CROSSCUTTING AREAS

Dynamic Mechanism Design with Budget-Constrained Buyers Under Limited Commitment

Santiago R. Balseiro,^a Omar Besbes,^a Gabriel Y. Weintraub^b

^a Graduate School of Business, Columbia University, New York, New York 10027; ^b Graduate School of Business, Stanford University, Stanford, California 94305

Received: June 14, 2016 **Abstract.** We study the dynamic mechanism design problem of a seller who repeatedly Revised: May 7, 2018; August 28, 2018 auctions independent items over a discrete time horizon to buyers who face a cumulative Accepted: November 1, 2018 budget constraint. A driving motivation behind our model is the emergence of real-time Published Online in Articles in Advance: bidding markets for online display advertising in which such budgets are prevalent. We May 1, 2019 assume the seller has a strong form of limited commitment: she commits to the rules of the current auction but cannot commit to those of future auctions. We show that the cele-Subject Classifications: bidding/auctions: brated Myersonian approach that leverages the envelope theorem fails in this setting, and games/group decisions; stochastic model therefore, characterizing the dynamic optimal mechanism seems intractable. Despite these applications: probability challenges, we derive and characterize a near-optimal dynamic mechanism. To do so, we Area of Review: Games and Networks show that the Myersonian approach is recovered in a corresponding *fluid* continuous time https://doi.org/10.1287/opre.2018.1830 model in which the time interval between consecutive items becomes negligible. Then we leverage this approach to characterize the optimal dynamic direct-revelation mechanism, Copyright: © 2019 INFORMS highlighting novel incentives at play in settings with buyers' budget constraints and seller's limited commitment. We show through a combination of theoretical and numerical results that the optimal mechanism arising from the fluid continuous time model approximately satisfies incentive compatibility for the buyers and is approximately sequentially rational for the seller in the original discrete time model.

Funding: S. R. Balseiro's research was partially supported by a Google Faculty Research Award. Supplemental Material: Supplemental material is available at https://doi.org/10.1287/opre.2018.1830.

Keywords: dynamic mechanism design • limited commitment • budget constraints • fluid approximation • display advertising • internet auctions • revenue management

1. Introduction

Consider the setting of a seller auctioning off a sequence of independent and perishable items to heterogeneous buyers facing *budget constraints* that limit their overall expenditure over a given time horizon. The presence of budgets couples the auctions buyers participate in because paying for an item today reduces a buyer's option value of waiting for potentially better future opportunities. Therefore, offering a sequence of one-shot static optimal auctions is not necessarily optimal, and the question of how to sell in such a dynamic environment arises. Thus motivated, in this paper, we study the problem of *dynamic mechanism design* in the presence of budget-constrained buyers.

An important assumption when studying dynamic mechanisms is the extent of commitment power endowed to the seller. In a prototypical setting of independent private valuations over time, when the seller can commit to a dynamic mechanism up front, she can typically achieve full surplus extraction because there is no information asymmetry at the time of contracting. For example, when budgets are abundant and do not bind, the seller can run efficient auctions and extract the entire surplus by charging appropriate participation fees at the beginning of the horizon. Similar ideas can be attempted when budgets are stringent. In contrast, in this paper we assume the seller has a strong form of limited commitment: she commits to the rules of the current auction, but she cannot commit to those of future auctions. Note that in this case, full surplus extraction is not possible anymore, and if budgets are abundant, the optimal dynamic mechanism consists of running a sequence of Myerson optimal static auctions (Myerson 1981). When budgets are stringent, however, the structure of an optimal mechanism is unknown. The presence of budgets together with the lack of commitment results in a novel and challenging dynamic mechanism design problem. Our work extends classic work in mechanism design by incorporating aggregate budget constraints.

Our model formulation is also motivated by practice, in particular, by online display advertising. In these markets, publishers sell impressions to advertisers using real-time auctions (see e.g., Muthukrishnan 2009, Korula et al. 2015). For their part, advertisers execute marketing *campaigns* based on a predetermined budget that extends for a fixed amount of time over which advertisers participate in a large number of such auctions. The budget constraints are typically exogenously imposed, at least in the short run. Finally, although these markets originated using second-price auctions, they have typically evolved to become much less transparent with respect to the auction rules. In particular, publishers are not always committed to the rules of the future auctions they will run (Yuan et al. 2013). Sequential auctions are also common in flower, wine, and art markets (see for example, Ashenfelter 1989, van den Berg 2001). In the context of this type of sequential auctions it may also be reasonable to expect that each buyer faces a budget constraint that limits her overall spending (Pitchik 2009). For example, art museums have special yearly funds destined for the acquisition of works of art (Dobrzynski 2012).

We now describe our model and main contributions in more detail. We consider the problem of a risk-neutral seller auctioning items arriving sequentially over a given time horizon. The seller has a cost for each item she sells. Given the time horizons involved in the main applications we have in mind (weeks to a few months), the seller does not discount future payoffs. There is a set of risk-neutral buyers, each one endowed with a positive budget, which constrains her total expenditure over all the auctions she participates in. We assume buyers have private values that are independent across items and buyers. At every period, buyers learn their current valuations but are uncertain about their value for future items.

We have purposely designed our model to be the simplest possible extension of a classic setting to be able to crisply highlight the specific effect that budgets have on an optimal mechanism. In particular, we have assumed the standard independent private valuation model. We also assume that budgets and their evolutions are common knowledge. As we discuss in Section 1.1, assuming private budgets introduces significant challenges even in a one-shot static auction. Further, we will see that even our "simple" dynamic model is already a very rich model that will present significant challenges in the analysis while yielding novel insights.

We provide a mechanism design formulation of the problem in which the seller aims to maximize profits subject to dynamic Bayesian incentive compatibility, dynamic Bayesian individual rationality, and budget feasibility. The seller is sequentially rational and cannot commit not to reoptimize her choice in future auctions. Further, we naturally restrict attention to dynamic mechanisms that are *Markov* with respect to the state of remaining budgets. In our setting, payments impact remaining budgets, and the dynamic incentive compatibility constraints incorporate the option value of future opportunities via continuation values, which are typically nonlinear functions of the current state. As a consequence, the celebrated Myersonian approach of using the envelope theorem to express expected interim payments as a linear functional of the allocation fails. Therefore, characterizing the dynamic optimal mechanism for our base model seems intractable.

Our main contribution lies in overcoming these challenges. In particular, we derive and characterize an easily computable near-optimal mechanism for this class of problems in the sense that the mechanism approximately satisfies incentive compatibility for the buyers and is approximately sequentially rational for the seller. To achieve this, a key building block is the introduction of a *fluid continuous time* model characterized by a set of Hamilton-Jacobi-Bellman (HJB) equations that can be interpreted as representing a limit in which the time interval between consecutive items offered becomes negligible (and values, payments, and costs are appropriately normalized). Notably, we show that in this regime one recovers the Myersonian approach. Intuitively, when the number of items is large, the payment in one auction is small relative to the budget.¹ Hence, a firstorder Taylor expansion around the current state is a good approximation, and the continuation values become approximately linear. We show that this approximation is exact in the fluid continuous time model. In turn, payments can be expressed as a linear functional of the allocation, recovering the main tool underlying the Myerson approach to characterize the structure of an optimal mechanism. The solution is more subtle than the static mechanism case, though, because it naturally includes value function derivatives that capture the budgets' future opportunity costs.

The continuous time fluid model yields a concrete prescription for the original discrete time model. We develop a numerical approach to test its performance and show through extensive experiments that the mechanism is indeed near-optimal. Furthermore, we establish a theoretical link between the fluid continuous time model and the original discrete time counterpart in the case of one buyer. First, we develop an approach that leverages state-of-the-art results in partial differential equations (PDEs) to show that such a system admits a solution, and therefore, an optimal fluid mechanism always exists. Second, we mathematically prove that this prescription indeed becomes a nearoptimal best response in the discrete time model as the number of auctions increases.

The fluid model, beyond offering a concrete prescription, also leads to insights on the structure of nearoptimal mechanisms in such dynamic environments. In the mechanism, a given object is allocated to the bidder with the highest "modified virtual value," provided that it is larger than the seller's cost. A modified virtual value is equal to the classic virtual value times an "allocation factor"; these factors capture the dynamics introduced by budgets. To illustrate, in the case of a single buyer, the optimal mechanism is a *twotier auction* that allocates the item whenever the report is greater than or equal to a *threshold value* and charges the buyer a *payment*, which is lower than the threshold value. The threshold and the payment change dynamically over time to balance the desire of the seller to extract the budget with as few items as possible with the threat of the buyer to not participate.

In summary, to the best of our knowledge, our paper is the first in the literature that (1) studies a new class of problems relevant to both theory and practice dynamic mechanism design when selling a sequence of items to buyers who face a cumulative budget constraint—and (2) shows the potential of a fluid continuous time model and how it recovers the ability to use an envelope approach and allows to obtain a near-optimal prescription in the original discrete model.

The rest of the paper is organized as follows. Section 1.1 positions the paper in the literature. Section 2 describes the discrete time stochastic model and solution concept and formalizes the dynamic mechanism design problem under limited commitment. Section 3 introduces the fluid continuous time model, and Section 4 provides a sharp characterization of an optimal solution. Section 5 connects the fluid and discrete models and establishes through a combination of numerics and theory the near-optimality of the fluid model prescription in the discrete model. Section 6 concludes with some final remarks. Owing to space considerations, all proofs are presented in the supplemental material.

1.1. Related Literature

In this section we discuss the connection of our work to several streams of literature. First, our paper naturally relates to problems of dynamic mechanism design. Pavan et al. (2014) and Kakade et al. (2013) provide necessary and sufficient conditions for optimal mechanisms in large classes of environments in which agents' information changes over time. However, our model differs from theirs in two important dimensions. First, they assume the designer can commit to a dynamic mechanism up front. Second, in our model the current state of bidders is determined by payments, whereas in theirs it is determined by allocations. Lewis and Yildirim (2002) derive an optimal mechanism in a setting with learning by doing with two suppliers when the auctioneer cannot commit to future auctions. Although it also focuses on the limited commitment case, in their setting the state is again determined by the

allocation and not by the payments. Under these dynamics, an envelope approach can be directly applied in their stylized discrete time model.

