Dynamic Double Auctions: Towards First Best

Santiago R. Balseiro
Columbia University
srb2155@columbia.edu

Vahab Mirrokni
Google Research
mirrokni@google.com

Renato Paes Leme
Google Research
renatoppl@google.com

Song Zuo
Google Research
szuo@google.com

November 15, 2018

Abstract

We study the problem of designing dynamic double auctions for two-sided markets in which a platform intermediates the trade between one seller offering independent items to multiple buyers, repeatedly over a finite horizon, when agents have private values. Motivated by online advertising and ride-hailing markets, we seek to design mechanisms satisfying the following properties: *no positive transfers*, i.e., the platform never asks the seller to make payments nor buyers are ever paid and *periodic individual rationality*, i.e., every agent should derive a non-negative utility from every trade opportunity. We provide mechanisms satisfying these requirements that are asymptotically efficient and budget-balanced with high probability as the number of trading opportunities grows. Moreover, we show that the average expected profit obtained by the platform under these mechanisms asymptotically approaches first best (the maximum possible welfare generated by the market).

Keywords: Double auctions, two-sided markets, dynamic mechanism design, internet advertising, revenue management.

1 Introduction

Two-sided markets that enable sellers and buyers to trade have received considerable attention in the last decade. Prominent examples include online advertising markets where publishers (the sellers) offer advertising opportunities to advertisers or ride-hailing markets where drivers (the service providers) provide transportation services to passengers (the customers). In these markets, trade is coordinated by an intermediating platform that determines which parties should trade, collects payments from buyers and transfers payments to sellers. When one buyer and one seller trade a single item, the seminal paper Myerson & Satterthwaite (1983) characterized a profit-maximizing mechanism for the platform when agents (the buyer and seller) have private valuations for the item. Because the platform needs to pay information rents to both parties for inducing truthful revelation of preferences, the platform is not able to extract the gains from trade. The same paper shows that, because of agents' incentives, a platform seeking to maximize social welfare cannot allocate efficiently without subsidizing trade.

In many markets, however, the same set of agents trade independent items repeatedly over time. For example, in internet advertising platforms, each user visiting a publisher's website triggers an auction in which multiple advertisers can bid to show an ad, which results in an abundant number of trading opportunities per day. Thus, instead of offering independent trading mechanisms for each item, the platform can design dynamic mechanisms that link decisions across time. It is well-known that by linking decisions across time, it is possible to mitigate the impact of agents' incentives and, when the number of trading opportunities is large, the platform can asymptotically achieve any Pareto efficient outcome (see, e.g., Jackson & Sonnenschein 2007).

In this paper, we study the problem of designing a dynamic double auction for two-sided markets in which a platform intermediates the trade between one seller offering independent items to multiple buyers, repeatedly over a finite horizon.[1] The seller and buyers have private values drawn independently over time from a common-knowledge joint distribution (we allow values to be arbitrarily correlated across agents). Agents' values are sequentially revealed as items become available for sale and agents are uncertain about their value for future items.

Motivated by the markets we study, we design incentive compatible mechanisms (in the dynamic sense) satisfying the following additional properties. Firstly, the platform can not withhold or create items, i.e., every item sold by the seller should be allocated to one buyer. In ride hailing, it is only

[1]We emphasize that our techniques do not rely on the finite horizon model. Our analysis and results can easily be adapted to a discounted infinite horizon model.

possible to allocate a ride to a driver if we also allocate a ride to a passenger and vice versa. In internet advertising it is only possible to fill an ad request from a publisher if we allocate an ad of a buyer. Secondly, mechanisms should have *no positive transfers*, i.e., the platform never asks the seller to make payments nor buyers are ever paid. This property is meant to exclude designs that are unappealing from the practical perspective – in ride-hailing, for example, drivers are typically not requested to make payments and buyers do not receive payments. Thirdly, the mechanism should be *periodic individually rational*, i.e., every agent should derive a non-negative utility from every trade opportunity. This is a common property enforced in the dynamic mechanism design literature. Krähmer & Strausz (2015) use it to model markets where regulations specify withdrawal rights, i.e., each party is able to cancel a transaction at any point in time. Ashlagi et al. (2016) and Mirrokni, Paes Leme, Tang & Zuo (2018) use it to exclude impractical designs such as demanding large participation fees at the beginning of the horizon. Balseiro, Mirrokni & Leme (2017) and Bergemann, Castro & Weintraub (2017) motivate it in the context of exchange markets for internet advertisement.

A profit-optimal mechanism satisfying these properties can be characterized, in theory, via dynamic programming using agents' promised utilities as state variables. The promised utility framework, which was introduced three decades ago by Green (1987), Spear & Srivastava (1987) and Thomas & Worrall (1990), provides an approach to design mechanisms that are dynamic incentive compatible when values are independent over time. Because of the so-called curse of dimensionality, this approach is not tractable because the state space grows exponentially with the size of the problem (given, for example, in terms of the number of players). Moreover, even determining an optimal mechanism when optimal value functions are provided is challenging since, in general, the traditional approach of optimizing the objective point-wise is not applicable.

Our goal is to design approximation mechanisms, i.e., incentive compatible mechanisms satisfying all our design requirements that are efficient to compute and, at the same time, have provable performance guarantees. We benchmark our mechanisms with "first best," the maximum possible welfare generated by the market when agents' private values are observed by the platform. In particular, we seek to design relatively simple mechanisms whose average expected profit (for the platform) achieve first best asymptotically as the number of time periods grows large.

1.1 Our results and techniques

A challenge faced in designing mechanisms using the promised utility framework is to guarantee that promise utilities are achievable by a feasible mechanism. Specifically, promise utilities should be non-negative and, because of the no positive transfer constraint, not too large. Our main strategy is to design mechanisms that carefully control the stochastic process followed by the promised utilities by appropriately coupling these processes.

We address this challenge in three steps. Firstly, in Section 3, we consider a setting with multiple buyers where the seller has a degenerate (deterministic) valuation. This case is simpler since there is no need to elicit the seller's private information. We design the mechanism so that the buyers' promise utilities are only coupled by the allocation, which, in turn, is only dependent on the competitors values (but not on their states). As a result, these processes are only weakly coupled and we can analyze each buyer independently. Secondly, in Section 4, we consider a market with only one seller and one buyer (this is the bilateral trade setting of Myerson & Satterthwaite 1983) both with non-degenerate valuations. Since the decision of allocating to the seller and to the buyer are completely entangled (we allocate to one if and only if we allocate to the other) it is no longer possible to decouple the stochastic processes of their promised utilities. Instead, we design the mechanism so that their promised utilities are perfectly coupled.

Thirdly, in Section 5, we consider the general case with multiple buyers and a seller, all with non-degenerate valuations. The challenge in designing mechanisms for two-sided settings is that if the promised utility of a certain agent ever becomes zero, then the mechanism has no alternative other than stop allocating altogether to that agent. This is particularly problematic if the agent whose promised utility hits zero is the seller, since this would force the mechanism not to allocate to any of the remaining buyers. In Section 5.1, we tackle this challenge by constructing a novel mechanism that perfectly couples the promised utilities of all agents (seller and buyers).

In all cases we show, using martingale concentration arguments, that the mechanisms are asymptotically efficient, i.e., they allocate optimally in all but sublinearly many periods. Moreover, the average expected profit obtained by the platform approaches first best asymptotically and, in the limit, the platform is able to fully extract the gains from trade.

We describe briefly the mechanisms for the first two cases. Our mechanisms attempt to maximize gains from trade by allocating efficiently as much as possible, where the efficient allocation involves allocating the item to the buyer that values it the most whenever the highest value exceeds the seller's cost. Initially, the platform extracts all gains from trade by implementing a first-price rule in which the buyer with the highest value pays his value, and the seller is paid his cost. Truthful reporting is guaranteed by promising agents higher rents in the future if they report "higher" types. When the cumulative payments of buyers exceed a certain threshold, the platform delivers promises by switching to a second-price rule in which the buyer with the highest value pays the second-highest value and the seller is paid the value of the highest buyer. Moreover, we prove that as the length of the horizon gets larger, the fraction of time periods in which the first-price rule is implemented can be made larger, and the platform's profit converges to first best.

Finally, our mechanisms are not weakly budget balanced per period in the sense that the platform might need to subsidize trade in some periods. This should not be surprising given the impossibility result of Myerson & Satterthwaite (1983). However, on aggregate over all time periods, our mechanisms are both budget balanced in expectation and with high probability. Thus, our results imply that, in repeated settings, it is asymptotically possible for multiple parties to trade efficiently without subsidizing trade.

The paper is structured as follows. In Section 2 we introduce our model and some preliminary results. In Section 3 we study one-sided markets with multiple buyers and a seller with deterministic cost, in Section 4 we consider the case of a single buyer and a single seller, and in Section 5 we study two-sided markets with multiple buyers and a single seller.

1.2 Related work

Our paper lies in the intersection of revenue management, dynamic mechanism design, and double auctions. The study of double auctions started with the seminal paper of Myerson & Satterthwaite (1983) and since then there has been a long stream of work trying to overcome their impossibility results. A way to counter the impossibility is to consider settings in which a large number of items are traded. McAfee (1992) shows that it is possible to obtain close to optimal efficiency as the number of items being traded simultaneously grows large. Azevedo & Budish (2017) argue that such mechanism are strategyproof in large markets. Another avenue to counter the impossibility is by designing approximation mechanisms and bounding the efficiency and revenue loss. This approach was started by McAfee (2008) and has generated a very fruitful stream of results: Segal-Halevi et al. (2016), Blumrosen & Mizrahi (2016), Colini-Baldeschi et al (2017), and Colini-Baldeschi et al (2017). In this paper, we counter the impossibility by considering a dynamic setting in which items arrive repeatedly over time and enforcing incentive constraints across time. Finally, McAfee

& Reny (1992) show that the impossibility result of Myerson-Satterthwaite can by bypassed, in the bilateral trade of a single item, when information is correlated and the joint distribution of types satisfies certain stochastic dominance and hazard rate conditions. In contrast, our paper requires no such assumptions on the distribution of values.

Our work is close to Guo & Hörner (2015), Gorokh et al. (2017), Balseiro et al. (2018), who overcome inefficiency in allocation problems by moving to repeated settings. These papers, however, focus on allocation problems without monetary transfers in which the objective of the platform is to maximize social welfare, while we focus on allocation problems with monetary transfers in which the objective of the platform is profit maximization, i.e., maximizing total transfers. Jackson & Sonnenschein (2007) show that by linking decisions across identical and independent decisions problems, the principal can asymptotically achieve any Pareto efficient outcome as the number of time periods grows large. They introduce a budget-based mechanism in which each agent can report each type a limited number of times. The information structure in Jackson & Sonnenschein (2007) is different to ours in that they assume that the agents know their types for all future periods at time zero. Nevertheless, their budget-based mechanism can be implemented in our dynamic setting in which values are revealed sequentially over time. This mechanism, however, is not feasible for our setting as it does not satisfy our more stringent participation constraints.

Designing optimal auctions for intermediaries have been an active area of research especially for static double auctions (Gomes & Mirrokni 2014, Niazadeh et al. 2014). These studies are partly motivated by applications in ad exchanges and identify optimal revenue sharing schemes, but are still restricted to the same limitations identified in the seminal work of Myerson & Satterthwaite. Recently, there has been a study of revenue sharing in repeated auctions by Balseiro et al. (2017). In their setting, they relax the budget balance constraint to hold in aggregate over time, but still enforce the incentive constraints in each period. A second main difference is that they do not account for the seller's incentives.

The other important line of related work is on dynamic mechanism design for revenue management problems. Kakade et al. (2013) and Pavan et al. (2014) provide a framework for designing optimal mechanisms when the value of agents change over time. Vulcano et al. (2002), Gallien (2006), Board & Skrzypacz (2015) and Gershkov & Moldovanu (2014) study the problem of selling non-perishable goods using dynamic pricing. For social welfare maximization, Bergemann & Välimäki (2010) provide a framework generalizing the VCG mechanism to dynamic environment. Athey & Segal (2013) show how efficiency can be obtained together with budget balance. Given our motivation of ride-hailing and internet advertising our setting is closer to papers that study revenue optimal auctions where goods are perishable and valuations are independent across periods (Ashlagi et al. 2016, Mirrokni et al. 2016*a*,*b*, Mirrokni, Paes Leme, Ren & Zuo 2018, Mirrokni, Paes Leme, Tang & Zuo 2018, Mirrokni et al. 2019). Those papers mentioned above provide approximation guarantees with respect to the optimal mechanism but do not compare against the first-best benchmark. Our benchmark is closer to the one of Balseiro, Mirrokni & Leme (2017), who also compare against the first best benchmark. However, their construction only works for a single buyer and a non-strategic seller. The other major difference is that they consider an infinite horizon model with discounting, which makes the decision problem stationary across periods, while we consider a finite horizon setting without discounting.

