supply chain in which the manufacturer needs to provide incentives for the retailer to report the realized price truthfully. Let $F(v)$ be the cumulative distribution function and $f(v)$ be the density of the retail price. Assume that the virtual value of the retail price $\phi(v) = v - (1 - F(v))/f(v)$ increasing in v. A standard application of the envelope formula yields that the average optimal profit is given by

$$\bar{\Pi}^{\mathrm{S}} = \mathbb{E}_v \left[\max_{q \geq 0} \{\phi(v)D(q) - cq\} \right] ,$$

that is, the manufacturer faces another newsvendor problem in which the marginal revenue is replaced by $\phi(v)$ to reflect the informational rent that needs to be paid to the retailer. When $\phi(v) < c$ the seller's objective is decreasing in the order quantity and the optimal solution is $q^{\mathrm{S}}(v) = 0$. Conversely, when $\phi(v) \geq c$ the optimal ordering quantity is again given by the critical fractile formula $q^{\mathrm{S}}(v) = G^{-1}\big((\phi(v) - c)/\phi(v)\big)$. The envelope formula yields that the optimal payment is given by $z^{\mathrm{S}}(v) = vD\left(q^{\mathrm{S}}(v)\right) - \int_0^v D\left(q^{\mathrm{S}}(x)\right) \mathrm{d}x$. Because the allocation is non-decreasing (from the monotonicity of the virtual value), the optimal mechanism can be implemented via the non-linear pricing contract $z^{\mathrm{S}}(q) = \inf_v \{z^{\mathrm{S}}(v) : q^{\mathrm{S}}(v) \geq q\}$.

Dynamic Individual Rationality The optimal dynamic mechanism under DIR supplies in the first time period the optimal order quantity of the decentralized supply chain $q^{\mathrm{S}}(v)$ and from then on supplies the optimal order quantity of the integrated supply chain $q^{\mathrm{FB}}(v)$. In the first time period the retailer is charged the optimal payment of the decentralized supply chain $z^{\mathrm{S}}(v)$ together with an additional upfront fee of $\beta/(1-\beta)\bar{\Pi}^{\mathrm{FB}}$. In return the retailer obtains all future goods at marginal cost c. By charging a large enough payment in the first time period and exploiting that both parties are informationally symmetric for future periods, the principal can achieve almost full surplus extraction. The average optimal profit of this mechanism is given by

$$\bar{\Pi}^{\mathrm{DIR}} = (1 - \beta)\bar{\Pi}^{S} + \beta\bar{\Pi}^{\mathrm{FB}} ,$$

and the convergence rate to first-best is $(1 - \beta)$.

Periodic Individual Rationality The Deferred Utility Mechanism has three phases as before. In the savings phase the manufacturer implements the optimal mechanism of the integrated supply chain: the manufacturer supplies $q^{\text{FB}}(v)$ and extracts the maximum possible retailer surplus by setting the payment equal to the retailer's revenue $z^{\text{FB}}(v)$. Truthful reporting during the savings phase is achieved exclusively via future promises. In the income phase the retailer is supplied the optimal order quantity of the integrated supply chain $q^{\text{FB}}(v)$ at marginal cost c. The mechanism initially starts with a promised utility in the low end of the saving phase. Because the PIR constraint is binding in the savings phase, the envelope formula (3) implies that $\mathbb{E}_v[u(v;w)] = w/\beta$ and the state drifts upwards. Thus after an initial period in the savings phase, the mechanism will likely switch to the income phase. Let $\bar{\Pi}^{\text{FB}}(v) = vD\big(q^{\text{FB}}(v)\big) - cq^{\text{FB}}(v)$ be the optimal profit of the integrated supply chain when the retail price is v. In the initial savings phase the retailer derives no rent and the manufacturer keeps the entire surplus of efficient trade $\bar{\Pi}^{\text{FB}}(v)$, while in the later income phase the whole surplus of the supply chain $\bar{\Pi}^{\text{FB}}(v)$ is collected by the retailer. Thus the retailer's utilities are backloaded. Theorem 4.3 implies that the rate of convergence to first-best of this mechanism is $(1-\beta)^{1/2}$.