A related stream of work including Vulcano et al. (2002), Gallien (2006), Gershkov and Moldovanu (2014), and Board and Skrzypacz (2015) studies dynamic pricing and revenue management problems using dynamic mechanism design. In these models the designer sells multiple items over a finite horizon after which the items perish, introducing the option value of waiting for better future opportunities for the seller. In contrast, in our case, an item can only be used in the current period, and the option value of waiting arises at the buyer level. Akan et al. (2015) use a mechanism design approach to characterize a firm's optimal screening strategy when consumers learn their valuations for future consumption over time. This line of work assumes that the seller has full commitment power. Deb and Said (2015) study a model in which consumers arrive over two time periods, and the firm cannot commit in advance to the contractual terms it offers in the second period. In dynamic pricing settings, the impact of the seller's inability to commit when facing strategic buyers who optimally time their purchases has been studied in Hörner and Samuelson (2011) and Dilme and Li (2019). Limited commitment also plays a critical role in the Coase conjecture in which a monopolist selling durable goods competes with future incarnations of herself (see, e.g., Gul et al. 1986). Our model differs in that goods sold are independent and perishable instead of durable, and buyers are budget constrained instead of unit demand.

Skreta (2006, 2015) study revenue-maximizing mechanisms in which the seller cannot commit not to propose a new auction if the object fails to sell in the current auction. However, different from our work, in these papers the seller auctions a *single* object, and many of the difficulties in the analysis arise because the revelation principle cannot be applied directly owing to a "ratchet" effect. In this same line of work, our paper is more related to that by Liu et al. (2019), which also leverages a continuous time limit to characterize the optimal mechanism (in their model, the seller's commitment power vanishes in this limit). Similar to them, we do not attempt to formalize the extensive form nor the equilibrium concept in the limit model because there are unresolved conceptual issues with these definitions in games for which actions are updated continuously (see, e.g., Simon and Stinchcombe 1989, Bergin and MacLeod 1993, Alós-Ferrer and Ritzberger 2008).

There is a stream of papers that study auction design when selling a *single* item with financially constrained bidders (see, e.g., Laffont and Robert 1996, Che and Gale 1998, Maskin 2000, Che and Gale 2000). In particular, Pai and Vohra (2014) characterize the optimal auction for selling one item in the presence of multiple budget-constrained buyers using linear programming. The optimal mechanism when values and budgets are both private is quite complex, involving several pooling regions in which the highest bidder may not win outright. Brusco and Lopomo (2008) study the bidding equilibrium in two-unit simultaneous ascending-bid auctions when each bidder has private information about her willingness to pay and her budget. On its part, Borgs et al. (2005) and Bhattacharya et al. (2010) study approximation algorithms for designing revenue-optimal incentivecompatible mechanisms in multiunit auctions with budget-constrained bidders. These papers, however, focus on static one-shot settings, whereas here we focus on a financial constraint over a number of sequential auctions. Finally, Benoît and Krishna (2001) consider a model with two objects with common values in a complete information setting and study the revenue performance of different standard auctions.

Another related stream of work pertains to the study of dynamic interactions of bidders under specific auction mechanisms using approximate notions of equilibria. Iyer et al. (2014) study such interactions when bidders learn about their valuations over time. Balseiro et al. (2015) and Gummadi et al. (2012) study budgetconstrained bidders in repeated second-price auctions. None of these papers, however, study optimal mechanism design. Finally, other studies, such as Nazerzadeh et al. (2013), have also considered notions of asymptotic incentive compatibility.

2. Model

2.1. Seller and Buyers

We consider the problem of a seller who has a number of perishable items arriving sequentially over a given time horizon [0, T]. We first analyze a discrete model in which there are N perishable items arriving at times $\delta, 2\delta, \ldots, N\delta = T$ to sell sequentially. This model is the natural extension of the single-shot auction to multiperiod auctions with budget-constrained buyers. To naturally scale the problem with the number of items, we assume that values, payments, and costs are scaled by the time period length δ .² Given that the only time points at which decisions are made are discrete, we index time periods backward by $n = N, \ldots, 1$.

The seller is risk neutral, does not discount future payoffs, and has a strictly positive $\cot c > 0$ for each item she sells.³ At every period of time, the seller runs an auction to allocate the item and determine payments among a set of buyers. The seller wishes to maximize profits from selling the items over the horizon to the buyers.

There are *K* risk-neutral buyers, who are budget constrained and remain in the market for the whole

time horizon. We index the buyers by k = 1, ..., K. Buyer *k* is endowed with a positive budget B_k , which constrains her total expenditure over the horizon. The presence of budget constraints introduces intertemporal substitution because buyers should ponder the option value of future opportunities in their decisions. At every period of time, buyers have independent private valuations for the item for sale. Specifically, a buyer's valuation for the item is drawn independently (across bidders and time periods) and at random from a cumulative distribution function $F_k(\cdot)$ with density $f_k(\cdot)$ and domain being an interval in \mathbb{R}_+ ⁴ We denote $\overline{F}_k(v) = 1 - F_k(v)$. At every period, buyers learn their current valuations but are uncertain about their values for future items. Throughout this paper, we make the standard assumption in the mechanism design literature that for each buyer k, the virtual valuation function $\phi_k(v) = v - F_k(v)/f_k(v)$ is increasing in v. Examples of distributions satisfying this condition are uniform, exponential, truncated normal, and Weibull, among others. Buyers have a quasi-linear utility function given by the difference between the sum of the values of items won minus the payments made to the seller during the length of the horizon. We summarize the buyer's characteristics by a vector $(B_k, F_k(\cdot))$.

The number of opportunities N and the cost c are assumed to be common knowledge. The buyers' vectors of characteristics are also assumed to be common knowledge, and so are the budgets' evolutions.

2.2. Timing of Events

The timing of the events is as follows. Initially, buyers arrive to the market at the beginning of the time horizon, and their characteristics, together with the length of the horizon, are publicly disclosed. Then the following steps are sequentially repeated as a new item becomes available for sale. First, buyers learn their valuations for the current item, and then the seller offers a mechanism to sell the current item. In turn, buyers submit a report to the seller. The seller publicly announces the outcome of the mechanism for that item and the payments buyers should make to the seller. Third, the item is transferred to the winner (if any), payments are made, and budgets are updated. We refer to the game associated with the sale of an individual item as the stage game. We assume that the entire history of auction outcomes is publicly observed.

2.3. Stage Game

We restrict attention to direct revelation mechanisms, which we then argue is done without loss of optimality by the revelation principle. At each time period, the seller chooses a *stage mechanism* $\mathbf{m} = (\mathbf{p}, \mathbf{z}) \in \mathcal{M}$, where we denote by \mathcal{M} the space of stage mechanisms

(we use boldface for vectors and vector functions). Buyers are asked to report their valuations, and the space of reports $\mathbb{R}_{\perp} \triangleq \mathbb{R}_{+} \cup \{\perp\}$ includes the element \perp , which represents the case in which the buyer opts not to participate in the mechanism. The mechanism for the stage game is characterized by the pair of functions (**p**, **z**), where $\mathbf{p} : \mathbb{R}^{K}_{\perp} \to \Delta$ is a probability allocation function from the space of reports to the probability simplex $\Delta \triangleq \{\mathbf{q} \in \mathbb{R}_+^K : \sum_{k=1}^K q_k \leq 1\}$ (and hence at most one good is allocated), $\mathbf{z} : \mathbb{R}_{\perp}^{K} \to \mathbb{R}^{K}$ is the payment function from the space of reports to the space of possible payments.⁵ To sum up, given reports $\mathbf{w} \in \mathbb{R}^{K}_{+}$, the function $p_{k}(\mathbf{w})$ determines the probability that the item is assigned to buyer *k*, whereas the function $z_k(\mathbf{w})$ determines the payment buyer k must make to the seller. We assume that when $w_k = \bot$ (i.e., buyer *k* does not participate in the stage mechanism), buyer *k* never gets allocated the item and in return pays nothing; in this case, $p_k(\mathbf{w}) = z_k(\mathbf{w}) = 0$.

Further, we assume that if one buyer does not participate in an auction, then transfers are not allowed at all to any bidder; that is, if $w_k = \bot$, then $\mathbf{z}_{-k}(\mathbf{w}) = 0$, where $\mathbf{z}_{-k}(\cdot)$ denotes the payments to bidders other than k. Later we will consider mechanisms that are incentive compatible and individually rational, hence, in the optimal path all buyers will participate and bid their own values truthfully. The previous assumption, however, limits threats (outside the path being played) that the seller can impose via payments to/from competitors to incentivize a given buyer to participate. We believe these threats are unreasonable and unrealistic. Our assumption eliminates them in a simple way, allowing us to focus on our main purpose, which is the applicability of the fluid formulation.⁶

2.4. Dynamic Game with Limited Commitment and Revelation Principle

We assume that the seller is unable to commit to future proposed mechanisms; she can only commit to the mechanism in the current period. We impose sequential rationality to reflect the seller's inability to commit. We also restrict attention to Markov strategies with respect to the natural state given by the vector of remaining budgets $\mathbf{x} \in \mathbb{R}^{K}_{+}$ and the number of items remaining for sale in the horizon $n \in \mathbb{N}$. The restriction to these types of Markov strategies has been used in related papers because of its behavioral appeal and its simplicity (see, e.g., Lewis and Yildirim 2002, Pavan et al. 2014).

If we allow for general mechanisms for the seller and strategies for the buyers and model the dynamic relation between the seller and buyers as a dynamic game, an optimal mechanism corresponds to a *Markov perfect equilibrium* of such a game. By the revelation principle, the payoffs for the seller in any Markov perfect equilibrium can be achieved by a direct revelation mechanism. Hence, the direct revelation optimal mechanism remains optimal among a larger class of mechanisms.⁷

2.5. Mechanism Design Problem and Solution Concept

A Markov strategy for the seller is a *dynamic mechanism* that specifies a stage mechanism for every possible market state $(\mathbf{x}, n) \in \mathcal{F}^{\delta} \subset \mathbb{R}_{+}^{K} \times \mathbb{N}$, where the set of feasible states is $\mathcal{F}^{\delta} \triangleq \prod_{k=1}^{K} [0, B_k] \times \{1, \dots, N\}$. We denote such a dynamic mechanism as $\mathbf{M} : \mathcal{F}^{\delta} \to \mathcal{M}$, where $\mathbf{M}[\mathbf{x}, n] = \mathbf{m} \in \mathcal{M}$ is the stage mechanism for market state (\mathbf{x}, n) . With some abuse of notation, we denote by $\mathbf{Z}[\mathbf{x}, n] = \mathbf{z}$ the payment function and by $\mathbf{P}[\mathbf{x}, n] = \mathbf{p}$ the allocation function for market state (\mathbf{x}, n) . We denote by \mathbb{M}^{δ} the set of all dynamic mechanisms.