2 Preliminaries

We consider the problem of a platform designing a dynamic mechanism for a two-sided market with one seller a_0 and multiple buyers $\{a_1, \ldots, a_n\}$ over a finite[2] horizon with periods $t = 1 \ldots T$. In each period t, the seller has an item with private opportunity cost $v_0^t \in \mathbb{R}_+$ for sale and each buyer has private value $v_i^t \in \mathbb{R}_+, \forall i \in [n]$ for getting this item (we refer to v_0^t as the opportunity cost or seller's value interchangeably). The vector of values $\boldsymbol{v}^t = (v_0^t, v_1^t, \ldots, v_n^t)$ for each period t is independently and identically distributed according to a joint cumulative distribution function $\mathcal{F}(\cdot)$. In other words, the cost and valuations are arbitrarily correlated among buyers and seller but are independent across time. Values are supported in the bounded set $\mathcal{V} = [0, \bar{v}]^{n+1}$ and the distribution of values is common-knowledge. The cost of the t-th item v_0^t is privately observed by the seller and the value v_i^t is privately observed by buyer i.

We assume the platform can commit at the beginning of the horizon to a dynamic mechanism that spans the whole time horizon. A (direct) mechanism, with the reports of the private values $\hat{\boldsymbol{v}}^t = (\hat{v}_0^t, \hat{v}_1^t, \ldots, \hat{v}_n^t)$ as input, determines the allocation of the item, $\boldsymbol{q}^t = (q_0^t, q_1^t, \ldots, q_n^t) \in [0,1]^{n+1}$, and the monetary transfer $\boldsymbol{z}^t = (z_0^t, z_1^t, \ldots, z_n^t) \in \mathbb{R}^{n+1}$. We further allow the mechanism to be *dynamic* in the sense that the allocation $\boldsymbol{q}^t$ and the monetary transfer $\boldsymbol{z}^t$ may depend on all the historical reports as well, i.e., $\hat{\boldsymbol{v}}^\tau$ for all $\tau \leq t$. In other words, $q_i^t = q_i^t(\hat{\boldsymbol{v}}^1, \ldots, \hat{\boldsymbol{v}}^t)$ and

[2]Our model can accommodate a random number of time periods by allowing dummy arrivals that are valued at zero by every buyer. For example, we could discretize the horizon into small periods (say, fractions of a second) so that with high probability there is at most one item arriving per period. Denoting by α the probability that an item arrives at a given period, we would set the distribution of values to $1 - \alpha + \alpha\mathcal{F}(x)$, where $\mathcal{F}(x)$ is the distribution of values in the original model. This modification does not impact our results.

$z_i^t = z_i^t(\hat{\boldsymbol{v}}^1, \dots, \hat{\boldsymbol{v}}^t)$. Often it will be convenient to abbreviate $\hat{\boldsymbol{v}}^{1..t} = (\hat{\boldsymbol{v}}^1, \dots, \hat{\boldsymbol{v}}^t)$ to denote a sequence of reports.

An allocation $q_i^t \in [0,1]$ corresponds to the probability that buyer i gets allocated the t-th item. For the seller, it will be convenient to define q_0^t as the probability that the seller *sells* the item. Since the platform cannot create items or destroy them, we impose the feasibility (FSB) constraint that the probability that item is sold is the same probability that the item is bought:

$$q_0^t = \sum_{i=1}^{n} q_i^t. \tag{FSB}$$

The constraint that the platform cannot create items is very natural in most contexts, since items are provided by sellers. The fact that a platform cannot destroy items is important in settings like ride hailing, where the probability that a driver is allocated a ride must be exactly the same as the sum of probabilities that riders are allocated. Similarly in internet advertising, the platform displays an ad from a buyer in the seller's webpage, so for this transaction to happen there must exist one seller and one buyer. This is in sharp contrast to one-sided markets where the mechanism designer can decide when to withhold items exclusively based on buyers reports.

The next constraint we establish is related to the payments z_i^t. For buyers, z_i^t will consist of payments from the buyers to the platform, while for the seller, z_0^t will consist of payments from the platform to the seller. We impose that mechanisms should have no positive transfers, i.e., the platform never asks the seller to make payments nor buyers are ever paid. The no positive transfers (NPT) constraint is formulated as follows:

$$z_i^t \geq 0, \forall i \in [n]; \qquad z_0^t \geq 0. \tag{NPT}$$

Next we will impose an individual rationality constraint. We will enforce this constraint periodically, i.e., for each t, agents are no worse off than their outside option (which we normalize to zero). Buyers measure their period utility as $v_i^t \cdot q_i^t - z_i^t$, while the seller measures his period utility as $z_0^t - v_0^t \cdot q_0^t$. Precisely, the seller's utility is the seller's gains over his opportunity cost. We can formulate the periodic individual rationality (PIR) as follows:

$$v_i^t \cdot q_i^t - z_i^t \geq 0, \forall i \in [n]; \quad z_0^t - v_0^t \cdot q_0^t \geq 0. \tag{PIR}$$

This constraint guarantees that each buyer pays at most his value for the item, the seller is pays at most his cost, and precludes the platform from charging upfront fees. The (PIR) constraint implies ex-post individual rationality, i.e., the cumulative payoff of each agent is at least his outside option in every realization; thus, our mechanisms still apply if we impose this weaker requirement.

We will consider direct incentive compatible mechanisms in which agents are incentivized to report their values truthfully. To this end, the final constraint we impose is dynamic incentive compatibility, i.e., buyers and the seller should be better off reporting truthfully in each period. Unlike the previous constraints that were defined in each time period, incentive compatibility is defined across time. Given a history of reports $\hat{\boldsymbol{v}}^1, \ldots, \hat{\boldsymbol{v}}^{t-1}$ until period t, agent i should be better off reporting his true value in the current period (i.e., $\hat{v}_i^t = v_i^t$) assuming all agents report truthfully in the future and regardless of the competitors' reports for the current time period. Denote by $\phi_i^t(v_i^t\,;\,\hat{\boldsymbol{v}}^{1..t})$ the utility of agent i in period t when his value is v_i^t and the reported values are $\hat{\boldsymbol{v}}^{1..t}$:

$$\begin{aligned}\phi_i^t(v_i^t\,;\,\hat{\boldsymbol{v}}^{1..t}) &= v_i^t \cdot q_i^t(\hat{\boldsymbol{v}}^{1..t}) - z_i^t(\hat{\boldsymbol{v}}^{1..t})\,, \quad \forall i \in [n]\,,\\ \phi_0^t(v_0^t\,;\,\hat{\boldsymbol{v}}^{1..t}) &= z_0^t(\hat{\boldsymbol{v}}^{1..t}) - v_0^t \cdot q_0^t(\hat{\boldsymbol{v}}^{1..t})\,.\end{aligned}$$

A mechanism satisfies dynamic incentive compatibility if:

$$\begin{aligned} v_i^t \in \operatorname*{argmax}_{\hat{v}_i^t} \Bigg\{ & \phi_i^t(v_i^t\,;\,\hat{\boldsymbol{v}}^{1..t-1}, (\hat{v}_i^t, \hat{\boldsymbol{v}}_{-i}^t)) + \\ & \mathop{\mathbb{E}}_{\boldsymbol{v}^{t+1..T}}\Big[\sum_{\tau=t+1}^{T} \phi_i^\tau(v_i^\tau\,;\,\hat{\boldsymbol{v}}^{1..t-1}, (\hat{v}_i^t, \hat{\boldsymbol{v}}_{-i}^t), \boldsymbol{v}^{t+1..\tau})\Big] \Bigg\}\,, \end{aligned} \tag{DIC'}$$

for every vector of competitions' reports $\hat{\boldsymbol{v}}_{-i}^t$, agent i, and time period t. We impose the incentive compatibility constraint in an ex-post sense over current reports: reporting truthfully is a dominant strategy for each agent regardless of the competitors' reports for the current time period. We believe this constraint is more appealing than interim incentive compatibility as agents do not need to form beliefs about the reports of competitors for the current time period (though they do need to hold beliefs for future time periods). Condition (DIC') guarantees, by backwards induction, that at any point in time, the optimal strategy for each agent is to report truthfully in the current period and keep reporting truthfully from that point on.

Finally, we write the utility of an agent as the sum of the utilities in each period:

$$U_i = \mathop{\mathbb{E}}_{\boldsymbol{v}}\left[\sum_{t=1}^{T} \phi_i^t\right]\,.$$

Similarly, the profit of the platform is simply the sum of all agents' payments:

$$\Pi = \mathop{\mathbb{E}}_{\boldsymbol{v}} \left[\sum_{t=1}^{T} \left(\sum_{i=1}^{n} z_i^t - z_0^t \right) \right].$$

Promised utility framework The task of designing dynamic mechanisms appears at first daunting given the enormous design space. When values are independently distributed across time, however, it is without loss of optimality[3] to formulate the mechanism design problem using the *promised utility framework*, which we will adopt here. The main idea of the promised utility framework is to restrict attention to mechanisms that employ the expected utility of each buyer for the remainder of the horizon as states variables (in the language of dynamic programming, the expected utility-to-go). The state, which is referred to as the promised utility, for buyer i at time period t is given by $w_i^t = \mathbb{E}_{\boldsymbol{v}^{t..T}} \left[\sum_{\tau=t}^{T} \phi_i^\tau \right]$.

We will redefine a dynamic mechanism as a mapping of the current report $\boldsymbol{v}^t$ and a state vector $\boldsymbol{w}^t \in \mathbb{R}_+^{n+1}$. We use the same notation for the actual valuations and the reported valuations since the mechanism is set to provide incentives for agents to report truthfully. A dynamic mechanism is then redefined as a tuple $(\boldsymbol{q}, \boldsymbol{z}, \boldsymbol{u}, \boldsymbol{w}^1)$ where $\boldsymbol{w}^1 \in \mathbb{R}_+^{n+1}$ is an initial state, $\boldsymbol{q} = (q_i^t)_{it}$, $\boldsymbol{z} = (z_i^t)_{it}$ and $\boldsymbol{u} = (u_i^t)_{it}$ are sequences of functions for each agent i and each period t, such that each of those functions maps $(\boldsymbol{v}^t; \boldsymbol{w}^t)$ where $\boldsymbol{v}^t = (v_0^t, \dots, v_n^t)$ is the vector of reported valuations and $\boldsymbol{w}^t = (w_0^t, \dots, w_n^t)$ is the vector of promised utilities (to be defined next) to the real line.

Given such functions, the mechanism is executed as follows: in period t and based on the state $\boldsymbol{w}^t$ of promised utilities, the mechanism starts by collecting valuations v_i^t from each agent, then allocates to agent i according to $q_i^t(\boldsymbol{v}^t; \boldsymbol{w}^t)$, executes transfers according to $z_i^t(\boldsymbol{v}^t; \boldsymbol{w}^t)$, and determines the state $\boldsymbol{w}^{t+1}$ for the next period according to:

$$w_i^{t+1} = u_i^t(\boldsymbol{v}^t; \boldsymbol{w}^t).$$

With some abuse of notation, we will use $\phi_i^t(v_i, \hat{\boldsymbol{v}}^t; \boldsymbol{w}^t)$ to denote the utility of agent i in period t when his value is v_i, the reported values are $\hat{v}_i$ and the promised utilities are $\boldsymbol{w}^t$:

$$\phi_i^t(v_i, \hat{\boldsymbol{v}}^t; \boldsymbol{w}^t) = v_i^t \cdot q_i^t(\hat{\boldsymbol{v}}^t; \boldsymbol{w}^t) - z_i^t(\hat{\boldsymbol{v}}^t; \boldsymbol{w}^t), \quad \forall i \in [n],$$
$$\phi_0^t(v_0, \hat{\boldsymbol{v}}^t; \boldsymbol{w}^t) = z_0^t(\hat{\boldsymbol{v}}^t; \boldsymbol{w}^t) - v_0^t \cdot q_0^t(\hat{\boldsymbol{v}}^t; \boldsymbol{w}^t).$$

[3]In the sense that for every dynamic mechanism there is a dynamic mechanism in the promised utility framework achieving at least the same revenue and welfare

We next define when a promised utility mechanism is dynamic incentive compatible. A mechanism is said to be incentive compatible if it satisfies two properties. Firstly, the mechanism should satisfy the promise keeping (PK) constraint:

$$w_i^t = \mathop{\mathbb{E}}_{\boldsymbol{v}^t}\left[\phi_i^t(v_i, \boldsymbol{v}^t; \boldsymbol{w}^t) + u_i^t(\boldsymbol{v}^t; \boldsymbol{w}^t)\right], \tag{PK}$$
$$\forall i \in \{0, 1, \dots, n\}, t \le T \text{ and } w_i^{T+1} = 0,$$

which says that w_i^t corresponds to the aggregate utility obtained by agent i from periods $t, t+1, \dots, T$ assuming that he and the other agents report truthfully. This guarantees that the mechanism delivers the utility promised to each agent. Secondly, the mechanism should satisfy the dynamic incentive compatibility (DIC) constraint:

$$\phi_i^t\left(v_i^t, (v_i^t, \hat{\boldsymbol{v}}_{-i}^t); \boldsymbol{w}^t\right) + u_i^t\left((v_i^t, \hat{\boldsymbol{v}}_{-i}^t); \boldsymbol{w}^t\right) \ge \tag{DIC}$$
$$\phi_i^t\left(v_i^t, (\hat{v}_i^t, \hat{\boldsymbol{v}}_{-i}^t); \boldsymbol{w}^t\right) + u_i^t\left((\hat{v}_i^t, \hat{\boldsymbol{v}}_{-i}^t); \boldsymbol{w}^t\right),$$

which says that each agent maximizes the current period utility plus the expected utility over future periods by reporting truthfully in this period.[4]

In this model, we are interested in designing $(\boldsymbol{q}, \boldsymbol{z}, \boldsymbol{u}, \boldsymbol{w}^1)$ to maximize the profit of the platform, which can be informally stated as follows:

$$\begin{aligned} \text{maximize} \quad & \Pi \\ \text{subject to} \quad & \text{(DIC) (PIR) (PK) (FSB)}. \end{aligned} \tag{OPT}$$

First-best benchmark The profit benchmark that we will consider is the "first best," which corresponds to the maximum possible welfare achievable when agents' private values are observed by the platform:

$$\Pi^{\mathsf{FB}} = T \cdot \mathbb{E}\left[\left(\max_{i=1..n} v_i - v_0\right)^+\right],$$

where we denote by $(x)^+ = \max(x, 0)$ the positive part of a number $x \in \mathbb{R}$. We will study mechanisms that have a sublinear additive approximation to the first best benchmark, i.e., $\Pi = \Pi^{\mathsf{FB}} - o(T)$ which implies that $\Pi/\Pi^{\mathsf{FB}} \to 1$ as $T \to \infty$.