Martingale Utilities The Martingale Utility Mechanism has the following structure in this setting. The retailer is supplied the optimal order quantity of the integrated supply chain $q^{\text{FB}}(v)$ and the payment is determined using a dynamic threshold s. Let $\bar{\Pi}^{\text{FB}}(v) = vD\big(q^{\text{FB}}(v)\big) - cq^{\text{FB}}(v)$ be the optimal profit of the integrated supply chain when the retail price is v. The manufacturer captures the entire surplus from efficient trade $\bar{\Pi}^{\text{FB}}(v)$ when the price is below this threshold, and the additional surplus $\bar{\Pi}^{\text{FB}}(v) - \bar{\Pi}^{\text{FB}}(s)$ generated when the price is above this threshold is provided to the retailer. That is, the retailer pays her revenue $z(v) = vD\big(q^{\text{FB}}(v)\big)$ when $v \le s$ and pays the manufacturer's marginal cost $z(v) = sD\big(q^{\text{FB}}(s)\big) + c\big(q^{\text{FB}}(v) - q^{\text{FB}}(s)\big)$ when $v > s$. Thus the retailer derives a positive flow utility only if the price is above the threshold s. The threshold is dynamic and set so that the average continuation utility is equal to the agent's expected flow utility of the current period, or equivalently $(1-\beta)w = \mathbb{E}_v\left[(\bar{\Pi}^{\text{FB}}(v) - \bar{\Pi}^{\text{FB}}(s))^+\right]$. Truthful reporting is achieved via a combination of instantaneous payments and promised future payments. Theorem 5.3 implies that the rate of convergence to first-best of this mechanism is $(1-\beta)^{1/2}$.

D.2 Principal-Agent Model with Private Cost Information

Consider a dynamic principal-agent problem as in Krishna et al. (2013). The model herein presented is general and captures many real-word examples such as retail franchising, labor contracts, and procurement contracts. A principal contracts with an agent to repeatedly produce output on her behalf. The principal obtains a revenue $R(q)$ when the agent produces q publicly observed units in a time period. The agent's marginal production cost in time period t, denoted by c_t, is privately observed by the agent. The realized costs $(c_t)_{t\ge 1}$ are assumed to be independent and identically distributed across time periods with support $[0, \bar{c}]$, cumulative distribution function $F(c)$. The agent is informed about her cost when making her production decision. We assume that the revenue function $R : \mathbb{R} \to \mathbb{R}_+$ is concave, increasing, differentiable and satisfies $R(0) = 0$.

The principal offers a contract to the agent that specifies a payment as a function of the number of units produced. Since the principal has commitment power, we can restrict attention to direct mechanism in which the agent reports her cost truthfully to the principal. In this setting a stage mechanism is a pair of functions $(q, z) \in \mathcal{M}$ as before. That is, when the agent reports a cost of c, $q(c)$ determines the number of units to be produced by the agent and $z(c)$ determines the payment to be made to the agent. Given a direct dynamic mechanism $\pi \in \mathbb{M}$ the discounted profit of the principal is $\Pi^\pi = \mathbb{E}\left[\sum_{t=1}^\infty \beta^{t-1}\big(R\left(q_t^\pi(c_t)\right) - z_t^\pi(c_t)\big)\right]$, and the discounted utility of the agent is given by $U^\pi = \mathbb{E}\left[\sum_{t=1}^{\infty} \beta^{t-1}\big(z_t^\pi(c_t) - c_t q_t^\pi(c_t)\big)\right]$. We study the principal's problem of designing a contract that maximizes his expected total discounted profits.

First-best mechanism In the first-best outcome (FB) there is no asymmetry of information between the principal and the agent. The average optimal expected profit is given by

$$\bar{\Pi}^{\mathrm{FB}} = \mathbb{E}_c \left[\max_{q \geq 0} \{R(q) - cq\} \right] .$$

We assume that $\lim_{q\to 0} R'(q) = \infty$ and $\lim_{q\to\infty} R'(0) = 0$ to guarantee that the solution is interior for all costs $c \in [0, \bar{c}]$ and the first-order method holds (see, e.g., Laffont and Martimort (2001)). The optimal production quantity is given by the $q^{\mathrm{FB}}(c) = (R')^{-1}(c)$ and the payment to be charged is $z^{\mathrm{FB}}(c) = cq^{\mathrm{FB}}(c)$.

Optimal static mechanism In the optimal static incentive compatible mechanism (S) the principal needs to provide incentives for the agent to report the realized cost truthfully. Let $F(c)$ be the cumulative distribution function and $f(c)$ be the density of the agent's production cost. Assume that the virtual cost $\psi(c) = c + F(c)/f(c)$ is increasing. A standard application of the envelope formula yields that the average optimal profit is given by

$$\bar{\Pi}^{\mathrm{S}} = \mathbb{E}_c \left[\max_{q \geq 0} \{R(q) - \psi(c)q\} \right] ,$$

where the marginal cost of the agent is replaced by $\psi(c)$ to reflect the informational rent that needs to be paid to the agent. The optimal production quantity is given by $q^{\mathrm{S}}(v) = (R')^{-1}\big(\psi(c)\big)$ and the envelope formula yields that the optimal payment is given by $z^{\mathrm{S}}(c) = cq^{\mathrm{S}}(c) + \int_c^{\bar{c}} q^{\mathrm{S}}(x)\mathrm{d}x$. When the principal's revenue function is linear with slope $r > 0$, the optimal static mechanism reduces to take-it-or-leave-it offer with price $\psi^{-1}(r)$.