Given a dynamic mechanism $\mathbf{M} \in \mathbb{M}^{\delta}$ and an initial state $(\mathbf{x}, n) \in \mathcal{G}^{\delta}$, when all buyers report their values truthfully, budgets evolve according to the stochastic process $\{\mathbf{x}_i\}_{i=1}^n$ with dynamics $\mathbf{x}_{i-1} = \mathbf{x}_i - \delta \mathbf{Z}[\mathbf{x}_i, i](\mathbf{v}_i)$ and initial state $\mathbf{x}_n = \mathbf{x}$, where $\mathbf{Z}[\mathbf{x}_i, i](\mathbf{v}_i)$ is the vector of payments in market state (\mathbf{x}_i, i) and value realizations $\mathbf{v}_i = (v_{i,k})_{k=1}^K$. We denote by $\mathbb{E}_{\mathbf{x},n}^M[-]$ the expectation with respect to this process and with respect to present and future buyers' valuations.

We denote the total seller's profit from an initial state (\mathbf{x}, n) until the end of the horizon, when all buyers report truthfully by

$$\Pi^{\mathbf{M}}(\mathbf{x}, n) = \mathbb{E}_{\mathbf{x}, n}^{\mathbf{M}} \left[\delta \sum_{i=1}^{n} \sum_{k=1}^{K} (Z_{k}[\mathbf{x}_{i}, i](\mathbf{v}_{i}) - cP_{k}[\mathbf{x}_{i}, i](\mathbf{v}_{i})) \right]$$
$$= \mathbb{E}_{\mathbf{x}, n}^{\mathbf{M}} \left[\sum_{i=1}^{n} \delta \pi(\mathbf{M}[\mathbf{x}_{i}, i]) \right],$$

where the expectation is taken with respect to the evolution of budgets when buyers bid truthfully and the dynamic mechanism **M** is implemented, and $\pi : \mathcal{M} \to \mathbb{R}$ is the expected seller's profit functional for a one-stage mechanism under truthful reporting. That is, when the seller selects a stage mechanism $\mathbf{m} \in \mathcal{M}$ and all buyers report their value truthfully, the seller's profit per period is given by $\pi(\mathbf{m}) = \sum_{k=1}^{K} \mathbb{E}_{\mathbf{v}} \cdot [z_k(\mathbf{v}) - cp_k(\mathbf{v})]$, where the expectation is taken with respect to the buyers random values \mathbf{v} . Therefore, the second equation follows from the fact that values are independent between themselves and independent of the mechanism for the stage. Similarly, we denote the total utility of buyer *k* from an initial state (\mathbf{x}, n) when all buyers report truthfully by

$$\begin{aligned} U_k^{\mathbf{M}}(\mathbf{x}, n) &= \mathbb{E}_{\mathbf{x}, n}^{\mathbf{M}} \bigg[\delta \sum_{i=1}^n (v_{i,k} P_k[\mathbf{x}_i, i](\mathbf{v}_i) - Z_k[\mathbf{x}_i, i](\mathbf{v}_i)) \bigg] \\ &= \mathbb{E}_{\mathbf{x}, n}^{\mathbf{M}} \bigg[\sum_{i=1}^n \delta u_k(\mathbf{M}[\mathbf{x}_i, i]) \bigg], \end{aligned}$$

where we denote by $u_k : \mathcal{M} \to \mathbb{R}$ the expected buyer's utility functional. That is, when the seller selects a stage mechanism $\mathbf{m} \in \mathcal{M}$, the expected utility per period for buyer *k* when all buyers report their values truthfully is given by $u_k(\mathbf{m}) = \mathbb{E}_{\mathbf{v}}[v_k p_k(\mathbf{v}) - z_k(\mathbf{v})]$.

2.6. Backward Induction Characterization

Using the principle of optimality of dynamic programming, we provide sufficient conditions that characterize a *direct revelation dynamic optimal mechanism* $\mathbf{M}^* \in \mathbb{M}^{\delta}$. To simplify notation, we will drop the dependence on the optimal dynamic mechanism \mathbf{M}^* , and with some abuse of notation, we denote by $\Pi(\mathbf{x}, n) \triangleq \Pi^{\mathbf{M}^*}(\mathbf{x}, n)$ the expected cumulative profit-to-go of the seller from state $(\mathbf{x}, n) \in \mathcal{F}^{\delta}$ under the optimal mechanism \mathbf{M}^* when all buyers report their values truthfully. Similarly, we denote by $U_k(\mathbf{x}, n) \triangleq U_k^{\mathbf{M}^*}(\mathbf{x}, n)$ the expected cumulative utility-to-go for buyer *k*.

For all states (\mathbf{x}, n) , the optimal stage mechanism $\mathbf{M}^*[\mathbf{x}, n]$ should be *budget feasible* and *dynamic incentive compatible* (DIC). Budget feasibility requires that budget constraints are satisfied for every buyer and every report. DIC requires that buyers are better off participating in the mechanism and reporting their true values in the current stage given the continuation values for future opportunities (and assuming that all buyers report truthfully onward). More formally, the optimal stage mechanism should satisfy, for all $(\mathbf{x}, n) \in \mathcal{P}^{\delta}$,

$$\mathbf{M}^{*}[\mathbf{x}, n] \in \arg \max_{\mathbf{m} \in \mathcal{M}} \delta \pi(\mathbf{m}) + \mathbb{E}_{\mathbf{v}} \Big[\Pi(\mathbf{x} - \delta \mathbf{z}(\mathbf{v}), n - 1) \Big]$$
(1a)
s.t. $\delta z_{k}(\mathbf{w}) \leq x_{k}, \quad \forall k, \mathbf{w} \in \mathbb{R}_{\perp}^{K},$ (1b)

$$v \in \arg \max_{w \in \mathbb{R}_{\perp}} \delta u_k(\mathbf{m}, v, w) + \mathbb{E}_{\mathbf{v}_{-k}} \Big[U_k(\mathbf{x} - \delta \mathbf{z}(w, \mathbf{v}_{-k}), n - 1) \Big] \forall v, k,$$
(1c)

where $\mathbb{E}_{\mathbf{v}_{-k}}[-]$ denotes the expectation with respect to the values of buyers different from *k*, and with some abuse of notation, we define the interim buyer utility as

$$u_k(\mathbf{m}, v, w) = \mathbb{E}_{\mathbf{v}_{-k}}[vp_k(w, \mathbf{v}_{-k}) - z_k(w, \mathbf{v}_{-k})].$$

The objective gives the profit-to-go of the seller under the mechanism, constraint (1b) imposes budget feasibility, and constraint (1c) imposes dynamic incentive compatibility. The DIC constraint implies that the mechanism is *individually rational* in the sense that every buyer has a (weak) incentive to participate in each auction because reporting her value truthfully dominates reporting \perp . Additionally, by the one-stage-deviation principle for sequential rationality (Hendon et al. 1996), DIC implies that bidding truthfully in every item is a buyer's best response over all dynamic bidding strategies when all other buyers report truthfully.

Additionally, the seller's optimal value function satisfies the following recursive equation:

$$\Pi(\mathbf{x}, n) = \delta \pi(\mathbf{M}^*[\mathbf{x}, n]) + \mathbb{E}_{\mathbf{v}} \Big[\Pi \Big(\mathbf{x} - \delta \mathbf{Z}^*[\mathbf{x}, n](\mathbf{v}), n - 1 \Big) \Big], \quad (2)$$

with boundary conditions $\Pi(\mathbf{x}, 0) = 0$ for all budgets $\mathbf{x} \in \mathbb{R}_+^K$ and $\Pi(\mathbf{0}, n) = 0$ for all $n \ge 1$. Similarly, the buyers' value function under the optimal dynamic mechanism satisfies the equation

$$U_{k}(\mathbf{x}, n) = \delta u_{k}(\mathbf{M}^{*}[\mathbf{x}, n]) + \mathbb{E}_{\mathbf{v}} \Big[U_{k}(\mathbf{x} - \delta \mathbf{Z}^{*}[\mathbf{x}, n](\mathbf{v}), n-1) \Big].$$
(3)

We assume that buyers are excluded after they run out of budget. In turn, the boundary conditions are given by $U_k(\mathbf{x}, 0) = 0$, for all $\mathbf{x} \in \mathbb{R}_+^K$, and $U_k(\mathbf{x}, n) = 0$, for all $n \ge 1$, and $\mathbf{x} \in \mathbb{R}_+^K$ with $x_k = 0$.

2.7. Envelope Condition in Dynamic Setting

In the static mechanism design problem, the celebrated optimal auction result of Myerson (1981) exploits the fact that expected interim payments can be expressed in closed form as a linear functional of the allocation; in particular, this follows from the envelope theorem applied to the incentive compatibility constraints. As a result, in that setting, one may eliminate payments from the problem, and the objective becomes a linear functional over the allocation, which can be easily solved by optimizing point-wise over values. Milgrom (2004) provides an extensive discussion of the centrality of envelope theorems in mechanism design for characterizing incentive compatibility, deriving revenue equivalences and ultimately characterizing optimal mechanisms. Pavan et al. (2014) also emphasize how an envelope condition plays a crucial role in their analysis of dynamic mechanisms with commitment.