[4] If a mechanism in the promised utility framework satisfies (PK) and (DIC) then its associated dynamic mechanism in usual form $\boldsymbol{q}^t(\hat{\boldsymbol{v}}^1, \dots, \hat{\boldsymbol{v}}^t) = \boldsymbol{q}^t(\hat{\boldsymbol{v}}^t; \boldsymbol{u}^{t-1}(\hat{\boldsymbol{v}}^{t-1}, \boldsymbol{u}^{t-2}(\hat{\boldsymbol{v}}^{t-2}, \dots)))$ and similarly for $\boldsymbol{z}^t$ satisfies (DIC').

2.1 Dynamic programming formulation

Invoking the Principle of Optimality, the platform's problem can be, in theory, solved recursively by backwards induction. Although this reduction is standard in the theory of dynamic programming, we supply the details here for completeness. Define $\Pi^t(\boldsymbol{w})$ to be the optimal profit-to-go that can be obtained from periods t to T if the promised utility at period t is equal to $\boldsymbol{w}$. We solve the following programs back from $t = T$ to $t = 1$. The Bellman equation is:

$$\begin{aligned} \Pi^t(\boldsymbol{w}^t) = \max_{q_t, z_t, u_t} \mathbb{E}_{\boldsymbol{v}^t} \Bigg[& \Pi^{t+1}(\boldsymbol{u}^t(\boldsymbol{v}^t; \boldsymbol{w}^t)) \\ & + \sum_{i=1}^n z_i^t(\boldsymbol{v}^t; \boldsymbol{w}^t) - z_0^t(\boldsymbol{v}^t; \boldsymbol{w}^t) \Bigg] \\ \text{s.t.} \quad & (\text{DIC}), (\text{PIR}), (\text{PK}), (\text{F}\textsc{sb}), (\text{NPT}), \\ & u^t(\boldsymbol{v}^t; \boldsymbol{w}^t) \geq 0, \quad \text{for period } t \end{aligned} \tag{DP}$$

with the boundary condition that $\Pi^{T+1}(\boldsymbol{w}) = -\boldsymbol{w}_0$ if $w_i = 0$ for all $i \in [n]$ and $\Pi^{T+1}(\boldsymbol{w}) = -\infty$ otherwise. This boundary condition reflects that, by (NPT), no transfers are allowed to the buyers at the end of the horizon and thus the promise utilities of all buyers should be zero at that point. The platform, however, can refund the seller at the end of the horizon. This is equivalent to adding the constraint $u_i^T(\boldsymbol{v}^T; \boldsymbol{w}^T) = 0$ for $i \in [n]$ and using the boundary condition $\Pi^{T+1}(\boldsymbol{w}) = -\sum_{i=0}^n w_i$.

Because the state space of the value function grows exponentially with the number of agents, this dynamic program can not be efficiently solved in practice. Thus motivated, we design approximation mechanisms that are easy to implement and have provable performance guarantees.

3 One-sided markets

Our strategy will be to propose a mechanism, show that it satisfies the constraints, and bound its profit with respect to first best. We will do in three stages of increasing complexity. We start in this section with a one-sided market where buyers have non-degenerate valuations (i.e., not deterministic) and the seller has a constant value. Then we move in Section 4 to a setting with a single buyer and a single seller where both the buyer and the seller have non-degenerate valuations. Finally, in Section 5 we study the general case with one seller and multiple buyers, all with non-degenerate valuations.

For the mechanisms in Sections 3 and 4, we will use the following notation. For buyers $i = 1..n$

we define:

$$\underline{w}_i = \mathop{\mathbb{E}}_{\boldsymbol{v}}\left[(v_i - \max \boldsymbol{v}_{-i})^+\right] , \quad \overline{w}_i^t = (T - t + 1)\underline{w}_i ,$$

where $\max \boldsymbol{v}_{-i} := \max_{j=0\ldots n: j\neq i} v_j$ stands for the maximum of all other valuations except v_i. Intuitively, $\underline{w}_i$ captures agent's i expected surplus from efficient trade achieved, for example, by the VCG auction. The upper bound $\overline{w}_i^t$ captures, in turn, the expected surplus from efficient trade achievable from period t until the end of the horizon. Those values define a desirable range $[\underline{w}_i, \overline{w}_i^t]$ for the promised utility w_i^t. We will set up the mechanism in such a way that we will be able to allocate efficiently whenever w_i is in this range and will argue that by setting the initial promise properly, each buyer will remain in the desirable range with high probability for all but a sublinear number of periods.

For simplicity, we start with the case where the valuation of the seller is equal to a constant v_0 with probability one. In that case, the platform is always able to purchase the item from the seller at a price v_0 whenever needed. There is no need to elicit information from the seller or worry about his incentive constraints, so we focus on providing the adequate incentives for the buyers.

3.1 Envelope formula

In designing our approximation mechanisms, we first design the allocation and payment rule and then construct a promise utility function to guarantee incentive compatibility. Before moving forward, we discuss briefly how this approach works in the one-sided case. Given a mechanism $(\boldsymbol{q}, \boldsymbol{z}, \boldsymbol{u})$, let $U_i(\boldsymbol{v}; \boldsymbol{w}) = v_i \cdot q_i(\boldsymbol{v}; \boldsymbol{w}) - z_i(\boldsymbol{v}; \boldsymbol{w}) + u_i(\boldsymbol{v}; \boldsymbol{w})$ the utility-to-go of buyer i when the state is $\boldsymbol{w}$ and agents report $\boldsymbol{v}$. Using the envelope formula (see, e.g., Myerson 1981) we obtain that a necessary and sufficient condition for a mechanism to be dynamic incentive compatible is that the allocation $q_i(\boldsymbol{v}; \boldsymbol{w})$ is monotone in v_i and

$$U_i(\boldsymbol{v}; \boldsymbol{w}) = U_i(x, \boldsymbol{v}_{-i}; \boldsymbol{w}) + \int_x^{v_i} q_i(v, \boldsymbol{v}_{-i}; \boldsymbol{w}) \mathrm{d}v.$$

where x is the "lowest" type, which is $x = 0$ for buyers $i \in [n]$. (A similar formula holds for the seller with the lowest type equal to $x = \bar{v}$ and the end points of the integral reversed). The integral in the right-hand side is typically referred to as the information rent. In the mechanisms given in Section 3 and Section 4 we will set the utility of the lowest type $U_i(x, \boldsymbol{v}_{-i}; \boldsymbol{w})$ to be independent of the realization of values of the other agents, i.e., we set $U_i(x, \boldsymbol{v}_{-i}; \boldsymbol{w}) = U_i(x; \boldsymbol{w})$ for all $\boldsymbol{v}_{-i}$.

Because the promise keeping constraint implies that $w_i = \mathbb{E}_{\boldsymbol{v}}\left[U_i(\boldsymbol{v};\boldsymbol{w})\right]$ we can solve for the utility of the lowest type $U_i(x;\boldsymbol{w})$ and obtain that

$$\begin{aligned} & v_i \cdot q_i(\boldsymbol{v};\boldsymbol{w}) - z_i(\boldsymbol{v};\boldsymbol{w}) + u_i(\boldsymbol{v};\boldsymbol{w}) \\ &= w_i + \int_x^{v_i} q_i(v,\boldsymbol{v}_{-i};\boldsymbol{w})\mathrm{d}v - \mathbb{E}_{\boldsymbol{v}}\left[\int_x^{v_i} q_i(v,\boldsymbol{v}_{-i};\boldsymbol{w})\mathrm{d}v\right]. \end{aligned} \tag{PI}$$

We shall use this formula to construct dynamic incentive compatible mechanisms: for any feasible, monotone allocation rule and payment rule, we can use (PI) to pin down a promise utility rule that makes the mechanism dynamic incentive compatible.

3.2 Mechanism definition

Since our goal is to achieve first best, we should try to allocate as efficiently as possible. In each period, we will identify the buyer with the largest valuation. If the highest valuation is lower than the seller's value v_0, the item is not allocated. If trade is possible, i.e. $\max_{i\in[n]} v_i > v_0$, then we will allocate to the top bidder if his promised utility is in the range $[\underline{w}_i, \overline{w}_i^t]$. If w_i is in the range $[0, \underline{w}_i)$ the allocation will be inefficient and we will choose a probability (depending on w_i) to allocate to the top bidder and with the remaining probability we will leave the item with the seller.

Formally, let i be the buyer with largest v_i (breaking ties lexicographically). If $v_i \le v_0$, we leave the item with the seller. If on the other hand gains from trade are possible (i.e. $v_i > v_0$), then we allocate using one of two different regimes determined by the promised utility w_i. The different allocation regions are shown in Figure 1a. Whenever the regions overlap, we use the mechanism of the low promised utility region.

High promised utility (second-price auction): If $\bar{w}_i^t - \bar{v} \le w_i^t \le \bar{w}_i^t$ (where $\bar{v}$ is the valuation of the highest type), we allocate the item to buyer i and charge the second highest price. For the winning buyer i:

$$q_i^t = 1\,, \qquad z_i^t = \max \boldsymbol{v}_{-i}^t\,, \qquad u_i^t = w_i - \underline{w}_i\,.$$

When the promised utility is high, the buyer with the highest value is charged the second-highest price so that its promise utility decreases by $\underline{w}_i$. This guarantees that no buyer reaches the end of the horizon with a positive promised utility, which is required because of the no positive transfers constraint.

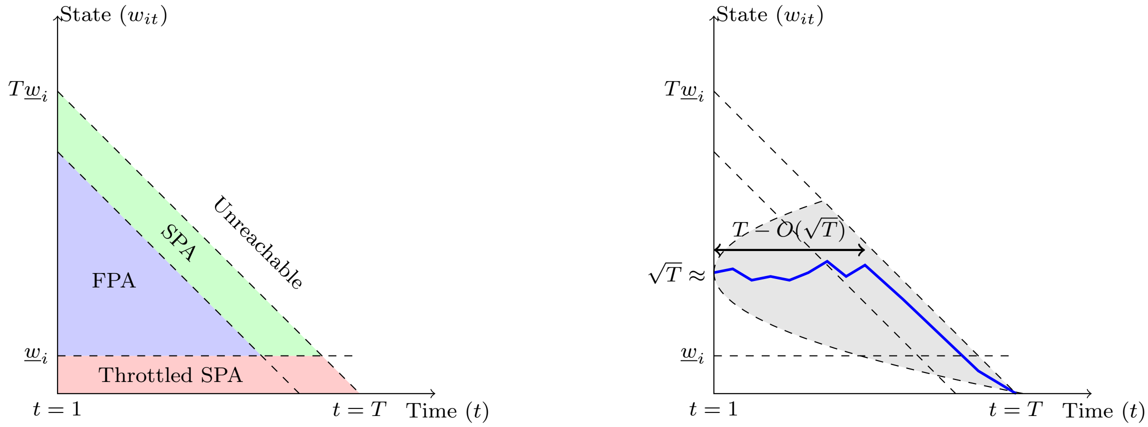

(a) Allocation rules for different states. States above $\bar{w}_i^t$ are not reachable by the mechanism.

(b) Typical path of the promised utility (in blue) and confidence bands showing mostly likely paths (in gray).

Figure 1: Illustration of the one-sided mechanism

Medium promised utility (first-price auction): If $\underline{w}_i \le w_i^t \le \bar{w}_i^t - \bar{v}$, we allocate the item to buyer i and charge the bid. For the winning buyer i:

$$q_i^t = 1\,, \qquad z_i^t = v_i^t\,, \qquad u_i^t = w_i^t - \underline{w}_i + (v_i^t - \max \boldsymbol{v}_{-i}^t)\,.$$

In this region the platform captures the entire buyer surplus.