Dynamic Individual Rationality The optimal dynamic mechanism under DIR produces in the first time period the optimal incentive compatible quantity $q^{\mathrm{S}}(c)$ and "pays" $z^{\mathrm{S}}(c) - \beta/(1-\beta)\bar{\Pi}^{\mathrm{FB}}$, that is, the difference between the optimal incentive compatible payment and an additional upfront fee. In return the principal transfers ownership of the firm to the agent: from then on the first-best quantity $q^{\mathrm{FB}}(c)$ is produced and the agent captures the entire benefit of production $R\left(q^{\mathrm{FB}}(c)\right)$. The average optimal profit of this mechanism is given by

$$\bar{\Pi}^{\mathrm{DIR}} = (1-\beta)\bar{\Pi}^{S} + \beta\bar{\Pi}^{\mathrm{FB}} ,$$

and the convergence rate to first-best is $(1-\beta)$.

Periodic Individual Rationality The Deferred Utility Mechanism has three phases as before. In the savings phase the principal implements the first-best mechanism: the agent produces $q^{\mathrm{FB}}(c)$ and gets paid $z^{\mathrm{FB}}(v)$, the minimum possible amount to cover her cost. Truthful reporting during the savings phase is achieved exclusively via future promises. In the income phase the principal transfers ownership of the firm to the agent: the first-best quantity $q^{\mathrm{FB}}(c)$ is produced and the agent captures the entire benefit of production $R\left(q^{\mathrm{FB}}(c)\right)$. The mechanism initially starts with a promised utility in the low end of the saving phase. Because the PIR constraint is binding in the savings phase, the envelope formula (3) implies that $\mathbb{E}_c[u(c;w)] = w/\beta$ and the state drifts upwards. Thus after an initial period in the savings phase, the mechanism will likely switch to the income phase. Let $\bar{\Pi}^{\mathrm{FB}}(c) = R\big(q^{\mathrm{FB}}(c)\big) - cq^{\mathrm{FB}}(c)$ be the first-best profit when the agent's cost is c. In the initial savings phase the agent derives no utility and the principal keeps the entire surplus of efficient trade $\bar{\Pi}^{\mathrm{FB}}(c)$, while in the later income phase the whole surplus $\bar{\Pi}^{\mathrm{FB}}(c)$ is collected by the agent. Thus the agent's utilities are backloaded. Theorem 4.3 implies that the rate of convergence to first-best of this mechanism is $(1-\beta)^{1/2}$.

Martingale Utilities The Martingale Utility Mechanism has the following structure in this setting. The first-best quantity $q^{\mathrm{FB}}(c)$ is produced and payments are determined using a dynamic threshold s. Let $\bar{\Pi}^{\mathrm{FB}}(c) = R\big(q^{\mathrm{FB}}(c)\big) - cq^{\mathrm{FB}}(c)$ be the first-best profit when the agent's cost is c. The principal captures the entire surplus from efficient trade $\bar{\Pi}^{\mathrm{FB}}(c)$ when the cost is above to this threshold, and the additional surplus $\bar{\Pi}^{\mathrm{FB}}(c) - \bar{\Pi}^{\mathrm{FB}}(s)$ generated when the cost is below this threshold is provided to the agent. That is, the agent is paid her cost $z(c) = cq^{\mathrm{FB}}(c)$ when $c \geq s$ and is paid the principal's marginal revenue $z(v) = sq^{\mathrm{FB}}(s) + R\big(q^{\mathrm{FB}}(c)\big) - R\big(q^{\mathrm{FB}}(s)\big)$ when $c < s$. Thus the agent derives a positive flow utility only if the cost is below the threshold s. The threshold is dynamic and set so that the average continuation utility is equal to the agent's expected flow utility of the current period, or equivalently $(1-\beta)w = \mathbb{E}_c\left[(\bar{\Pi}^{\mathrm{FB}}(c) - \bar{\Pi}^{\mathrm{FB}}(s))^{+}\right]$. Truthful reporting is achieved via a combination of instantaneous payments and promised future payments. Theorem 5.3 implies that the rate of convergence to first-best of this mechanism is $(1-\beta)^{1/2}$.