2.7.1. Envelope Condition When the Payments Impact the State. We next explain how this approach fails in the present setting and highlight where it does so. With some abuse of notation, we denote by $\widetilde{U}_k(v, w)$ the interim utility-to-go of buyer k in state (\mathbf{x}, n) when she has observed a value of v, reports w to the seller, the stage mechanism is $\mathbf{m} \in \mathcal{M}$, all other buyers report truthfully, and the dynamic mechanism \mathbf{M}^* is used thereafter under truthful reporting. The latter satisfies the following equation:

$$\begin{split} \widetilde{U}_k(v,w) &= \mathbb{E}_{\mathbf{v}_{-k}} \Big[\delta v \ p_k(w,\mathbf{v}_{-k}) - \delta z_k(w,\mathbf{v}_{-k}) \\ &+ U_k(\mathbf{x} - \delta \mathbf{z}(w,\mathbf{v}_{-k}), n-1) \Big], \end{split}$$

where the argument of the value function on the righthand side denotes the budget evolution. The DIC conditions imply that for the given state the stage mechanism $\mathbf{m} \in \mathcal{M}$ should satisfy $\widetilde{U}_k(v, v) = \max_{w \ge 0} \widetilde{U}_k(v, w)$. Using the envelope theorem, we obtain

$$\frac{\mathrm{d}\widetilde{U}_k(v,v)}{\mathrm{d}v} = \delta \mathbb{E}_{\mathbf{v}_{-k}}[p_k(v,\mathbf{v}_{-k})],$$

and integrating, one obtains

$$\widetilde{U}_k(v,v) = \widetilde{U}_k(0,0) + \delta \int_0^v \mathbb{E}_{\mathbf{v}_{-k}}[p_k(v,\mathbf{v}_{-k})] \mathrm{d}v.$$

Using our definition of the utility-to-go, we get that the envelope condition corresponding to the DIC constraint is given by

$$\mathbb{E}_{\mathbf{v}_{-k}} \left[v p_k(v, \mathbf{v}_{-k}) - z_k(v, \mathbf{v}_{-k}) + \delta^{-1} U_k(\mathbf{x} - \delta \mathbf{z}(v, \mathbf{v}_{-k}), n-1) \right]$$

= $\mathbb{E}_{\mathbf{v}_{-k}} \left[-z_k(0, \mathbf{v}_{-k}) + \delta^{-1} U_k(\mathbf{x} - \delta \mathbf{z}(0, \mathbf{v}_{-k}), n-1) \right]$
+ $\int_0^v \mathbb{E}_{\mathbf{v}_{-k}} \left[p_k(v, \mathbf{v}_{-k}) \right] dv.$ (4)

Note that in a static mechanism design problem, there are no continuation values, and the envelope condition can be used to derive a linear equation of interim expected payments on allocations. As mentioned previously, Myerson (1981) uses this equation to write the optimal mechanism design problem as a linear program on allocations only.

However, in the present case, the continuation value $U_k(\mathbf{x}, n)$ is typically nonlinear in the budget \mathbf{x} , and hence one cannot solve for the payments as in the static mechanism design problem, and the Myerson approach fails. In turn, it seems intractable to derive the form that an optimal mechanism would take (or any insights on it) and how such a mechanism could be implemented.

3. A Tractable Fluid Model

3.1. Toward a Tractable Formulation

As seen earlier, the key difficulty stems from the fact that the payments affect the state, and the continuation values may be nonlinear. We next provide a heuristic argument showing that this issue disappears in a corresponding fluid continuous time model. In particular, the Myerson approach based on the envelope condition can be applied again in such a setting. This will motivate the formal introduction of such a model in Sections 3.2 and 3.3. We will then leverage this fluid approach to characterize the structure of an optimal mechanism.

If continuation values were linear, then one could again solve for the payments as a linear functional of the allocation and as a result characterize the dynamic optimal mechanism. Although continuation values are nonlinear in general, intuitively, when the number of items is large and the gains of trade per auction are small relative to the budget (δ is small), then the payment in one auction is small relative to the budget, and the continuation values are approximately linear in the region of achievable states from the current state. We next provide a heuristic derivation of the envelope condition in this case. Performing a first-order expansion of the discrete model buyer's value function around the current state for buyer *k*, we obtain

$$U_k(\mathbf{x} - \delta \mathbf{z}(\mathbf{w}), n-1) \approx U_k(\mathbf{x}, n-1) - \nabla_{\mathbf{x}} U_k(\mathbf{x}, n-1) \\ \cdot \delta \mathbf{z}(\mathbf{w}).$$

Applying this approximation to the envelope condition (4) and canceling terms, we get that

$$\mathbb{E}_{\mathbf{v}_{-k}} [vp_k(v, \mathbf{v}_{-k}) - z_k(v, \mathbf{v}_{-k}) - \nabla_{\mathbf{x}} U_k(\mathbf{x}, n-1) \cdot \mathbf{z}(v, \mathbf{v}_{-k})]$$

$$\approx \mathbb{E}_{\mathbf{v}_{-k}} [-z_k(0, \mathbf{v}_{-k}) - \nabla_{\mathbf{x}} U_k(\mathbf{x}, n-1) \cdot \mathbf{z}(0, \mathbf{v}_{-k})]$$

$$+ \int_0^v \mathbb{E}_{\mathbf{v}_{-k}} [p_k(v, \mathbf{v}_k)] dv.$$
(5)

Comparing the latter with the corresponding envelope condition in the discrete model (4), we see how the central difficulty disappears when the number of items is large: in Equation (5), the resulting equation is linear in the payments $\mathbf{z}(\cdot)$, and now one can solve a system of linear equations to obtain the payments as a function of allocations and replace it in the seller's objective. With this heuristic argument, one recovers the ability to use the main tool underlying the Myerson approach to characterize the structure of an optimal mechanism. One should note that in the present setting, however, the system obtained is more complex than in the one-shot mechanism design case because terms associated with the gradients of the value function are present. This is to be expected, though; these terms capture the dynamics of the problem, in particular, the budgets' future opportunity costs.

We next introduce a *fluid* model in which an infinite number of infinitesimal items are available to sell. In such a formulation, the approximation in Equation (5) becomes exact. We will show how to leverage the fluid model to derive a near-optimal mechanism for the original problem (1).

3.2. Fluid Model

In this section, we consider a continuous time model in which bidders spend an infinitesimal amount of their budgets in each auction. Here we have a *fluid* of items arriving sequentially, and budgets are depleted deterministically. We refer to this as a *continuous time fluid* mechanism design problem. Heuristically, the fluid model can be interpreted as the limit of the discrete model as $\delta \downarrow 0$ (in which the number of items scales as T/δ). Budgets are not scaled because values, payments, and costs were assumed to scale with the time period length δ .

We now consider the problem of a seller who has items arriving continuously at unit rate to sell sequentially to Kpossible budget-constrained buyers. The seller receives the items one at a time, and the sequence of events is assumed to be identical to the one in the discrete time setting. We consider a continuous time formulation over a finite time horizon of length T. Time is indexed backward by $t \in [0, T]$. The seller wishes to maximize revenues from the sequential sale of items over the horizon to the buyers. We denote the budget of buyer k in the fluid model by B_k .

It is worth noting that *we do not aim to formally define the extensive form game nor the equilibrium concept* because these concepts remain generally elusive in games for which actions are updated continuously, as in our continuous time model (see, e.g., Bergin and MacLeod 1993, Alós-Ferrer and Ritzberger 2008). In our model, even defining the extensive form game remains elusive because to do so, the ordinary differential equation that describes the budgets' state evolution must admit a unique solution for every allowable strategy profile. This is hard to guarantee except if strong assumptions over the set of allowed mechanisms and budget evolutions are imposed, such as Lipschitz continuity, which are not clear to hold a priori (see, e.g., Başar and Olsder 1999).

Instead, we take as a starting point the HJB equations of the continuous time fluid model that we informally derive. We will define a *dynamic optimal mechanism* as a solution arising from these equations. Such a solution can be thought of as a prescription to be used in the discrete time model. In fact, we will show in Section 5 that this prescription becomes near optimal and near DIC in the discrete time model as the number of periods increases.⁸

3.3. Our Solution Concept: System of HJB Equations

As before, a Markov strategy for the seller is a *dynamic mechanism* that specifies a stage mechanism for every possible market state $(\mathbf{x}, t) \in \mathcal{G} \triangleq \prod_{k=1}^{K} [0, B_k] \times \mathbb{R}_+$. We denote such a dynamic mechanism as $\mathbf{M} : \mathcal{G} \to \mathcal{M}$, where $\mathbf{M}[\mathbf{x}, t] = \mathbf{m} \in \mathcal{M}$ is the stage mechanism for market state (\mathbf{x}, t) . We denote by \mathbb{M} the set of all dynamic mechanisms. In the fluid model, the evolution of budgets is dictated by a system of ordinary differential equations (ODEs) instead of a discrete time stochastic process. Given a dynamic mechanism $\mathbf{M} \in \mathbb{M}$ and an initial state $(\mathbf{x}, t) \in \mathcal{G}$, the budgets $\bar{\mathbf{x}} : [0, t] \to \mathbb{R}_+^K$ evolve deterministically according to the system

of ODEs $\bar{x}_k(s) = e_k(\mathbf{M}[\bar{\mathbf{x}}(s), s])$ with initial condition $\bar{\mathbf{x}}(t) = \mathbf{x}$. In the latter, we denote by $e_k : \mathcal{M} \to \mathbb{R}$ the expected buyer's expenditure rate functional under truthful reporting; that is, $e_k(\mathbf{m}) = \mathbb{E}_{\mathbf{v}}[z_k(\mathbf{v})]$.

Motivated by the argument developed in Section 2.7, we define a *direct revelation dynamic optimal mechanism* for the continuous time fluid model as the solution of a system of coupled HJB equations. We provide an informal derivation of the HJB equations in Appendix B of the supplemental material. Let $\overline{\Pi} : \mathcal{G} \to \mathbb{R}$ be the cumulative profit-to-go of the seller from state $(\mathbf{x}, t) \in \mathcal{G}$ under the optimal mechanism \mathbf{M}^* when all buyers report their values truthfully. Similarly, let $\overline{U}_k(\mathbf{x}, t) : \mathcal{G} \to \mathbb{R}$ be the expected cumulative utility-to-go for buyer *k*.

Suppose that the value functions $\Pi(\mathbf{x}, t)$ and $U_k(\mathbf{x}, t)$ are differentiable everywhere. For all states (\mathbf{x}, t) , the optimal stage mechanism $\mathbf{M}^*[\mathbf{x}, t]$ should be budget feasible and DIC. More formally, the optimal stage mechanism should satisfy for all $(\mathbf{x}, t) \in \mathcal{G}$ (these are the conditions analogous to (1a)–(1c) in the original discrete time formulation)

$$\mathbf{M}^{*}[\mathbf{x}, t] \in \arg \max_{\mathbf{m} \in \mathcal{M}} \pi(\mathbf{m}) - \nabla_{\mathbf{x}} \overline{\Pi}(\mathbf{x}, t) \cdot \mathbf{e}(\mathbf{m})$$
(6a)
s.t. $z_{k}(\mathbf{w}) = 0$, $\forall \mathbf{w} \in \mathbb{R}_{\perp}^{K}$ if $x_{k} = 0$,
(6b)
 $v \in \arg \max_{w \in \mathbb{R}_{\perp}} u_{k}(\mathbf{m}, v, w) - \nabla_{\mathbf{x}} \overline{U}_{k}(\mathbf{x}, t)$
 $\cdot \mathbb{E}_{\mathbf{v}_{-k}}[\mathbf{z}(w, \mathbf{v}_{-k})] \quad \forall v, k,$
(6c)

where the objective (6a) imposes sequential rationality for the seller through the corresponding HJB equation, constraint (6b) imposes budget feasibility, and constraint (6c) imposes dynamic incentive compatibility through the corresponding HJB equations.