Low promised utility (throttled second-price auction): If $0 \le w_i^t < \underline{w}_i$, then we allocate to buyer i with probability $q_i = w_i^t/\underline{w}_i$ and charge the second-highest price if allocated. In such case we update the promised utility to zero.

$$q_i^t = w_i^t/\underline{w}_i\,, \qquad z_i^t = q_i^t \cdot \max \boldsymbol{v}_{-i}^t\,, \qquad u_i^t = 0\,.$$

In this region the allocation is throttled to prevent the promise utility from being negative.

Losing buyers: For all buyers other than the winner (or all buyers in case $v_0^t \ge \max_{i\in[n]} v_i^t$), we update the promised utility as follows:

$$q_j^t = 0\,, \qquad z_j^t = 0\,, \qquad u_j^t = (w_j^t - \underline{w}_j)^+\,.$$

3.3 Mechanism analysis

Note that we only defined the mechanism for $w_i^t \leq \overline{w}_i^t$. We start by showing that if the initial promises $\boldsymbol{w}^1$ are not too high, then the promise utility of agent i stays within the interval $[0, \overline{w}_i^t]$.

Lemma 3.1. *If $w_i^1 \leq \overline{w}_i^1 = T\underline{w}_i$, then $w_i^t \leq \overline{w}_i^t$ for all t. Moreover, $w_i^t \geq 0$ for all t.*

Proof. It follows by induction on t that $w_i^t \leq \overline{w}_i^t = (T-t+1)\underline{w}_i$. In all but one case, the promised utility changes to $w_i^{t+1} = w_i^t - \underline{w}_i$ or zero, in which case the induction holds trivially. The only remaining case is the winning buyer with medium promised utility, in which case:

$$\begin{aligned} w_i^{t+1} &= w_i^t - \underline{w}_i + v_i^t - \max \boldsymbol{v}_{-i}^t \\ &\leq \bar{w}_i^t - \bar{v} + \bar{v} - \underline{w}_i = \overline{w}_i^{t+1}, \end{aligned}$$

where the inequality follows because $w_i^t \leq \overline{w}_i^t - \overline{v}$ in this region and $0 \leq v_j^t \leq \bar{v}$ for all $j \in [n]$. The claim that $w_i^t \geq 0$ for all t follows trivially for all cases, other perhaps when the winning buyer has medium promised utility, in which case

$$w_i^{t+1} = w_i^t - \underline{w}_i + v_i^t - \max \boldsymbol{v}_{-i}^t \geq v_i^t - \max \boldsymbol{v}_{-i}^t \geq 0,$$

where the first inequality follows because $w_i^t \geq \underline{w}_i$ in this region, and the last inequality because i is the winner. □

From taking $t = T$ in the previous lemma, we get as a corollary that $u_i^T(\boldsymbol{v}^T) = 0$ for all $\boldsymbol{v}^T$, implying that all promise utilities are zero at the end of the horizon, as required. Next we argue that this dynamic mechanism satisfies the constraints imposed by problem (Opt). Above we defined the q_i, z_i, u_i for the buyers. Since the seller has a trivial valuation, we can use the following trivial definitions for the seller:

$$q_0^t = \sum_{i=1}^n q_i^t, \qquad z_0^t = v_0^t \cdot q_0^t, \qquad w_0^1 = u_0^t = 0.$$

Lemma 3.2. *If $0 \leq w_i^1 \leq \overline{w}_i^1$, then the mechanism described satisfies (*DIC*), (*PIR*), (*PK*), (*NPT*) and (*Fsb*).*

Proof. Feasibility (Fsb) follows from the fact we only allocate to one buyer in each period. Periodic individual rationality (PIR) is satisfied since we charge the winner either his value or the second-highest value. Since buyer payments are always non-negative, (NPT) is satisfied. Finally, to check dynamic incentive compatibility (DIC) and promise keeping (PK), it suffices to check that the mechanism satisfies (PI) since the allocations are trivially monotone. This follows by construction. We provide details for completeness. Dropping the dependence on time and the state to simplify the notation, we obtain that the promise utility functions should satisfy:

$$u_i(\boldsymbol{v}) = w_i - v_i \cdot q_i(\boldsymbol{v}) + \int_0^{v_i} q_i(v, \boldsymbol{v}_{-i}) \mathrm{d}v + z_i(\boldsymbol{v}) \\ - \mathop{\mathbb{E}}_{\boldsymbol{v}} \left[\int_0^{v_i} q_i(v, \boldsymbol{v}_{-i}) \mathrm{d}v \right] .$$

For the buyers in the low promised utility region it follows that $u_i(\boldsymbol{v}) = 0$ because $q_i(\boldsymbol{v}) = w_i/\underline{w}_i \mathbf{1}\{v_i \geq \max \boldsymbol{v}_{-i}\}$, $z_i(\boldsymbol{v}) = (\max \boldsymbol{v}_{-i}) q_i(\boldsymbol{v})$, $\int_0^{v_i} q_i(v, \boldsymbol{v}_{-i}) \mathrm{d}v = w_i/\underline{w}_i (v_i - \max \boldsymbol{v}_{-i})^+ = (v_i - \max \boldsymbol{v}_{-i}) q_i(\boldsymbol{v})$, and $\mathbb{E}_{\boldsymbol{v}} \left[\int_0^{v_i} q_i(v, \boldsymbol{v}_{-i}) \mathrm{d}v \right] = w_i$. For the buyers in the medium promised utility region it follows that $u_i(\boldsymbol{v}) = w_i - \underline{w}_i + (v_i - \max \boldsymbol{v}_{-i})^+$ from the facts that $v_i q_i(\boldsymbol{v}) - z_i(\boldsymbol{v}) = 0$ since (PIR) is binding and $\int_0^{v_i} q_i(v, \boldsymbol{v}_{-i}) \mathrm{d}v = (v_i - \max \boldsymbol{v}_{-i})^+$ since $q_i(\boldsymbol{v}) = \mathbf{1}\{v_i \geq \max \boldsymbol{v}_{-i}\}$. The case for the buyers in the high promised utility region follows similarly. □

3.4 Profit approximation

Finally, we argue that the revenue of the mechanism approaches first best. We will analyze the stochastic process defined by the promised utility w_i^t of each buyer i. A typical path for this process is depicted in Figure 1b. Before w_i^t hits the second-price auction region (green region in Figure 1a) or the throttled region (red region in the figure), this process satisfies $w_i^{t+1} = w_i^t + (v_i^t - \max \boldsymbol{v}_{-i}^t)^+ - \underline{w}_i$ and thus behaves like a martingale from our definition of $\underline{w}_i$. By standard martingale concentration inequalities, w_i^t does not deviate by more than $O(\sqrt{t \log t})$ from the initial promised utility with high probability. Recall that the platform captures the entire buyer surplus when the promise utilities are at the medium level. Therefore, we need to set the initial promised utility not so small to avoid reaching the throttled region too early but not so large that it will hit the second-price region. The sweet spot turns out to be around $O(\sqrt{T \log T})$, which will cause the mechanism to stay in the first-price region for around $\Omega(T - \sqrt{T})$ periods.

Theorem 3.3. *If the initial promised utilities are set to* $w_i^1 = \bar{v}\sqrt{8T \log T} + \underline{w}_i$*, then the expected profit of the mechanism, denoted by* Π*, satisfies:*

$$\Pi \geq \Pi^{\mathsf{FB}} - O\left(\sqrt{T \log T}\right) ,$$

whenever $w_i^1 \le \overline{w}_i^1 - \bar{v}$. *In particular, the mechanism is asymptotically optimal, i.e.,* $\Pi/\Pi^{\mathsf{FB}} \to 1$ *as* $T \to \infty$.

We conjecture that this bound is tight up to, perhaps, log factors. We now present the proof of Theorem 3.3:

Proof. The mechanism is feasible by Lemma 3.2 because $0 \le w_i^1 \le \overline{w}_i^1$. Let S be the expected social welfare of the mechanism, which in virtue of the promised utility framework is given by $S = \Pi + \sum_{i=0}^n w_i$. We focus on bounding S in terms of the optimal social welfare Π^{FB}.

Step 1. Let $\tau_i = \inf\{t \ge 1 : w_i^t < \underline{w}_i\}$ be the stopping time measuring the first time in which the promised utility w_i^t falls in the throttled region, i.e. $w_i^t < \underline{w}_i$. Let also $\hat{q}_i^t$ be the efficient allocation, which allocates the item to buyer if $v_i \ge \max \boldsymbol{v}_{-i}$ (breaking ties lexicographically). Because the mechanism allocates "efficiently" to agent i until time τ_i, we can write:

$$\begin{aligned}
\Pi^{\mathsf{FB}} - S &= \mathbb{E}_{\boldsymbol{v}}\left[\sum_{t=1}^{T}\sum_{i\in[n]}(v_i^t - v_0^t)\cdot \hat{q}_i^t - \sum_{i\in[n]}(v_i^t - v_0^t) q_i^t\right] \\
&= \sum_{i\in[n]} \mathbb{E}_{\boldsymbol{v}}\left[\sum_{t=\tau_i}^{T}(v_i^t - v_0^t)\cdot(\hat{q}_i^t - q_i^t)\right] \\
&\le \bar{v}\sum_{i\in[n]}\mathbb{E}_{\boldsymbol{v}}\left[(T-\tau_i+1)^+\right] ,
\end{aligned}$$

where the second equation follows from the definition of the stopping time τ_i and linearity of expectation, and the inequality because values satisfy $0 \le v_i^t \le \bar{v}$. We proceed by bounding the terms $\mathbb{E}_{\boldsymbol{v}}\left[(T-\tau_i+1)^+\right]$.

Step 2. Now define

$$\mu_i^t = w_i^t + \sum_{\tau=1}^{t-1}(v_i^\tau \cdot q_i^\tau - z_i^\tau).$$

and observe that by (PIR), $\mu_i^t \ge w_i^t$. Let $\mathcal{H}_t = \sigma(\boldsymbol{v}_1, \ldots, \boldsymbol{v}_{t-1})$ be the natural filtration. By (PK) we have $\mathbb{E}[\mu_i^{t+1}|\mathcal{H}_t] = \mu_i^t + \mathbb{E}[v_i^t \cdot q_i^t - z_i^t + w_i^{t+1}|\mathcal{H}_t] - w_i^t = \mu_i^t$. Therefore, $\{\mu_i^t\}_{t=1}^T$ is a martingale with respect to the natural filtration.

Because the second-price region and the throttling regions are absorbing we have that $w_i^\tau = \mu_i^\tau,\ \forall \tau \le t$ as long as $w_i^t \le \overline{w}_i^t - \bar{v}$, and once $w_i^t > \overline{w}_i^t - \bar{v}$ we have $w_i^\tau > \overline{w}_i^\tau - \bar{v},\ \forall \tau \ge t$. Let

$T_i^* = T - \overline{v}/\underline{w}_i$ and consider a time $1 \le t \le T_i^*$. Note that $T_i^* \ge 1$ because $\underline{w}_i \le w_i^1 \le \overline{w}_i^1 - \bar{v}$. Because in these time periods the second-price region and the throttling region do not intersect (i.e., $\underline{w}_i \le \bar{w}_i^t - \bar{v}$) we obtain that

$$\begin{aligned}\Pr[\tau_i \le t] = \Pr\left[\min_{s \le t} \mu_i^s \le \underline{w}_i\right] &\le \exp\left(-\frac{(w_i^1 - \underline{w}_i)^2}{8\sum_{k=1}^t \overline{v}^2}\right) \\ &\le \exp\left(-\frac{(w_i^1 - \underline{w}_i)^2}{8t\overline{v}^2}\right),\end{aligned}$$

where the inequality follows from the maximal version of Azuma's concentration inequality[5], because $\{\mu_i^t\}_{t=1}^T$ is a martingale with increments bounded by $|\mu_i^t - \mu_i^{t+1}| \le 2\bar{v}$ and $\mu_i^1 = w_i^1$. This implies that

$$\begin{aligned}&\mathbb{E}_{\boldsymbol{v}}[(T - \tau_i + 1)^+] \\ &= \sum_{t=1}^T \Pr[\tau_i \le t] \le \sum_{t=1}^{T_i^*} \Pr[\tau_i \le T_i^*] + T - T_i^* \\ &\le T_i^* \exp\left(-\frac{(w_i^1 - \underline{w}_i)^2}{8T_i^*\overline{v}^2}\right) + T - T_i^*,\end{aligned}$$

where the first equation follows from summation by parts and the first inequality because probabilities are at most one. By letting $w_i^1 = \overline{v}\sqrt{8T\log T} + \underline{w}_i$ and using that $T_i^* \le T$, we have

$$\begin{aligned}\sum_{t=1}^T \Pr[\tau_i \le t] &\le T \cdot \exp(-\log T) + T - T_i^* \\ &= 1 + T - T_i^* = \overline{v}/\underline{w}_i + 1.\end{aligned}$$

Hence,

$$\begin{aligned}\Pi^{\mathsf{FB}} - S &\le \bar{v} \sum_{i \in [n]} \mathbb{E}_{\boldsymbol{v}}\left[(T - \tau_i + 1)^+\right] \\ &\le \bar{v} \sum_{i \in [n]} (\overline{v}/\underline{w}_i + 1) = O(1).\end{aligned}$$

[5]For a martingale $\{M_k\}_{k=0}^K$ with $|M_k - M_{k-1}| \le c_k$, then for any $X < M_0$, $\Pr\left[\min_{k \le K} M_k \le X\right] \le \exp\left(-\frac{(M_0 - X)^2}{2\sum_{k=1}^K c_k^2}\right)$.