Finally, the seller's value function should satisfy the PDE (which is the fluid continuous time counterpart of Equation (2); see Appendix B of the supplemental material)

$$\frac{\partial \bar{\Pi}}{\partial t}(\mathbf{x},t) = \pi(\mathbf{M}^*[\mathbf{x},t]) - \nabla_{\mathbf{x}} \bar{\Pi}(\mathbf{x},t) \cdot \mathbf{e}(\mathbf{M}^*[\mathbf{x},t]), \quad (7)$$

with boundary conditions $\overline{\Pi}(\mathbf{x}, 0) = 0$ for all budgets $\mathbf{x} \in \mathbb{R}_+^K$ and $\overline{\Pi}(\mathbf{0}, t) = 0$ for all $t \ge 0$. Similarly, the buyers' value function satisfies the PDE (which is the fluid continuous time counterpart of Equation (3); see Appendix B of the supplemental material)

$$\frac{\partial U_k}{\partial t}(\mathbf{x},t) = u_k(\mathbf{M}^*[\mathbf{x},t]) - \nabla_{\mathbf{x}} \bar{U}_k(\mathbf{x},t) \cdot \mathbf{e}(\mathbf{M}^*[\mathbf{x},t]), \quad (8)$$

with $\overline{U}_k(\mathbf{x}, 0) = 0$ for all budgets $\mathbf{x} \in \mathbb{R}_+^K$ and $\overline{U}_k(\mathbf{x}, t) = 0$ for all $t \ge 0$ and $\mathbf{x} \in \mathbb{R}_+^K$ with $x_k = 0$.

4. Solution of HJB Equations and Economic Insights

In this section, we first characterize the solution of the system of coupled HJB equations and via its structure provide economic insights regarding the dynamic optimal mechanism in the continuous time fluid model. Then we provide further insights using numerical experiments in the case of multiple buyers and additional analytical results in the case of a single-buyer.

4.1. Characterization of HJB Equations Solution

First, we characterize the solution of the coupled HJB equations under the assumption that the value functions are well behaved.

Definition 1. The value functions $\Pi, \overline{\mathbf{U}} : \mathcal{G} \to \mathbb{R}^{K}$ are said to be well behaved if for every state (\mathbf{x}, t) with t > 0 (i) they are differentiable for the buyers with positive budget at the state $\mathcal{H} = \{k = 1, ..., K : x_k > 0\}$, (ii) the Jacobian matrix $D\overline{\mathbf{U}}|_{\mathcal{H}}$ satisfies that $\mathbf{I} + D\overline{\mathbf{U}}|_{\mathcal{H}}$ is a nonsingular matrix, and (iii) $(1 - \nabla_x \overline{\Pi}|_{\mathcal{H}})(\mathbf{I} + D\overline{\mathbf{U}}|_{\mathcal{H}})^{-1} \ge 0.9$

We later present numerical evidence that these conditions are typically satisfied in instances of interest. We have the following result.

Theorem 1 (HJB Equations Solution Structure). Let $\eta = (\mathbf{1} - \nabla_{\mathbf{x}} \Pi|_{\mathcal{X}})(\mathbf{I} + D \overline{\mathbf{U}}|_{\mathcal{X}})^{-1}$. Additionally, suppose that the value functions $\Pi, \overline{\mathbf{U}} : \mathcal{G} \to \mathbb{R}$ are well behaved and satisfy

$$\frac{\partial \bar{\Pi}}{\partial t} = \mathbb{E}_{\mathbf{v}} \bigg[\bigg(\max_{k \in \mathcal{H}} \big\{ \eta_k \phi_k(v_k) - c \big\} \bigg)^+ \bigg], \tag{9a}$$

$$\frac{\partial \bar{U}_k}{\partial t} = \mathbb{E}_{\mathbf{v}} \left[\mathbf{1} \left\{ \eta_k \phi_k(v_k) \ge \max \left\{ c, \max_{i \in \mathcal{H} \setminus k} \eta_i \phi_i(v_i) \right\} \right\} \\
\cdot \frac{\bar{F}_k(v_k)}{f_k(v_k)} \right], \quad \text{for } k \in \mathcal{H},$$
(9b)

and $\overline{U}_k(\mathbf{x},t) = 0$ for $k \notin \mathcal{K}$, with boundary conditions as before. Let $\mathbf{M}^*[\mathbf{x},t] = (\mathbf{P}^*[\mathbf{x},t], \mathbf{Z}^*[\mathbf{x},t])$ be given by

$$\begin{aligned} P_k^*[\mathbf{x},t](\mathbf{v}) &= \mathbf{1}\{v_k > y_k(\mathbf{v}_{-k})\}, \quad k = 1, \dots, K, \\ Z_k^*[\mathbf{x},t](\mathbf{v}) &= \sum_{i \in \mathcal{H}} a_{ki} P_i^*[\mathbf{x},t](\mathbf{v}) y_i(\mathbf{v}_{-i}), \quad k = 1, \dots, K, \end{aligned}$$

where we denote by $y_k(\mathbf{v}_{-k}) = \inf \{ v : \eta_k \phi_k(v) > \max_{i \in \mathcal{H} \setminus k}, \eta_i \phi_i(v_i), \eta_k \phi_k(v) > c \}$ the smallest value for the kth buyer that wins against reports \mathbf{v}_{-k} from the competitors, and $\mathbf{A} = (a_{ki}) = (\mathbf{I} + D\bar{\mathbf{U}}|_{\mathcal{H}})^{-1}.$

Then the functions Π , $\overline{\mathbf{U}}$, together with dynamic mechanism \mathbf{M}^* , satisfy HJB equations (6)–(8).

All proofs are available in the supplemental material. We briefly describe the main steps of the proof of Theorem 1. In this result, we leverage the continuous

time fluid formulation to adapt the Myerson approach (Myerson 1981), which involves solving for payments as a linear functional of the allocation and then reformulating the seller's optimization problem. In particular, building on the intuition outlined in Section 2.7, we show that in the fluid model the envelope condition yields an equation on payments and allocations that is linear in the former (see Equation (C-10) in the supplemental material). Using this expression to replace the seller's objective function results in an objective that, for given value functions' derivatives, only depends linearly on the allocation and the payment of the lowest type (see Equation (C-12) in the supplemental material). This objective function can then be optimized point-wise to characterize the optimal allocation rule. In turn, we first show that when the value functions satisfy the conditions in the statement, the mechanism in the statement is optimal at any given state; that is, it is an optimal solution of problem (6). We then show that evaluating HJB Equations (7) and (8) at this optimal mechanism yields Equations (9a) and (9b) in the statement. Thus, the value functions and mechanism in the statement satisfy HJB Equations (6)–(8).

To highlight in a simpler setting the economic insights, it is helpful to specialize the preceding result to the case of a single buyer, in which case we obtain a sharper characterization in quasi-closed form. In addition, we later mathematically prove that in the case of a single buyer, the value functions are indeed well behaved. In the following, we denote by $y^+ = \max(y, 0)$ the positive part of a number $y \in \mathbb{R}$.

Corollary 1 (HJB Equations Solution Structure for Single Buyer). Suppose that there is a single buyer. Let $\eta(x,t) = (1 - \overline{\Pi}_x(x,t))/(1 + \overline{U}_x(x,t))$. Suppose that there exists value functions $\overline{\Pi}, \overline{U} : \mathbb{R}_+ \times [0,T] \to \mathbb{R}$ that are well behaved (differentiable and satisfying $0 \le \overline{\Pi}_x(x,t) \le 1$, $\overline{U}_x(x,t) \ge 0$ for x > 0 and t > 0) and satisfy

$$\bar{\Pi}_t(x,t) = \mathbb{E}_v[(\eta(x,t)\phi(v) - c)^+],$$
(10a)

$$\bar{U}_t(x,t) = \mathbb{E}_v \Big[\mathbf{1} \big\{ \eta(x,t)\phi(v) \ge c \big\} \bar{F}(v) / f(v) \Big],$$
(10b)

for all x > 0 and t > 0 with boundary conditions $\overline{\Pi}(x, 0) = \overline{U}(x, 0) = 0$ for all $x \ge 0$ and $\overline{\Pi}(0, t) = \overline{U}(0, t) = 0$ for all $t \ge 0$. Let $M^*[x, t] = (P^*[x, t], Z^*[x, t])$ be given by

$$P^{*}[x,t](v) = \mathbf{1}\{v > r(x,t)\},\$$

$$Z^{*}[x,t](v) = \frac{r(x,t)}{1 + \bar{U}_{x}(x,t)}\mathbf{1}\{v > r(x,t)\},\$$

where the threshold value is given by $r(x,t) = \phi^{-1} \cdot (c/\eta(x,t))$. Then the functions $\overline{\Pi}$, \overline{U} , together with dynamic mechanism M^* , satisfy HJB equations (6)–(8).

First, we discuss the single-buyer result. Note that when the budget is appropriately large, the derivatives of the value functions with respect to the budget are zero, and $\eta(x,t) = 1$. Thus, the auctions are effectively decoupled, and the dynamic optimal mechanism implements the static Myerson optimal auction with reserve $\phi^{-1}(c)$ for every item. Now, when the budget is stringent, dynamics play an important role in the optimal mechanism, as characterized by the corollary. The optimal mechanism is a *two-tier auction* that allocates the item whenever the report is greater than or equal to a threshold value of $r = \phi^{-1}(c/\eta(x,t))$, where $\eta(x,t) = (1 - \Pi_x(x,t))/(1 + U_x(x,t))$ is an allocation factor, and charges the buyer a payment of $r(x,t)/(1 + U_x(x,t))$ whenever the item is allocated. Notably, the buyer's payment is lower than the threshold value in the optimal dynamic mechanism. Additionally, note that both the threshold and payment are state dependent.

To better understand the intuition underlying the two-tier structure, consider the case of a stringent budget, for which, if the seller would implement the static Myerson optimal auction for every item, the buyer would run out of budget before the end of the horizon. In this case, the seller could implement the Myerson auction until the buyer would deplete her budget. One can show that the seller would have an incentive to unilaterally deviate from such a mechanism. In particular, the seller has an incentive to increase the reserve price so as to deplete the budget of the buyer by using as few items as possible. However, if the payment when the item is allocated is too high, the buyer can always decide, for example, to wait and not participate until the seller applies the Myerson auction close to the end of the horizon (given the seller's lack of commitment power). The seller, however, does not have an incentive to wait to implement Myerson later and would prefer to trade early.¹⁰ The optimal dynamic mechanism in Corollary 1 aims to balance the desire of the seller to extract the budget with as few items as possible with the threat of the buyer to not participate. In particular, the two-tier mechanism exposes the buyer to all items, but the buyer only acquires those that are very valuable for her (above the threshold value). The payment is smaller than the threshold, so the buyer has an incentive to report truthfully.

We note that the mechanism in Corollary 1 can also be implemented via a second-price auction with dynamically adjusted reserve (or, alternatively, a dynamic posted price mechanism). Here the seller sets the posted price to $r(x, t)/(1 + \bar{U}_x(x, t))$, and, in turn, the buyer will shade bids to $v/(1 + \bar{U}_x(x, t))$, where v is her value. Thus, the buyer wins the item when her value satisfies v > r(x, t) and pays the posted price $r(x, t)/(1 + \bar{U}_x(x, t))$ in the case of winning, which results in the same allocation and payment as before.

Theorem 1 provides a generalization for multiple buyers. Similar to the earlier discussion, when budgets are ample for all bidders and do not play a role, the derivatives of the value functions are zero and $\eta_k = 1$ for every buyer *k*. In this case, as expected, the Myerson optimal auction is implemented for every item. In general, when budgets might be stringent, under the optimal mechanism, a given object is allocated to the bidder with the highest *modified virtual value* $\eta_k \phi_k(v_k)$, when the latter is larger than the seller's cost *c*. The *allocation factors* η_k capture the dynamics introduced by budgets. Our numerical experiments suggest that typically buyers with lower budgets are assigned lower allocation factors, leading their bids to be ranked relatively lower. This prevents a buyer with a low budget from winning too often and depleting her budget before the end of the horizon, thus maintaining competition throughout the horizon.

In terms of payments, when no value is above the buyer-specific threshold value $\phi_k^{-1}(c/\eta_k)$, the item is not allocated, and no transfers are made. When buyer *k* is the winner, she pays $a_{kk}y_k(\mathbf{v}_{-k})$, where $y_k(\mathbf{v}_{-k})$ is the minimum report that wins against reports \mathbf{v}_{-k} from the competitors. The payment factors a_{kk} are typically less than or equal to one and nondecreasing with time. A similar intuition as that discussed for the single-buyer case applies: they provide a discount to the buyers to incentivize them to report truthfully and participate in early auctions. Additionally, every losing buyer $i \neq k$ pays a fraction of the winner payment given by $a_{ik}y_k(\mathbf{v}_{-k})$. These cross terms capture the externality on losing buyers' utilities associated with the winning buyer's budget decreasing after paying for the item. For example, the optimal mechanism takes advantage of the fact that a losing buyer might be better off when the winning buyer's budget decreases and charges this losing buyer a payment proportional to the winner's payment ($a_{ik} > 0$). Conversely, when a losing buyer is worse off when the winning buyer's budget decreases, the mechanism needs to provide monetary incentives for this losing buyer to bid truthfully and participate in the auction ($a_{ik} < 0$). We will later show via numerical experiments that for multiple bidders, both cases can arise. More broadly, in the next subsection we provide further results about the structure of the optimal mechanism. First, we provide an existence result for the case of a single buyer.

4.1.1. Existence of a Solution. We now show that Equation (10), which characterizes the fluid optimal dynamic mechanism, admits a solution in the case of a single buyer. We introduce the following assumption.

Assumption 1. *The following hold:*

i. The distribution of values has compact support $[0, \overline{v}]$, where $\overline{v} > c$.

ii. The distribution of values has a continuously differentiable density f(v) such that f(v) > 0 for all $v \in [0, \bar{v}]$. iii. The inverse virtual valuation $\psi(y) = \phi^{-1}(y)$ is twice differentiable with Lipschitz derivatives. The derivative satisfies $\psi'(y) > 0$ for all $y \ge 0$.

iv. The inverse virtual valuation satisfies $\psi'(y)(1 + \psi'(y)) + (y - \psi(y))\psi''(y) > 0$ for all $y \ge 0$.

The first two assumptions are technical and are imposed to simplify the analysis. The third assumption imposes that the inverse virtual valuation function is twice continuously differentiable and strictly increasing. Assuming that the inverse virtual valuation is increasing is equivalent to the typical assumption in the mechanism design literature that the virtual valuation is increasing. The fourth assumption controls the curvature of the inverse virtual valuation function. This condition is a common assumption in the literature on PDEs that we use to prove the next existence theorem (see Evans, 2010, p. 626). A sufficient condition for this assumption to hold is that the inverse virtual valuation function is concave. These assumptions are not very restrictive and are satisfied by distributions such as uniform, truncated exponential, truncated normal, truncated Pareto, and truncated Weibull, among others.

The next result shows that an optimal mechanism always exists in the case of a single buyer.

Theorem 2 (Existence). Suppose that Assumption 1 holds Then, the system of coupled PDEs in the statement of Corollary 1 admits a well-behaved solution. In particular, the solution satisfies $0 \le \overline{\Pi}_x(x,t) \le 1$ and $\overline{U}_x(x,t) \ge 0$ for all x > 0 and t > 0 and the HJB equations (6)–(8).

It is important to note that the existence of solutions (let alone characterization of such solutions) of nonlinear systems of PDEs is notably hard to show in general (see, e.g., Evans 2010). Theorem 2 shows that the system of coupled HJB Equations (6)–(8) admits a solution by transforming the system of PDEs into a hyperbolic system of conservation laws and exploiting the theory of rarefaction waves. The proof is constructive and provides an efficient procedure to determine the solution by solving a system of ODEs. The extension of existence to multiple buyers seems at this stage to be beyond the state of the art in the analysis of systems of PDEs. As in the case of a single buyer, with multiple buyers, the PDEs can also be transformed to a hyperbolic system of conservation laws, but little is known about leveraging this for existence because the state is multidimensional (Bressan and Shen 2004). However, in our numerical experiments with multiple buyers, the computational procedure we used always found a solution to the system of PDEs, which is reassuring.

4.2. Further Economic Insights and Numerical Experiments

In this section, we provide further economic insights derived from the solution of the HJB equations. We start with the case of a single buyer for which we are able to provide analytical results. Then we provide numerical results for the general case.

4.2.1. Single-Buyer Problem. With some abuse of terminology, we refer to the *equilibrium path* to the actual evolution under the direct revelation dynamic optimal mechanism. We study the behavior of the direct revelation dynamic optimal mechanism along the equilibrium path. Let $\gamma(x,t) \triangleq \overline{\Pi}_x(x,t)$ and $\mu(x,t) \triangleq \overline{U}_x(x,t)$ be the seller's and buyer's marginal payoff for an additional unit of buyer's budget, respectively. We refer to these quantities as the *shadow prices*. The expenditure rate at state (x,t) under the optimal mechanism is given by $e(x,t) \triangleq \mathbb{E}_v[Z^*[x,t](v)] = \frac{r(x,t)}{1+\mu(x,t)}\overline{F}(r(x,t))$, and the budget along the equilibrium path evolves according to

$$r(x,t)/(1+\mu(x,t))\overline{F}(r(x,t)),$$
 (11)

with the initial conditions $\bar{x}(T) = B$. Lemma C.3 in the supplemental material shows that this ODE always admits a unique solution. In addition, in Appendix C.3 of the supplemental material we show that the seller's shadow price $\gamma(t) \triangleq \gamma(\bar{x}(t), t)$, the buyer's shadow price $\mu(t) \triangleq \mu(\bar{x}(t), t)$, the threshold value $r(t) \triangleq r(\bar{x}(t), t)$, and the payment $r(t)/(1 + \mu(t))$ along the equilibrium path are nondecreasing in t; that is, they are nonincreasing as time progresses toward the end of the horizon (recall that t represents time to go). Further, we show that the threshold value and payment along the equilibrium path satisfy $r(t) \geq \frac{r(t)}{1+\mu(t)} \geq \psi(c)$; that is, the threshold value and the payment are never lower than the Myerson optimal reserve price $\psi(c)$ throughout the time horizon.

Figure 1(a) illustrates a typical equilibrium path when budget is abundant, whereas Figure 1(b) illustrates a typical equilibrium path when budget is stringent. We observe that when the budget is stringent, the threshold value and payments are dynamic over the time horizon and that the gap between the two can be substantial. The seller cannot charge too high a payment to ensure that the buyer is willing to participate in the auctions from the start of the horizon.

4.2.1.1. Budget Depletion on the Equilibrium Path. The following result shows that there are two distinct regimes for the optimal mechanism. When the initial budget-to-time ratio is above a certain threshold, the optimal dynamic mechanism is equivalent to a repeated application of the static Myerson optimal auction without budget constraints. When the initial budget-to-time ratio is below the same threshold, the shadow prices $\mu(t)$ and $\gamma(t)$ are strictly positive until some time $t_0 > 0$, and from then on the seller implements the Myerson auction.



Figure 1. (Color online) Equilibrium Paths Under the Fluid Optimal Mechanism M^* in the Case of a Single Buyer

Notes. Values are U[0, 1], the seller's cost is c = 0.25, and the horizon length is T = 10. In the left figures, the initial budget is abundant $(B/T \ge \rho_0)$, and the seller offers the static optimal auction with reserve price $\psi(c)$. In the right figures, the initial budget is stringent ($B/T < \rho_0$), and the seller offers a dynamic two-tier auction until time t_{0} and the static optimal auction with reserve price $\psi(c)$ from then on.

Proposition 1. Let $\rho_0 \triangleq (\psi(c) + c\dot{\psi}(c))\bar{F}(\psi(c))$ be the *threshold budget-to-time ratio, and let* $t_0 = \sup\{t \in [0, T] :$ $\mu(t) = \gamma(t) = 0$ be the first time that the shadow prices are zero. Then the following holds

1. If $B/T \ge \rho_0$, then $t_0 = T$. 2. If $B/T < \rho_0$, then $t_0 \ge T \left(\frac{B/T}{\rho_0}\right)^{\frac{1}{\alpha}}$ for some $\alpha \in (0, 1)$ independent of the initial state. In addition,

3. The budget never depletes before the end of the ho*rizon; that is,* $\bar{x}(t) > 0$ *for all* $t \ge 0$.

4. The ratio of budget to time remaining along the equilibrium path $\rho(t) \triangleq \bar{x}(t)/t$ is decreasing in t, that is, is increasing as time progresses toward the end of the horizon.