Step 3. To conclude, we observe that $\Pi = S - \sum_i w_i^1$. Since the initial promise utilities satisfy $w_i^1 = O(\sqrt{T \log T})$, we get that the expected profit of our mechanism satisfies:

$$\begin{aligned}\Pi = S - \sum_i w_i^1 &= \Pi^{\mathsf{FB}} - O(1) - O\left(\sqrt{T \log T}\right) \\ &= \Pi^{\mathsf{FB}} - O\left(\sqrt{T \log T}\right) ,\end{aligned}$$

and the result follows. □

4 Bilateral trade

The next setting we consider is a repeated version of the bilateral trade model of Myerson and Satterthwaite where there is a single buyer and a single seller and both parties have private, non-degenerate information. The values (v_0^t, v_1^t) are drawn from a joint distribution $\mathcal{F}$. The values can be correlated between the buyer and the seller but are otherwise independent across time.

We will keep the same desirable interval for the range of promised utilities of the buyer. Namely,

$$\underline{w}_1 = \mathbb{E}[(v_1 - v_0)^+] , \qquad \overline{w}_1^t = (T - t + 1) \cdot \underline{w}_1 .$$

Here $\underline{w}_1$ captures the buyer and seller surplus from efficient trade achieved, for example, by the VCG auction. In this setting the VCG auction involves trading the item whenever the buyer's value exceeds the seller's value, the buyer pays the seller's value and the seller is paid the buyer's value.

It will be convenient to define a mechanism in which the promised utilities of the buyer and the seller are perfectly coupled as follows:

$$w_0^t = w_1^t .$$

We next discuss the motivation behind this design choice. Recall that the platform can neither create items by (FSB) nor pay back to the buyer by (NPT). Thus, the platform can only deliver promises to the buyer by allocating the item to him. At the same time, because the seller is strategic, he needs to be given proper incentives to trade. By coupling the promise utilities we avoid a situation in which the promised utility of the seller is zero but that of the buyer is positive. In this case, the platform would be forced not to trade with the seller and bound to deliver some value to the buyer, which leads to an impossibility.

4.1 Mechanism definition

The mechanism will have a similar format to the previous one: we will allocate efficiently whenever $w_1^t \in [\underline{w}_1, \overline{w}_1^t]$ and throttle the allocation whenever w_1^t is below $\underline{w}_1$.

No trade: If $v_1^t \le v_0^t$, no trade happens and the mechanism rules are as follows:

$$q_0^t = 1\,, \qquad q_1^t = 0\,, \qquad z_0^t = z_1^t = 0\,,$$
$$u_1^t = (w_1^t - \underline{w}_1)^+\,, \qquad u_0^t = (w_0^t - \underline{w}_1)^+\,.$$

In the following cases, we assume that efficient trade is always possible, i.e., $v_1^t > v_0^t$.

High promised utility (VCG prices): If $\overline{w}_1^t - \bar{v} \le w_1^t \le \overline{w}_1^t$, the platform executes the trade and prices according to the VCG mechanism:

$$q_1^t = 1\,, \qquad z_1^t = v_0^t\,, \qquad u_1^t = w_1^t - \underline{w}_1\,,$$

$$q_0^t = 1\,, \qquad z_0^t = v_1^t\,, \qquad u_0^t = w_0^t - \underline{w}_1\,.$$

In this case the buyer pays the seller's value and the seller is paid the buyer's value. Thus, the platform needs to subsidize trade. As in the one-sided case, promise utilities decrease by $\underline{w}_1$ to guarantee that no agent reaches the end of the horizon with a positive promised utility.

Medium promised utility (first-price auction): If $\underline{w}_1 \le w_1^t \le \overline{w}_1^t - \bar{v}$, the platform executes the trade and prices at the bid:

$$q_1^t = 1\,, \qquad z_1^t = v_1^t\,, \qquad u_1^t = w_1^t - \underline{w}_1 + v_1^t - v_0^t\,,$$

$$q_0^t = 1\,, \qquad z_0^t = v_0^t\,, \qquad u_0^t = w_0^t - \underline{w}_1 + v_1^t - v_0^t\,.$$

In this case agents pay their values and the platform extracts all the gains from trade.

Low promised utility (throttled allocation): If $w_1^t < \underline{w}_1$ the mechanism will execute the trade with probability $w_1^t / \underline{w}_1^t$ if $v_1^t > v_0^t$. In such case, the parties are paid the VCG prices for that transaction, i.e., the buyer is charged v_0^t and the seller's payment is v_1^t. The mechanism is described as follows:

$$q_1^t = w_1^t / \underline{w}_1^t\,, \qquad z_1^t = v_0^t \cdot q_1^t\,, \qquad u_1^t = 0\,,$$

$$q_0^t = w_1^t / \underline{w}_1^t\,, \qquad z_0^t = v_1^t \cdot q_0^t\,, \qquad u_0^t = 0\,.$$

4.2 Mechanism analysis

We again only defined the mechanism for $w_1^t \le \overline{w}_1^t$. Using the same argument as in Lemma 3.1, we can argue that if $w_1^1 \le \overline{w}_1^1$, then $w_1^t \le \overline{w}_1^t$ for all t. The next step is to argue that the mechanism proposed satisfies the constraints imposed.

Lemma 4.1. *If $0 \le w_1^1 \le \overline{w}_1^1$, then the bilateral trade mechanism described satisfies (*DIC*), (*PIR*), (*NPT*), (*PK*) and (*FSB*). Moreover, promise utilities are coupled, i.e., $w_0^t = w_1^t$ for all t.*

Proof. From the buyer's perspective, the mechanism is identical to the one-sided mechanism in Section 3, so it follows from Lemma 3.2 that the buyer-side constraints are satisfied. For the seller, constraints (PIR), (NPT) and (FSB) can be trivially checked. For dynamic incentive compatibility (DIC) and (PK) it suffices to check that the mechanism satisfies (PI) since the allocation is trivially monotone. This follows by construction. □

4.3 Profit approximation and budget balance

Finally, because the promised utilities of the buyer and the seller are coupled, the profit approximation follows from the same argument as in the previous section.

Theorem 4.2. *If the initial promised utility is set to $w_1^1 = \bar{v}\sqrt{8T \log T} + \underline{w}_1$, then the expected profit of the mechanism, denoted by Π, satisfies:*

$$\Pi \ge \Pi^{\mathsf{FB}} - O\left(\sqrt{T \log T}\right)\,,$$

whenever $w_1^1 \le \overline{w}_1^1 - \bar{v}$. In particular, the mechanism is asymptotically optimal, i.e., $\Pi/\Pi^{\mathsf{FB}} \to 1$ as $T \to \infty$.

The proof is omitted as it follows by the same argument as the one in Theorem 3.3. A corollary of the previous theorem is that the mechanism is asymptotically efficient.

Corollary 4.3. *The mechanism is asymptotically efficient, i.e, in expectation, it performs efficient trade in all but in a sublinear number of time periods.*

Finally, we discuss the issue of budget balance. Note that our mechanism is not budget balanced in the sense that $\sum_{i=1}^n z_i^t - z_0^t \ge 0$ for all periods t. In repeated settings, however, the platform can subsidize trade in some periods using proceedings from other periods. Thus, we introduce the notion of aggregated budget balance, which means that $\sum_{t=1}^T \left(\sum_{i=1}^n z_i^t - z_0^t\right) \geq 0$. Although this property is not satisfied almost surely by the mechanism, the probability that the mechanism is not budget balanced in aggregate vanishes as the length of the time horizon goes to infinity. Additionally, because the profit of the platform is given by the sum of all agents' payments, Theorem 4.2 implies that, in expectation, the platform needs to make no subsides. Moreover, it follows from the proof of the theorem that budget balancedness can only be violated near the end of the horizon. Therefore, with high probability the platform can cover the subsidization costs using its proceedings from earlier periods.

Theorem 4.4. *The probability that the aggregate budget balance constraint is violated goes to zero as time goes to infinity:*

$$\mathbb{P}\left[\sum_{t=1}^T \left(\sum_{i=1}^n z_i^t - z_0^t\right) \geq 0\right] \to 1 \text{ as } T \to \infty\,.$$

Proof. Let $\sigma(\boldsymbol{v}^t) = (\max \boldsymbol{v}_{-0}^t - v_0^t)^+$ be the social welfare for period t when the agents' values are $\boldsymbol{v}^t$. Let $\boldsymbol{v} = \left(\boldsymbol{v}^t\right)_{t=1}^T$ be the vector of values for the whole horizon. We denote by $\Pi(\boldsymbol{v})$ and $\Pi^{\mathsf{FB}}(\boldsymbol{v}) = \sum_{t=1}^T \sigma(\boldsymbol{v}^t)$ the sample-path profit of the platform under our mechanism and first-best social welfare, respectively, when the realized vector of values is $\boldsymbol{v}$. We first argue that $\frac{1}{T}\mathbb{E}\left[\left|\Pi^{\mathsf{FB}} - \Pi(\boldsymbol{v})\right|\right]$ converges to zero. We have

$$\begin{aligned}
&\mathbb{E}\left[\left|\Pi^{\mathsf{FB}} - \Pi(\boldsymbol{v})\right|\right] \\
&\leq \mathbb{E}\left[\left|\Pi^{\mathsf{FB}} - \Pi^{\mathsf{FB}}(\boldsymbol{v})\right|\right] + \mathbb{E}\left[\Pi^{\mathsf{FB}}(\boldsymbol{v}) - \Pi(\boldsymbol{v})\right] \\
&\leq \sqrt{\mathrm{Var}\left(\Pi^{\mathsf{FB}}(\boldsymbol{v})\right)} + \Pi^{\mathsf{FB}} - \Pi = \tilde{O}\left(\sqrt{T}\right)\,,
\end{aligned}$$

where the first inequality follows from Minkowski's inequality and using that $\Pi^{\mathsf{FB}}(\boldsymbol{v}) \geq \Pi(\boldsymbol{v})$ almost surely, the second from Jensen's inequality, and the last from Theorem 4.2 and using that $\mathrm{Var}\left(\Pi^{\mathsf{FB}}(\boldsymbol{v})\right) = T\mathrm{Var}\left(\sigma(\boldsymbol{v}^1)\right)$ because values are independent and using that $\mathrm{Var}\left(\sigma(\boldsymbol{v}^1)\right) \leq \bar{v}^2/4$ because $\sigma(\boldsymbol{v}^1) \in [0, \bar{v}]$. We obtain by Markov's inequality that

$$\begin{aligned}
\mathbb{P}\left[\Pi(\boldsymbol{v}) < 0\right] &= \mathbb{P}\left[\Pi^{\mathsf{FB}} - \Pi(\boldsymbol{v}) > \Pi^{\mathsf{FB}}\right] \\
&\leq \mathbb{P}\left[\left|\Pi^{\mathsf{FB}} - \Pi(\boldsymbol{v})\right| \geq \Pi^{\mathsf{FB}}\right] = \tilde{O}\left(\sqrt{T}\right),
\end{aligned}$$

because $\Pi^{\mathsf{FB}} = T\,\mathbb{E}\left[\sigma(\boldsymbol{v}^1)\right]$. □

5 Two-sided markets

We finally consider a setting with multiple buyers and one seller, where we combine the ideas developed in previous sections. The challenge in designing mechanisms for two-sided settings with (NPT) is that the promised utilities must always be non-negative and bounded by the expected surplus of the optimal allocation in the upcoming periods. If the promised utility of a certain agent ever becomes zero, then the mechanism has no alternative other than only allocating to the extreme type of that agent (highest type for buyers and lowest type for sellers), which only allows for a non-trivial allocations in settings for which there is non-zero mass on the extreme type. Assuming there is no positive mass on the extreme type of an agent, the mechanism needs to stop allocating altogether to that agent if his promised utility hits zero. This is particularly problematic if the agent whose promised utility hits zero is the seller, since this would force the mechanism not to allocate to any of the remaining buyers. As a consequence, the main challenge in designing mechanisms for two-sided markets is to prevent the seller's promised utility from hitting zero before all buyers.

The main idea for the bilateral trade setting (Section 4) was to design the mechanism so that the promise utilities of the seller and buyer are perfectly coupled. We extend this idea to the two-sided setting by designing a mechanism that couples the promise utilities of the seller and buyers. In contrast to the bilateral trade setting, coupling the promised utilities is not straightforward in general two-sided markets. We next describe the two challenges we face.