The result highlights several interesting properties of the optimal dynamic mechanism that we explain next. A key feature of the dynamic optimal mechanism is that the buyer gets exposed to all items; to achieve this, the budget must never be depleted (part 3 of the proposition), and therefore, the Myerson optimal static auction is always implemented toward the end. The latter is reflected in part 2 of the proposition, which shows that multipliers become zero in the interval $[t_0, 0]$. Furthermore, consistent with the previous properties, the buyer becomes less budget

constrained over time (part 4). An interesting feature is that the budget never gets depleted. We would like to emphasize that this is not a consequence of the assumed differentiability of the value functions in the fluid model. In numerical experiments in the discrete time model (without differentiability assumptions imposed), we find that budgets are not depleted at the end of the horizon under the optimal mechanism either.

To understand better the intuition behind these results, consider a single buyer with budget-to-time ratio $e_0 = \psi(c)F(\psi(c))$. In this case, if the Myerson auction is applied throughout the horizon (and the buyer bids truthfully in every auction), the budget would be depleted exactly at the end of the horizon. This would be the seller's optimal selling strategy if she would run second-price auctions and could commit to the reserve price over the entire horizon. However, in our setting with limited commitment, this strategy cannot be sustained, as suggested by part 1 of the proposition because $\rho_0 > e_0$. One can actually verify that this mechanism does not satisfy the DIC constraint because the buyer has incentives to shade his bids and underspend, so in the future he is less budget constrained and his future utility improves. In fact, with a higher budget-to-time ratio, owing to the seller's lack of commitment power, the seller will offer lower prices and lower reserve thresholds. The seller is only able to offer the Myerson optimal static auction when the buyer is "sufficiently" budget unconstrained.

Finally, it is interesting to observe that in the case c = 0 "waiting" and then implementing, the Myerson auction is an optimal dynamic mechanism.¹¹ When c = 0, the seller is only concerned with extracting the buyer's budget without any consideration for the number of items sold to her. Implementing the Myerson auction for a fraction of the horizon achieves this goal. It is interesting to observe that the case c = 0 is degenerate in that there are multiple optimal mechanisms. For example, implementing the static Myerson auction until the budget depletes and then offering the items for free to the buyer (which yields much higher utility to the buyer) is also optimal.

4.2.2. Multiple Buyers. In this section, we solve the HJB equations for multiple buyers numerically via finite differences and study the equilibrium paths for different initials budgets.

4.2.2.1. Setting. We consider the case with K = 2 buyers with values distributed as U[0,1], the seller's cost

is c = 0.5, and the horizon length is T = 10. We solve the HJB equations in the domain $[0, \overline{B}]^2 \times [0, T]$, where $\overline{B} = 2.5$. We discretize budgets by setting a uniform grid with 25 points and discretize time by setting a uniform grid with 1,000 points. We observe that the value functions are well behaved at every point of the grid, in the sense that the matrix $\mathbf{I} + D\mathbf{U}$ is nonsingular and the allocation factors satisfy $\eta \ge 0$.

4.2.2.2. Equilibrium Paths and Allocations. Figure 2 shows the equilibrium path of budgets when initial budgets are identical for the buyers and equal to $B_1 = B_2 = 2$. Here the equilibrium is symmetric: both buyers are offered the same mechanism at each point in time, and budgets deplete at the same rate. The item is allocated to the buyer with the highest value whenever that value is higher than the dynamic threshold $\phi^{-1}(c/\eta)$. As in the case of one buyer, when initial budgets are stringent, the optimal mechanism uses an allocation factor $\eta < 1$ to prevent buyers from depleting budgets before the end of the horizon and then switches to the Myerson auction closer to the end. In our numerical simulations, we find that, as in the case of one buyer, budgets are never totally depleted. Similarly, the marginal utility of an additional unit of a buyer's budget $\frac{\partial U_k}{\partial x_k}$ is nonnegative and decreasing, because

Figure 2. (Color online) Equilibrium Paths Under the Fluid Optimal Mechanism \mathbf{M}^* in the Case of Two Buyers with Symmetric Budgets $B_1 = B_2 = 2$



Note. Plots coincide for both buyers because these are symmetric.

budgets become relatively less stringent with time. Interestingly, the marginal utility of an additional unit of a competitor's budget $\frac{\partial \tilde{U}_k}{\partial x_i}$ is initially positive when the budget-to-time ratio is small, later negative as the budget-to-time ratio increases, and finally zero. We discuss these effects later in this section.

Figure 3 shows the equilibrium path of budgets when buyer 1 (the low buyer) has an initial lower budget of $B_1 = 1.5$ and buyer 2 (the high buyer) has an initial higher budget of $B_2 = 2.5$. Here the low buyer is assigned a lower allocation factor than the high buyer, that is, $\eta_1 \leq \eta_2$, leading to the high buyer's bids being ranked higher. This prevents the low buyer from winning too often and depleting her budget before the end of the horizon, thus maintaining competition throughout the horizon. As expected, the marginal utility of an additional unit of a buyer's budget $\frac{\partial U_k}{\partial x_k}$ is larger for the low buyer because she is more budget constrained. Interestingly, the marginal utility of an additional unit of a competitor's budget $\frac{\partial \bar{U}_k}{\partial x_i}$ is positive for the low buyer and negative for the high buyer, which implies that the low buyer is better off when her competitor has an even larger budget and that the high buyer is better off when her competitor has an even smaller budget. We discuss this effect at the end of this subsection.

4.2.2.3. Payments. In terms of payments, recall that when buyer *k* is the winner, she pays $a_{kk}y_k(v_i)$, where $y_k(v_i) = \phi_k^{-1}(\max(\eta_i \phi_i(v_i), c) / \eta_k)$ is the minimum report that wins against reports v_i from her competitor. Additionally, the loser *i* pays a fraction of the winner payment given by $a_{ik}y_k(v_i)$, where the off-diagonal *payment factors a_{ik}* are typically small compared with the diagonal terms a_{kk} . In the symmetric case, we have that $y_k(v_i) = \max(v_i, \phi^{-1}(c/\eta))$, and thus, the winner's payments have a second-price auction structure modulo the payment factors $A = (a_{ki})$. As in the twotier mechanism in the case of one buyer, the diagonal payment factors a_{kk} are less than or equal to one and nondecreasing with time. These provide a discount to the buyers to incentivize them to report truthfully and participate in early auctions when budgets are stringent.

When buyers have different initial budgets, we observe that the low-budget buyer is given a deeper discount, that is, $a_{11} < a_{22}$. Additionally, for the low-budget buyer the off-diagonal payment factor is negative $a_{12} < 0$, whereas for the high-budget buyer the off-diagonal payment factor is positive $a_{21} > 0$ (even though both factors are close to zero). Because $\frac{\partial \tilde{U}_1}{\partial x_2} > 0$, the low-budget buyer is better off when her competitor has a larger budget, and thus the seller





rewards the low buyer when the high buyer wins the item and depletes her budget. Conversely, because $\frac{\partial U_2}{\partial x_1} < 0$, the high buyer is better off when her competitor has a lower budget, and thus the seller charges the high buyer when the low buyer wins the item and depletes her budget.

To understand why the low buyer is better off when her competitor has a larger budget, we first observe that when a competitor has a larger budget, there are two effects. First, because the low buyer becomes "weaker," her allocation factor in the optimal mechanism increases to even out the competition. As a consequence, the low buyer's effective threshold goes down, and in isolation, she would win more often and pay less when she wins. Second, because the competitor has more budget to spend, the competitor's allocation factor increases as well, and through this effect, the low buyer wins less often and pays more when she wins. The overall effect is a combination of these two, and the net effect on flow utility depends on which of these effects dominates. As we saw earlier in our numerical results, the first effect is stronger when the competitor has a relatively larger budget, and in this case, it dominates the second.

5. Near Optimality in the Discrete Model

In this section, we aim to connect the optimal mechanism in the fluid continuous time model to the original discrete problem we initially laid out in Equation (1). In particular, in Section 5.1, we first define the notions of approximate incentive compatibility for the buyers and approximate sequential rationality for the seller. In Section 5.2, we develop a numerical framework to show that the prescription obtained from the continuous time fluid model is approximately incentive compatible and approximately sequentially rational in the discrete time model. Finally, in Section 5.3, we establish theoretically such results for the case of a single buyer.

When attempting to establish the connection above, we note that the optimal mechanism \mathbf{M}^* from Theorem 1, however, is not guaranteed to be budget feasible for every sample path of the discrete model. Thus motivated, we introduce the *adjusted mechanism* $\hat{\mathbf{M}}[\mathbf{x}, t] = (\hat{\mathbf{P}}[\mathbf{x}, t], \hat{\mathbf{Z}}[\mathbf{x}, t])$ given by

$$\hat{P}_{k}[\mathbf{x}, t](\mathbf{v}) = \mathbf{1}\{v_{k} > y_{k}(\mathbf{v}_{-k})\}, \qquad k = 1, \dots, K,$$
$$\hat{Z}_{k}[\mathbf{x}, t](\mathbf{v}) = \min\left\{\frac{x_{k}}{\delta}, \sum_{i \in \mathcal{H}} a_{ki}y_{i}(\mathbf{v}_{-i})\hat{P}_{i}[\mathbf{x}, t](\mathbf{v})\right\},$$
$$k = 1, \dots, K,$$

for all $\mathbf{v} \in \mathbb{R}^{K}$, where we denote by $y_{k}(\mathbf{v}_{-k}) = \inf \{v : \eta_{k}\phi_{k}(v) > \max_{i \in \mathcal{H} \setminus k} \eta_{i}\phi_{i}(v_{i}), \eta_{k}\phi_{k}(v) > c\}$ the smallest value for the *k*th buyer who wins against reports \mathbf{v}_{-k} from the competitors, and $\mathbf{A} = (a_{ki}) = (\mathbf{I} + D\bar{\mathbf{U}}|_{\mathcal{H}})^{-1}$.

Additionally, if $v_k = \bot$ for some buyer k, then $\hat{P}_k[\mathbf{x}, t](\mathbf{v}) = 0$, $\hat{Z}_k[\mathbf{x}, t](\mathbf{v}) = 0$, k = 1, ..., K. Note that the modified mechanism coincides with \mathbf{M}^* whenever the budgets of buyers are not too small. The mechanism for the discrete stochastic model is given by setting $\hat{\mathbf{M}}[x, t] \triangleq \hat{\mathbf{M}}[x, t_i]$, where $t_i = \delta i$.