First, to maximize profits, the platform should try to allocate efficiently so as to maximize the gains of trade and charge agents their values (i.e., run a first-price auction) as long as possible. Equation (PI) gives that promised utilities evolve as a random walk with jumps equal to the information rents of each agent. The jumps are $(v_i - \max \boldsymbol{v}_{-i})^+$ for the buyer and $(\max \boldsymbol{v}_{-0} - v_0)^+$ for the seller. When there is only one buyer these jumps are equal and the promised utility processes of the seller and buyer are coupled. With multiple buyers, however, only the winner's promised utility increases and the magnitude of the winner's jump is smaller than the seller's whenever the second-highest buyer value exceeds the seller's value. Thus, the processes are not longer coupled. Recall that in the derivation of (PI) we set the utility for the lowest type to be independent of the realization of values of the other agents. By making the buyers utility for the lowest type a function of the reported types of the other agents, we can obtain a stronger version of (PI) that allows us to perfectly couple the agent's promised utilities (see equation (PI') in the proof of Lemma 5.2).

Second, when the promised utility is either low or high the platform should implement a mechanism that depletes the promised utility of the agents deterministically and at the same rate. Mechanisms such as the VCG auction or simple lotteries are not suitable because they lead to the seller's promised utility depleting at a faster rate than the buyers. We tackle this challenge by implementing a mechanism that offers a posted price to the seller and, whenever the seller trades, allocates the item for free to the buyers using a lottery. We show that by suitably choosing the lottery probabilities and the posted price of the seller we can guarantee that all agents deplete their utilities deterministically and at the same rate. Combining these two ideas, we can perfectly couple the agents promised utilities.

5.1 Mechanism definition

We construct our mechanism under the following assumption.

Assumption 5.1. There exists some reserve price $r \in [0, \bar{v}]$ and probabilities $\boldsymbol{\alpha} \in [0,1]^n$ such that $\alpha_i \mathbb{E}[v_i \mathbf{1}\{v_0 \le r\}] = \mathbb{E}[(r - v_0)^+]$ for all $i \in [n]$, $\sum_{i\in[n]} \alpha_i \le 1$, and $\mathbb{E}[(\max \boldsymbol{v}_{-0} - v_0)^+] \ge \mathbb{E}[(r - v_0)^+] > 0$.

Assumption 5.1 can be shown to hold when then the seller's values are independent of the buyers. To see this, note that

$$\frac{\mathbb{E}[(r - v_0)^+]}{\mathbb{E}[v_i \mathbf{1}\{v_0 \le r\}]} = \frac{1}{\mathbb{E}[v_i]\mathcal{F}_0(r)} \int_0^r \mathcal{F}_0(v) \mathrm{d}v \le \frac{r}{\mathbb{E}[v_i]} \,,$$

where the first equation follows from independence and using integration by parts, and the second because the seller's distribution of values $\mathcal{F}_0(v)$ is non-decreasing. Therefore, by picking a sufficiently small reserve price r, we can always guarantee that Assumption 5.1 holds.

As before we will define ranges of desirable promised utilities for each agent. Let $\underline{w}_0 = \mathbb{E}[(\max \boldsymbol{v}_{-0} - v_0)^+]$ be the seller's expected surplus from efficient trade achieved by the VCG auction. Additionally, let $\overline{w}^t = (T - t + 1)\mu$ where $\mu := \mathbb{E}[(r - v_0)^+]$. Those values define a desirable range $[\underline{w}_0, \overline{w}^t]$ for the promised utility w_i^t. We will set up the mechanism in such a way that we will be able to allocate efficiently whenever w_i^t is in this range and will argue that by setting the initial promise properly, each buyer will remain in the desirable range with high probability for all but a sublinear number of periods. Moreover, the mechanism will guarantee that the promise utility is equal for all agents, i.e., $w_0^t = w_1^t = \ldots = w_n^t$.

Below we describe the mechanism. We will use i to denote an arbitrary buyer and 0 the seller. We break ties lexicographically. We consider the following cases:

Medium promised utility (first-price auction): If $\underline{w}_0 \le w_0^t \le \overline{w}^t - \bar{v}$, the platform executes the trade, the winning buyer pays his value, and the seller is paid his cost.

$$
\begin{aligned}
q_i^t &= \mathbf{1}\{v_i^t \ge \max \boldsymbol{v}_{-i}^t\}\,, \qquad z_i^t = v_i^t q_i^t\,,\\
u_i^t &= w_i^t + \left(\max \boldsymbol{v}_{-0}^t - v_0^t\right)^+ - \underline{w}_0\,,\\
q_0^t &= \mathbf{1}\{v_0^t \le \max \boldsymbol{v}_{-0}^t\}\,, \qquad z_0^t = v_0^t q_0^t\,,\\
u_0^t &= w_0^t + \left(\max \boldsymbol{v}_{-0}^t - v_0^t\right)^+ - \underline{w}_0\,.
\end{aligned}
$$

As in the one-sided case and the bilateral case, the platform is able to completely extract the gains from trade.

Low and high promised utility (inefficient allocation): If $w_0^t \le \underline{w}_0$ or $\overline{w}^t - \bar{v} \le w_0^t$, the platform executes the trade whenever the seller's value is below r, allocates to each buyer with probability α_i, and charges to each agent the lowest type that guarantees winning. Additionally, to guarantee that the promised utility remains positive we throttle the allocation with probability:

$$p_i^t = \min(1, w_i^t/\mu)\,.$$

The mechanism is:

$$
\begin{aligned}
q_i^t &= \alpha_i p_i^t \mathbf{1}\{v_0^t \le r\}\,, && z_i^t = 0\,, && u_i^t = \left(w_i^t - \mu\right)^+\,,\\
q_0^t &= p_i^t \mathbf{1}\{v_0^t \le r\}\,, && z_0^t = r q_0^t\,, && u_0^t = \left(w_0^t - \mu\right)^+\,.
\end{aligned}
$$

5.2 Mechanism analysis

Throughout this section we assume that the initial promise utilities of all agents are the same and equal to w^1, where the initial state is chosen so that $0 \le w^1 \le \overline{w}^1 = \mu T$.

Lemma 5.2. *If the initial promised utility satisfies $0 \le w^1 \le \overline{w}^1$, then $0 \le w_i^t \le \overline{w}^t$ for all t and agent i, and the mechanism described satisfies* (DIC)*,* (PIR)*,* (NPT)*,* (PK)*, and $q_0^t \ge \sum_{i=1}^n q_i^t$. Moreover, promise utilities are coupled, i.e., $w_0^t = w_1^t = \ldots = w_n^t$ for all t.*

When the expected value of an agent is very small relatively to the seller's, the platform needs to set a low reserve price to satisfy Assumption 5.1; thus, limiting the range of attainable promised utilities. We remark, however, that our mechanism is always feasible when the number of time periods is sufficiently large. Alternatively, when the number of time periods is small, dropping buyers with low valuations might have a negligible impact on the platform's objective. We conjecture that Assumption 5.1 can be relaxed by choosing alternative mechanisms for the boundary region. Finally, our mechanism for the boundary region satisfies a weaker version of (FSB) that might require the platform to withhold an item from the buyers. In the context of internet advertising, for example, the platform could simply leave the ad slot empty or show a house ad. We believe it should be possible to construct mechanisms for two-sided markets that satisfy (FSB).

Proof. By construction, it is easy to see that the promise utilities are coupled when the initial state is the same for all agents. We prove by induction on t that $0 \le w_0^t \le \overline{w}^t$ for all t. The result for the other agents follows because promised utilities are coupled. We first argue that $w_0^t \ge 0$. When the state lies in the high or low promised utility region, the claim follows because $u_0^t = \left(w_0^t - \mu\right)^+ \ge 0$. When the state lies in the medium promised utility region, we have $u_0^t \ge w_0^t - \underline{w}_0 \ge 0$ since the state satisfies $w_0^t \ge \underline{w}_0$. We next argue that $w_0^t \le \overline{w}^t = (T-t+1)\mu$. When the state lies in the high or low promised utility region, the claim follows because $u_0^t = \left(w_0^t - \mu\right)^+ \le (\overline{w}^t - \mu)^+ = \overline{w}^{t+1}$ by induction. When the state lies in the medium promised utility region, we have $u_0^t \le w_0^t + \bar{v} - \underline{w}_0 \le \overline{w}^t - \mu = \overline{w}^{t+1}$ since the state satisfies $w_0^t \le \overline{w}^t - \bar{v}$ and $\mu \le \underline{w}_0$ by Assumption 5.1.

The feasibility constraint $q_0^t \ge \sum_{i=1}^n q_i^t$ follows from the fact we allocate to at most one buyer in each period (breaking ties lexicographically) only when we acquire the item from the seller. In the medium region, the feasibility constraint holds with equality. In the high and low regions, the constraint follows because $p_i^t \le 1$ and $\sum_i \alpha_i \le 1$ from Assumption 5.1.

Periodic individual rationality (PIR) is satisfied since we charge each buyer at most his value and pay the seller at least his cost. Since buyer payments are always non-negative, (NPT) is satisfied.

Finally, to check dynamic incentive compatibility (DIC) and promise keeping (PK), it suffices to check that the mechanism satisfies (PI) since the allocations are trivially monotone. We drop the dependence on time to simplify the notation. For the buyers in the low and high region, we have

$$
\begin{aligned}
u_i(\boldsymbol{v}) &= w_i - v_i \cdot q_i(\boldsymbol{v}) + \int_0^{v_i} q_i(v, \boldsymbol{v}_{-i})\mathrm{d}v + z_i(\boldsymbol{v}) \\
&\quad - \mathop{\mathbb{E}}_{\boldsymbol{v}}\left[\int_0^{v_i} q_i(v, \boldsymbol{v}_{-i})\mathrm{d}v\right] \\
&= w_i - p_i\alpha_i \mathop{\mathbb{E}}_{\boldsymbol{v}}\left[v_i \mathbf{1}\{v_0 \le r\}\right] \\
&= w_i - p_i\mu = (w_i - \mu)^+ \,,
\end{aligned}
$$

where the second equation follows from $\int_0^{v_i} q_i(v, \boldsymbol{v}_{-i})\mathrm{d}v = v_i q_i(\boldsymbol{v}) = \alpha_i p_i v_i \mathbf{1}\{v_0 \le r\}$ because the allocation is independent of the buyer's value and $z_i(\boldsymbol{v}) = 0$, the third equation follows from Assumption 5.1 and the definition of μ, and the last because $p_i = \min(1, w_i/\mu)$. For the seller in the low and high region, we have

$$
\begin{aligned}
u_0(\boldsymbol{v}) &= w_0 + v_0 \cdot q_0(\boldsymbol{v}) + \int_{v_0}^{\bar{v}} q_0(v, \boldsymbol{v}_{-0})\mathrm{d}v - z_0(\boldsymbol{v}) \\
&\quad - \mathop{\mathbb{E}}_{\boldsymbol{v}}\left[\int_{v_0}^{\bar{v}} q_0(v, \boldsymbol{v}_{-0})\mathrm{d}v\right] \\
&= w_0 - p_0 \mathop{\mathbb{E}}_{\boldsymbol{v}}\left[(v_0 - r)^+\right] \\
&= w_0 - p_0\mu = (w_0 - \mu)^+ \,,
\end{aligned}
$$

where the second equation follows from $\int_{v_0}^{\bar{v}} q_0(v, \boldsymbol{v}_{-0})\mathrm{d}v = p_0(r - v_0)^+$ and $z_0(\boldsymbol{v}) = rq_0(\boldsymbol{v})$, the third equation follows the definition of μ, and the last because $p_0 = \min(1, w_0/\mu)$. Equation (PI) for the seller in the medium range follows using similar arguments because $\int_{v_0}^{\bar{v}} q_0(v, \boldsymbol{v}_{-0})\mathrm{d}v = \left(\max \boldsymbol{v}_{-0}^t - v_0^t\right)^+$, using our definition of $\underline{w}_0$, and the fact that $z_0(\boldsymbol{v}) - v_0 q_0(\boldsymbol{v}) = 0$ because the (PIR) constraint is binding.

We conclude by showing that dynamic incentive compatibility (DIC) and promise keeping (PK) hold for the buyers in the medium range. In this case, however, (PI) does not hold. In order to guarantee that the promise utilities remain coupled, we design the utility of the lowest type of each buyer $U_i(0, \boldsymbol{v}_{-i})$ to be dependent on the report of the other agents. Recall that the promise keeping constraint and the envelope formula imply that the interim utility of the lowest type should satisfy

$$
\mathop{\mathbb{E}}_{\boldsymbol{v}_{-i}}\left[U_i(0, \boldsymbol{v}_{-i})\right] = w_i - \mathop{\mathbb{E}}_{\boldsymbol{v}}\left[\int_0^{v_i} q_i(v, \boldsymbol{v}_{-i})\mathrm{d}v\right] =: U_i(0) \,.
$$

Because this constraint is at the interim level, for any function $g_i(\boldsymbol{v}_{-i})$ we have that $U_i(0, \boldsymbol{v}_{-i}) =$

$U_i(0) + g_i(\boldsymbol{v}_{-i}) - \mathbb{E}_{\boldsymbol{v}_{-i}}[g_i(\boldsymbol{v}_{-i})]$ satisfies our requirements. We can rewrite the (PI) constraint as follows:

$$
\begin{aligned}
u_i(\boldsymbol{v}) = w_i - v_i q_i(\boldsymbol{v}) &+ \int_0^{v_i} q_i(v, \boldsymbol{v}_{-i}) \mathrm{d}v \qquad \text{(PI')} \\
&+ z_i(\boldsymbol{v}) - \mathop{\mathbb{E}}_{\boldsymbol{v}}\left[\int_0^{v_i} q_i(v, \boldsymbol{v}_{-i}) \mathrm{d}v\right] \\
&+ g_i(\boldsymbol{v}_{-i}) - \mathop{\mathbb{E}}_{\boldsymbol{v}_{-i}}[g_i(\boldsymbol{v}_{-i})] \,.
\end{aligned}
$$

Using that $\int_0^{v_i} q_i(v, \boldsymbol{v}_{-i}) \mathrm{d}v = (v_i - \max \boldsymbol{v}_{-i})^+$ together with the fact that $v_i \cdot q_i(\boldsymbol{v}) - z_i(\boldsymbol{v}) = 0$ because the (PIR) constraint is binding, we obtain that

$$
\begin{aligned}
u_i(\boldsymbol{v}) = w_i + (v_i - \max \boldsymbol{v}_{-i})^+ &- \mathop{\mathbb{E}}_{v}\left[(v_i - \max \boldsymbol{v}_{-i})^+\right] \\
&+ g_i(\boldsymbol{v}_{-i}) - \mathop{\mathbb{E}}_{\boldsymbol{v}_{-i}}[g_i(\boldsymbol{v}_{-i})] \,.
\end{aligned}
$$

Let $g_i(\boldsymbol{v}_{-i}) = (\max_{j \in [n]: j \neq i} v_j - v_0)^+$. Let $v_{(1)}$ denote the highest value of the buyers and $v_{(2)}$ denote the second-highest value of the buyers. Using these expressions and conditioning on whether i is the winner (i.e., $v_i \geq \max \boldsymbol{v}_{-i}$) we obtain that

$$
\begin{aligned}
&(v_i - \max \boldsymbol{v}_{-i})^+ \\
&= \begin{cases} (v_{(1)} - v_0)^+ - (v_{(2)} - v_0)^+ & \text{if } v_i \geq \max \boldsymbol{v}_{-i} \,, \\ 0 & \text{otherwise} \,. \end{cases}
\end{aligned}
$$

and

$$
g_i(\boldsymbol{v}_{-i}) = \begin{cases} (v_{(2)} - v_0)^+ & \text{if } v_i \geq \max \boldsymbol{v}_{-i} \,, \\ (v_{(1)} - v_0)^+ & \text{otherwise} \,. \end{cases}
$$

Combining both expressions we obtain $(v_i - \max \boldsymbol{v}_{-i})^+ + g_i(\boldsymbol{v}_{-i}) = (v_{(1)} - v_0)^+$. This implies that

$$
u_i(\boldsymbol{v}) = w_i + (v_{(1)} - v_0)^+ - \mathbb{E}_{\boldsymbol{v}}\left[(v_{(1)} - v_0)^+\right] \,,
$$

and the result follows because $v_{(1)} = \max \boldsymbol{v}_{-0}$. □

5.3 Profit approximation

In this section, we prove that the constructed mechanism approaches first best as T increases. The proof is similar to the proof of Theorem 3.3 except that the mechanism is not efficient when the promised utility goes above $\overline{w}^t - \overline{v}$. Such a difference introduces another case of efficiency loss, which we need to bound as well.

Theorem 5.3. *If the initial promised utility of all agents is set to* $w^1 = \underline{w}_0 + \overline{v}\sqrt{8T\log T}$*, then the expected profit of the mechanism, denoted by* Π*, satisfies:*

$$\Pi \geq \Pi^{\mathsf{FB}} - O\left(\sqrt{T\log T}\right),$$

whenever $w^1 \leq \overline{w}^1 - \bar{v}\left(1 + \sqrt{8T\log T}\right)$*. In particular, the mechanism is asymptotically optimal, i.e.,* $\Pi/\Pi^{\mathsf{FB}} \to 1$ *as* $T \to \infty$*.*

Proof. The mechanism is feasible by Lemma 5.2 because $0 \leq w^1 \leq \overline{w}^1$. Because the states of all agents are coupled, it is sufficient to consider the seller's promised utility. Let τ be the stopping time measuring the first time in which the seller's promised utility w_0^t falls in the inefficient regions (including both the high and low promised utility regions). In particular, let $\underline{\tau} = \inf\{t \geq 1 : w_0^t \leq \underline{w}_0\}$ be the stopping time of falling into the throttled region and $\overline{\tau} = \inf\{t \geq 1 : w_0^t \geq \overline{w}^t - \overline{v}\}$ be the stopping time of falling into the inefficient allocation region. Then $\tau = \min\{\underline{\tau}, \overline{\tau}\}$.

Let $S = \Pi + \sum_{i=0}^n w_i = \Pi + (n+1)w^1$ be the expected social welfare of the mechanism. Because the mechanism allocates efficiently in the medium promised utility region, by a similar argument as in the proof of Theorem 3.3, we have

$$\Pi^{\mathsf{FB}} - S \leq \overline{v}\,\mathbb{E}_{\boldsymbol{v}}[(T - \tau + 1)^+].$$

Next, we will bound $\mathbb{E}_{\boldsymbol{v}}[(T-\tau+1)^+]$. Let μ^t be defined similarly as in the proof of Theorem 3.3:

$$\mu^t = \mu^{t-1} + \left(\max \boldsymbol{v}_{-0}^t - v_0^t\right)^+ - \mathbb{E}_{\boldsymbol{v}}\left[(\max \boldsymbol{v}_{-0} - v_0)^+\right],$$

with $\mu^1 = w^1$. Therefore, $\{\mu_i^t\}_{t=1}^T$ is a martingale with respect to the natural filtration with increments bounded by $|\mu^t - \mu^{t+1}| \leq 2\overline{v}$. Additionally, we have that μ^t and the seller's promised utility w_0^t coincide until time τ.

Let $T^* = T + 1 - (\overline{v} + \underline{w}_0 + 2\overline{v}\sqrt{8T\log T})/\mu$ and consider a time $1 \leq t \leq T^*$. Note that $T^* \geq 1$

because $w^1 \le \overline{w}^1 - \bar{v}\left(1+\sqrt{8T\log T}\right)$ and $T^* \le T$ because $\mu \le \bar{v}$. In these time periods the high and low promised utility region do not intersect (i.e., $\underline{w}_0 \le \overline{w}^t - \bar{v}$). Because the initial promise utility lies in the medium region, we obtain for any time period $1 \le t \le T^*$ using the maximal version of Azuma's concentration inequality:

$$\Pr[\underline{\tau} \le t] = \Pr\left[\min_{s\le t} \mu^s \le \underline{w}_0\right] \le \exp\left(-\frac{(w^1-\underline{w}_0)^2}{8t\overline{v}^2}\right).$$

Similarly, using that $\overline{w}^t$ is decreasing with time:

$$\begin{aligned}\Pr[\overline{\tau} \le t] &\le \Pr\left[\max_{s\le t} \mu^s \ge \overline{w}^t - \overline{v}\right] \\ &\le \exp\left(-\frac{(\overline{w}^t - \overline{v} - w^1)^2}{8t\overline{v}^2}\right).\end{aligned}$$

Hence we have

$$\begin{aligned}
&\mathop{\mathbb{E}}_{\boldsymbol{v}}[(T-\tau+1)^+] = \sum_{t=1}^{T} \Pr[\tau \le t] \\
&\le \sum_{t=1}^{T^*} \Pr[\tau \le t] + T - T^* \\
&\le T^* \exp\left(-\frac{(w^1-\underline{w}_0)^2}{8T^*\overline{v}^2}\right) \\
&\quad + T^* \exp\left(-\frac{(\overline{w}^{T^*} - \overline{v} - w^1)^2}{8T^*\overline{v}^2}\right) + T - T^* \\
&\le T \exp\left(-\frac{(w^1-\underline{w}_0)^2}{8T\overline{v}^2}\right) \\
&\quad + T \exp\left(-\frac{(\overline{w}^{T^*} - \overline{v} - w^1)^2}{8T\overline{v}^2}\right) + T - T^* \\
&= 2 + T - T^* = 1 + \underline{w}_0/\mu + (2\sqrt{8T\log T}+1)\overline{v}/\mu\,,
\end{aligned}$$

where the first equation follows from summation by parts, the first inequality follows because probabilities are at most one, the second inequality because $\Pr[\tau \le t] \le \Pr[\underline{\tau} \le t] + \Pr[\overline{\tau} \le t]$ from the union bound, the third inequality because $T \le T^*$, the second equality because $w^1 - \underline{w}_0 = \overline{w}^{T^*} - \overline{v} - w^1 = \overline{v}\sqrt{8T\log T}$ from the definitions of T^* and w^1 together with $\overline{w}^{T^*} = \mu(T - T^* + 1)$, and the last equation from our definition of T^*.

Finally, combining with the definition of S, we conclude that

$$\begin{aligned}\Pi &= S - (n+1)w^1 \\ &= \Pi^{\mathsf{FB}} - \overline{v}\,\mathbb{E}_{\boldsymbol{v}}[(T-\tau+1)^+] - (n+1)w^1 \\ &= \Pi^{\mathsf{FB}} - O\left(\sqrt{T\log T}\right).\end{aligned}$$

□

References

Ashlagi, I., Daskalakis, C. & 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.

Athey, S. & Segal, I. (2013), 'An efficient dynamic mechanism', *Econometrica* **81**(6), 2463–2485.

Azevedo, E. M. & Budish, E. (2017), Strategy-proofness in the large, Technical report, National Bureau of Economic Research.

Balseiro, S., Gurkan, H. & Sun, P. (2018), Multi-agent mechanism design without money. forthcoming, *Operations Research*.

Balseiro, S., Lin, M., Mirrokni, V., Paes Leme, R. & Zuo, S. (2017), Dynamic revenue sharing, *in* 'Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA', pp. 2678–2686.

Balseiro, S. R., Mirrokni, V. S. & Leme, R. P. (2017), Dynamic mechanisms with martingale utilities. forthcoming, *Management Science*.

Bergemann, D., Castro, F. & Weintraub, G. (2017), The scope of sequential screening with ex post participation constraints, *in* 'Proceedings of the 2017 ACM Conference on Economics and Computation', ACM, pp. 163–164.

Bergemann, D. & Välimäki, J. (2010), 'The dynamic pivot mechanism', *Econometrica* **78**(2), 771–789.

Blumrosen, L. & Mizrahi, Y. (2016), Approximating gains-from-trade in bilateral trading, *in* 'International Conference on Web and Internet Economics', Springer, pp. 400–413.

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

Colini-Baldeschi, R., Goldberg, P., de Keijzer, B., Leonardi, S. & Turchetta, S. (2017), Fixed price approximability of the optimal gain from trade, *in* 'International Conference on Web and Internet Economics', Springer, pp. 146–160.

Colini-Baldeschi, R., Goldberg, P. W., de Keijzer, B., Leonardi, S., Roughgarden, T. & Turchetta, S. (2017), Approximately efficient two-sided combinatorial auctions, *in* 'Proceedings of the 2017 ACM Conference on Economics and Computation', ACM, pp. 591–608.

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

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

Gomes, R. & Mirrokni, V. (2014), Optimal revenue-sharing double auctions with applications to ad exchanges, *in* '23rd International World Wide Web Conference, WWW '14, Seoul, Republic of Korea, April 7-11, 2014', pp. 19–28.

Gorokh, A., Banerjee, S. & Iyer, K. (2017), From monetary to non-monetary mechanism design via artificial currencies, *in* 'Proceedings of the 2017 ACM Conference on Economics and Computation', ACM, pp. 563–564.

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

Guo, Y. & Hörner, J. (2015), Dynamic mechanisms without money, Technical report, Cowles Foundation for Research in Economics, Yale University.

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

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

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

McAfee, R. P. (1992), 'A dominant strategy double auction', *Journal of economic Theory* **56**(2), 434–450.

McAfee, R. P. (2008), 'The gains from trade under fixed price mechanisms', *Applied Economics Research Bulletin* **1**(1), 1–10.

McAfee, R. & Reny, P. (1992), 'Correlated information and mechanism design', *Econometrica* **60**(2), 395–421.

Mirrokni, V., Paes Leme, R., Ren, R. & Zuo, S. (2018), Dynamic mechanism design in the field, *in* 'Proceedings of the 2018 World Wide Web Conference', WWW '18, International World Wide Web Conferences Steering Committee, pp. 1359–1368.

Mirrokni, V., Paes Leme, R., Tang, P. & Zuo, S. (2016*a*), Dynamic auctions with bank accounts, *in* 'IJCAI', pp. 387–393.

Mirrokni, V., Paes Leme, R., Tang, P. & Zuo, S. (2016*b*), 'Optimal dynamic mechanisms with ex-post IR via bank accounts'.

Mirrokni, V., Paes Leme, R., Tang, P. & Zuo, S. (2018), Non-clairvoyant dynamic mechanism design, *in* 'Proceedings of the 2018 ACM Conference on Economics and Computation', ACM, pp. 169–169.

Mirrokni, V., Paes Leme, R., Tang, P. & Zuo, S. (2019), Optimal dynamic auctions are virtual welfare maximizers, *in* 'AAAI'.

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

Myerson, R. B. & Satterthwaite, M. A. (1983), 'Efficient mechanisms for bilateral trading', *Journal of Economic Theory* **29**(2), 265 – 281.

Niazadeh, R., Yuan, Y. & Kleinberg, R. (2014), Simple and near-optimal mechanisms for market intermediation, *in* 'Web and Internet Economics - WINE', pp. 386–399.

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

Segal-Halevi, E., Hassidim, A. & Aumann, Y. (2016), 'A random-sampling double-auction mechanism'.

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

Thomas, J. & 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. & Maglaras, C. (2002), 'Optimal dynamic auctions for revenue management', *Management Science* **48**(11), 1388–1407.

A Lower bounds

We now argue that no mechanism can achieve an approximation better than $O(\sqrt{T})$ of the first best benchmark. The lower bound holds even for a setting with a single buyer and a seller with constant value, which is a special case of all settings considered in the paper. This result follows from adapting a result of Balseiro, Mirrokni & Leme (2017) for the infinite discounted time model to the finite undiscounted model. The strategy in their paper (which is the one we will adopt here) is to look at a very simple setting where $v_0 = 0$ and there is only a single buyer with a two point distribution. In this very simple settings, the incentive and promise keeping constraints greatly simplify and the problem can be transformed to a stochastic control problem. A common technique in stochastic control is to consider the perfect information relaxation in which the platform is revealed the entire stochastic process of buyer's valuations at the beginning of the horizon. Under this relaxation, the problem reduces to analyzing a reflected random walk. Using standard concentration inequalities, we can bound this process and as a result, obtain a bound for our original problem.

Theorem A.1. *There is a setting where $v_0 = 0$ and there is a single buyer with a two point distribution where $\Pi^{\mathsf{FB}} - \Pi^* \geq \Omega(\sqrt{T})$, where Π^* is the profit of the optimal mechanism satisfying* Opt.

Proof. We consider an instance with a single buyer with valuation $v = v_1$ with probability f_1 and $v = v_2$ with probability $f_2 = 1 - f_1$ and $0 < v_1 < v_2$. We note here that subscripts i will denote different points in the support of the buyer's valuation instead of the buyer's identity.

Using the promised keeping constraint, we can eliminate payments from the objective in the dynamic program (DP) and restate the dynamic program in terms of the social welfare $S^t(w) = w + \Pi^t(w)$. For the case of one buyer, the payment variable can be easily eliminated using the discrete analog of the envelope formula (see equation (PI) in Section 2.1), simplifying the program to:

$$
\begin{aligned}
S^t(w) = \max_{(q_i,u_i)_i} \; & \sum\nolimits_i f_i \left(v_i q(v_i) + S_{t+1}(u(v_i))\right) \\
\text{s.t } & U_1 = w - f_2(v_2 - v_1)q(v_1)\,, \\
& U_2 = w + f_1(v_2 - v_1)q(v_1)\,, \\
& 0 \leq u(v_i) \leq U_i\,, \quad \forall i\,, \\
& 0 \leq q(v_1) \leq q(v_2) \leq 1\,, \quad \forall i\,.
\end{aligned}
$$

with boundary constraints that $S_{T+1}(w) = 0$. The expressions for U_i are obtained by combining the (PK) constraint $w = f_1 U_1 + f_2 U_2$ with $U_2 = U_1 + (v_2 - v_1)q(v_1)$ which comes out of the simplification of the incentive constraint using the envelope formula (PI).

Increasing $q(v_2)$ can only improve the objective function without any impact in the constraints, so we can assume $q(v_2) = 1$. Now, for any fixed allocation $q(v_1)$, the seller would like to make the promised utility as high as possible, since $S^t(w)$ is non-decreasing in w (this follows because the constraint set of the inner optimization problem expands as w increases). Therefore we can set $u(v_i) = U_i$. This implies that in an optimal mechanism, the state evolves according to:

$$w^{t+1} = w^t + \xi^t q^t(v_1) \quad \text{for}$$
$$\xi^t = \begin{cases} -f_2(v_2 - v_1)\,, & \text{w.p. } f_1\,, \\ f_1(v_2 - v_1)\,, & \text{w.p. } f_2\,. \end{cases}$$

This allows us to phrase the seller's problem as the following stochastic control problem:

$$\begin{aligned} S^1(w) = f_2 v_2 T + \max_{q \in \mathcal{Q}} \mathbb{E}\left[\sum_{t=1}^{T} f_1 v_1 q^t(v_1)\right] & \\ \text{s.t } w^{t+1} = w^t + \xi^t q^t(v_1)\,, & \\ w^1 = w\,, & \\ w^t \geq 0\,,\ 0 \leq q^t(v_1) \leq 1\,, & \end{aligned}$$

where $\mathcal{Q}$ denotes the set of all adaptive, non-anticipative policies $q = (q^t(v_1))_{t=1}^T$ that map a history at time t to an allocation $q^t(v_1)$ of the lowest type. Because $\mathbb{E}[\xi_t] = 0$, the latter is equivalent to the problem of controlling the steps of a random walk with an absorbing barrier at zero.

Consider a perfect information relaxation in which the decision maker has access to all realizations of the random variables $\xi = (\xi_t)_{t=1}^T$. Given a sample path $\xi \in \mathbb{R}^T$ 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 $S^1(w)$. More formally, we denote by $S^{\mathsf{HS}}(\xi; w)$ the optimal (deterministic) vale of the perfect information problem for sample path ξ, where H stands for hindsight. We then have $S^1(w) \leq \mathbb{E}\left[S^{\mathsf{HS}}(\xi; w)\right]$, where the perfect information problem is given by

$$
\begin{aligned}
S^{\mathsf{HS}}(\xi; w) = f_2 v_2 T + \max_{(w_t, q_t)_{t=1}^T} & \sum_{t=1}^{T} f_1 v_1 q^t(v_1) \\
\text{s.t } & w^{t+1} = w^t + \xi_t q^t(v_1)\,, \\
& w^1 = w\,, \\
& w^t \geq 0\,,\ 0 \leq q^t(v_1) \leq 1\,.
\end{aligned}
$$

It is not hard to see that the perfect information problem admits a simple optimal solution: greedily set $q^t(v_1)$ as large as possible. This implies that $q^t(v_1) = 1$ whenever $\xi^t > 0$ and $q^t(v_1) = \min\{1, w_t/(f_2(v_2 - v_1))\}$ when $\xi^t < 0$. Additionally, the state evolves according to the reflected random walk

$$
w^{t+1} = \max\{w^t + \xi^t, 0\}\,.
$$

We next provide a closed-form expression for the optimal objective value of the perfect information problem using Skohorod's map for reflected random walks. Using that $q^t(v_1) = 1$ whenever $\xi^t > 0$, we obtain that the optimal objective value is given by

$$
\begin{aligned}
S^{\mathsf{HS}}(\xi; w) = f_2 v_2 T + f_1 v_1 & \sum_{t=1}^{T} \mathbf{1}\{\xi^t > 0\} \\
+ f_1 v_1 & \sum_{t=1}^{T} q^t(v_1) \mathbf{1}\{\xi^t > 0\}\,.
\end{aligned}
$$

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

$$
\begin{aligned}
w^{T+1} - w = & \sum_{t=1}^{T} q^t(v_1)\xi^t \\
& = f_1(v_2 - v_1) \sum_{t=1}^{T} \mathbf{1}\{\xi^t > 0\} \\
& \quad - f_2(v_2 - v_1) \sum_{t=1}^{T} q^t(v_1) \mathbf{1}\{\xi^t < 0\}\,.
\end{aligned}
$$

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

$$S^{\mathsf{HS}}(\xi;w) = f_2 v_2 T + \frac{f_1 v_1}{f_2} \sum_{t=1}^{T} \mathbf{1}\{\xi^t > 0\} \\ - \rho \left(w^{T+1} - w\right).$$

Let $X_t = -\sum_{s=1}^{t} \xi^s$ be the state of the random walk by time t. Skohorod's map implies that the reflected random walk satisfies $w^{t+1} = w - X_t + \max_{1\le s\le t}(X_s - w, 0)^+$. Taking expectations and using that $\mathbb{E}[\xi^t] = 0$ we obtain that

$$\mathbb{E}\Big[S^{\mathsf{HS}}(\xi;w)\Big] = \Pi^{\mathsf{FB}} - \rho\,\mathbb{E}\Big[\max_{1\le s\le T}(X_s - w)^+\Big], \tag{A.1}$$

where we also used that $\mathbb{P}\{\xi_t > 0\} = f_2$ and $\Pi^{\mathsf{FB}} = T(f_1 v_1 + f_2 v_2)$. The second term can be further bounded from below as follows

$$\begin{aligned} \mathbb{E}\left[\max_{1\le s\le T}(X_s - w)^+\right] &\ge \mathbb{E}\left[(X_T - w)^+\right] \\ &\ge \left(\mathbb{E}\left[X_T^+\right] - w\right)^+, \end{aligned}$$

where the second inequality follows because $(X_T - w)^+ = (X_T^+ - w)^+$ because $w \ge 0$ together with Jensen's inequality. Because the random variable X_T is mean zero, we obtain that $\mathbb{E}\left[X_T^+\right] = \mathbb{E}\left|X_T\right|/2$. Let $Q_t = \sum_{s=1}^{t} \xi_s^2$ be the quadratic variation process. Marcinkiewicz-Zygmund inequality implies that there exists a constant $c_1 > 0$ such that $\mathbb{E}\big|X_t\big| \ge c_1 \mathbb{E}\big|Q_t^{1/2}\big|$. Because the random steps are almost surely lower bounded by $|\xi_s| \ge (v_2 - v_1)\min(f_1, f_2)$, we have that the quadratic variation process is lower bounded by $Q_t^{1/2} \ge (v_2 - v_1)\min(f_1, f_2)t^{1/2}$. Therefore, there exists some constant $c_2 > 0$ such that $\mathbb{E}\left[X_T^+\right] \ge c_2 T^{1/2}$.

Finally, we use the bound on $S^1(w)$ to bound Π^* as follows:

$$\begin{aligned} \Pi^{\mathsf{FB}} - \Pi^* &= \max_{w\ge 0}\left[\Pi^{\mathsf{FB}} + w - S^1(w)\right] \\ &\ge \max_{w\ge 0}\left[\Pi^{\mathsf{FB}} + w - \mathbb{E}\left[S^{\mathsf{HS}}(\xi;w)\right]\right] \\ &\ge \max_{w\ge 0}\left[w + \rho(\mathbb{E}[X_T^+] - w)^+\right] \\ &= \min(\rho, 1)\cdot\mathbb{E}\left[X_T^+\right] \ge \Omega(\sqrt{T}). \end{aligned}$$

□

B Implementation without knowing $\bar{v}$

Both the mechanism description and the analysis depend on the quantity $\bar{v}$ which is the highest value in the support of the valuation distribution. The choice of restricting to bounded support distribution is mostly to simplify the presentation and the analysis. It is possible to define and analyze a similar mechanisms that do not require $\bar{v}$. Instead of having two clearly marked first- and second-price regions, the alternative mechanism will "interpolate" between a first-price and second-price auction depending on $\overline{w}_i^t - w_i^t$. This mechanism is identical to the one described earlier in this section for $w_i^t \leq \underline{w}_i$ and for the losing buyers. For the winning buyer i if his promised utility w_i^t is in the range $[\underline{w}_i, \overline{w}_i^t]$, we allocate efficiently and update the promises as follows:

$$q_i^t = 1\,, \quad z_i^t = v_{-i} + \min(v_i - \max \boldsymbol{v}_{-i}, \overline{w}_i^t - w_i^t)\,,$$
$$u_i^t = w_i - \underline{w}_i + \min(v_i - \max \boldsymbol{v}_{-i}, \overline{w}_i^t - w_i^t)\,.$$

Note that when $w_i = \overline{w}_i^t$ the mechanism runs a second-price auction and for any $w_i < \overline{w}_i^t$ the mechanism interpolates between a first-price and second-price where the interpolation is given by the cap $\overline{w}_i^t - w_i^t$ on the gap $v_i - \max \boldsymbol{v}_{-i}$ that we are allowed to consume. In particular, when the promise utility satisfies $w_i^t \leq \overline{w}_i^t - \bar{v}$ the mechanism runs a first-price auction as before.