5.1. Approximate Incentive Compatibility and Sequential Rationality

A stage mechanism $\mathbf{m} \in \mathcal{M}$ is e^{IC} -incentive compatible for the buyers at state $(\mathbf{x}, n) \in \mathcal{F}^{\delta}$ with respect to dynamic mechanism $\mathbf{M} \in \mathbb{M}^{\delta}$ if at state (\mathbf{x}, n) any buyers' incentive to misreport *after* learning her value is at most δe^{IC} when the seller offers the stage mechanism \mathbf{m} at the current state, and the seller offers dynamic mechanism \mathbf{M} and the buyer reports truthfully onward. We formalize this concept next. In the following definition, $U_k^{\mathbf{M}}(\cdot)$ denotes the expected utility of buyer k under dynamic mechanism \mathbf{M} and truthful reporting in the discrete stochastic model.

Definition 2. Stage mechanism $\mathbf{m} \in \mathcal{M}$ is e^{IC} -incentive compatible at state $(\mathbf{x}, n) \in \mathcal{G}^{\delta}$ with respect to dynamic mechanism $\mathbf{M} \in \mathbb{M}^{\delta}$ if

$$\begin{split} \delta u_{k}(\mathbf{m}, v, w) &+ \mathbb{E}_{\mathbf{v}_{-k}} \Big[U_{k}^{\mathbf{M}}(\mathbf{x} - \delta \mathbf{z}(w, \mathbf{v}_{-k}), n-1) \Big] \\ &\leq \delta u_{k}(\mathbf{m}, v, v) + \mathbb{E}_{\mathbf{v}_{-k}} \Big[U_{k}^{\mathbf{M}}(\mathbf{x} - \delta \mathbf{z}(v, \mathbf{v}_{-k}), n-1) \Big] + \delta \epsilon^{\mathrm{IC}}, \end{split}$$

for every value $v \in \mathbb{R}_+$, report $w \in \mathbb{R}_\perp$, and buyer *k*.

Now we provide our definition of approximate seller sequential rationality (this is related to the concept of contemporaneous perfect ϵ -equilibria of Mailath et al. 2005). In the following definition, $\Pi^{M}(\cdot)$ denotes the expected profit of the seller under dynamic mechanism **M** and truthful reporting in the discrete stochastic model.

Definition 3. A dynamic mechanism $\mathbf{M} \in \mathbb{M}$ is said to be $(\epsilon^{SR}, \epsilon^{IC})$ -sequentially rational for the seller at state $(\mathbf{x}, n) \in \mathcal{G}^{\delta}$ if

$$\begin{split} \delta \pi(\mathbf{m}) + \mathbb{E}_{\mathbf{v}} \Big[\Pi^{\mathbf{M}}(\mathbf{x} - \delta \mathbf{z}(\mathbf{v}), n-1) \Big] &\leq \delta \pi(\mathbf{M}[\mathbf{x}, n]) \\ &+ \mathbb{E}_{\mathbf{v}} \Big[\Pi^{\mathbf{M}}(\mathbf{x} - \delta \mathbf{Z}[\mathbf{x}, n](\mathbf{v}), n-1) \Big] + \delta \epsilon^{\mathrm{SR}}, \end{split}$$

for every stage mechanism $\mathbf{m} \in \mathcal{M}$ such that

1. the mechanism is budget feasible;

2. the mechanism is ϵ^{IC} -incentive compatible for buyers at state $(\mathbf{x}, n) \in \mathcal{G}^{\delta}$ with respect to dynamic mechanism **M**;

3. payments are bounded by $0 \le z(v) \le \overline{v}$.

In particular, when a dynamic mechanism is ϵ^{IC} -incentive compatible and $(\epsilon^{\text{SR}}, \epsilon^{\text{IC}})$ -sequentially

rational with appropriately small values for ϵ^{IC} and ϵ^{SR} , we will say that the mechanism is near optimal.

The previous definitions consider sellers' and buyers' incentives to deviate in one period. However, if these incentives are shown to be small, by the one-stagedeviation principle for finite-horizon games, these results can be shown to imply that their incentive to deviate to a dynamic strategy spanning multiple periods is also small.

5.2. Numerical Framework to Evaluate Near Optimality

In this section, we describe our approach to numerically validate approximate sequential rationality and approximate incentive compatibility of the fluid dynamic mechanism [that solves the HJB Equations (7) and (8)] with multiple buyers in the discrete time model when the time interval between successive auctions δ is small.

First, we select a set of representative instances with multiple buyers. For each instance, we follow these steps:

Step 1. Solve the system of PDEs in Equation (9) via finite differences to obtain the fluid mechanism. Check that the value functions are well behaved.

Step 2. Implement the prescribed *adjusted* fluid mechanism $\hat{\mathbf{M}}[\mathbf{x}, t]$ provided by the solution of the system of PDEs in the discrete time model, and evaluate its performance both for buyers and the seller using Monte Carlo simulation.

Step 3. Compute, in the discrete time model and for every point in a fine grid of the state space, the maximum utility buyers can obtain by a single-stage unilateral deviation from truthful reporting.

Step 4. Compute, in the discrete time model and for every point in a fine grid of the state space, the maximum profit the seller can obtain by a single-stage deviation from the optimal fluid mechanism to a different budgetfeasible and approximately IC mechanism.

A detailed description of these steps above and the approach we take is presented in Appendix A of the supplemental material. We next describe the results.

Our evaluation test bed is one where the values of buyers are uniformly distributed in [0,1] and the time horizon is T = 10.¹² We initially vary the number of buyers $K = \{1, 2\}$, the opportunity cost $c = \{0.25, 0.5, 0.75\}$, and the number of items $N = \{10, 10^2, 10^3\}$. (Recall that $\delta = T/N$.) The maximum initial budget is set to B = 2.5. We discretize the budgets by setting a uniform grid with 50 points. For Step 1, we discretize time by setting a uniform grid with 10^4 points. For Step 2, value functions were calculated using Monte Carlo simulation with 300 sample paths; the resulting mean standard errors are small in all cases. For Step 4, we convert the single-stage optimization problem of the seller to a linear program by allowing for randomized mechanisms over a grid of possible payments. We discretize the possible payments by setting a uniform grid over the support of values $[0, \bar{v}]$ with 50 and 40 when the number of buyers is one and two, respectively. When testing for approximate incentive compatibility and sequential rationality in the discrete time model, we coarsen the grid as the number of buyers increases for computational tractability. For three buyers and more, the computational complexity associated with the discrete case explodes. Recall that this is in stark contrast to the prescription stemming from the continuous time fluid model that is very simple to obtain numerically (it only requires solving numerically a system of PDEs). For three buyers, the discrete time model can only be handled with significantly coarser grids. Results testing for approximate incentive compatibility and sequential rationality in the discrete time model with such grids for three buyers are presented in Appendix A.5 of the supplemental material. In this case, despite the discretization error, the optimality gap associated with the fluid prescription is still small.

We evaluate the fluid-based mechanism in the discrete model in the sense of approximate incentive compatibility and approximate sequential rationality. First, as a preliminary check, we confirm that in all our instances the performance of the adjusted dynamic fluid mechanism in the discrete model approaches the predicted performance in the fluid model when δ shrinks to zero. We next present a summary of our main numerical results regarding near optimality.

5.2.1. Approximate Incentive Compatibility. We denote the IC error at state (\mathbf{x}, n) as follows:

$$E_{\infty}(\mathbf{x},n) = \frac{1}{\delta} \max_{k} \max_{v} \max_{w \in \mathbb{R}_{\perp}} \left\{ \hat{U}_{k}(v,w) - \hat{U}_{k}(v,v) \right\},$$

where, with some abuse of notation, $\hat{U}_k(v, w)$ denotes the interim utility-to-go of buyer *i* under dynamic mechanism $\hat{\mathbf{M}}$ when she reports *w* and her value is *v* in period *n* and all competitors report truthfully. Note that we divide by δ because flow utility scales with δ . The top half of Table 1 reports the utility a buyer can gain from misreporting her value in the first period of the horizon, averaged over all possible initial budgets in the grid, for *c* = 0.5 (results for other costs are similar and not reported). The results show that the optimal fluid mechanism is ϵ^{IC} -incentive compatible in the discrete time model, where ϵ^{IC} approaches zero as the number of auctions *N* becomes large (or the length between periods δ becomes small).

5.2.2. Approximate Sequential Rationality. We compare the seller's profit when all buyers report truthfully under our mechanism, denoted by $\Pi^{\hat{\mathbf{M}}}(\mathbf{x}, n)$, with the profit the seller could obtain from a one-shot

Table 1. Approximate Incentive Compatibility (Step 3) and

 Approximate Sequential Rationality (Step 4)

	K	
Ν	1	2
Approximate in	centive compatibility (Step 3)	
10	0.094901	0.145863
10 ²	0.024100	0.041523
10 ³	0.000328	0.000442
Approximate se	quential rationality (Step 4)	
10	0.069487	0.141417
10 ²	0.018171	0.032844
10^{3}	0.003747	0.009116

Notes. Top half of table reports the maximum incentive compatible error $E_{\infty}(\mathbf{x}, N)$ across buyers, averaged over all budgets \mathbf{x} , evaluated at the first period (Step 3). The bottom half reports the difference $(\Pi^*(\mathbf{x}, N) - \Pi^{\hat{\mathbf{M}}}(\mathbf{x}, N))/\delta$ averaged over all budgets \mathbf{x} , evaluated at the first period (Step 4).

deviation denoted by $\Pi^*(\mathbf{x}, n)$. Here $\Pi^*(\mathbf{x}, n)$ is the profit obtained from offering an optimal approximately incentive-compatible and budget-feasible mechanism for the current period and then offering mechanism $\hat{\mathbf{M}}$ in the following periods. In $\Pi^*(\mathbf{x}, n)$, we allow for mechanisms that are $E_{\infty}(\mathbf{x}, n)$ -incentive compatible to guarantee that the candidate stage mechanism $\hat{\mathbf{M}}[\mathbf{x}, n]$ is feasible for the seller's problem. The bottom half of Table 1 reports the profit the seller can gain from offering an optimal mechanism in the first time period and then offering our mechanism in the following time periods. The results show that the optimal fluid mechanism is ($\epsilon^{\text{SR}}, \epsilon^{\text{IC}}$)-sequentially rational in the discrete time model, where ϵ^{SR} approaches zero as the number of auctions *N* becomes large.

Overall, the analysis is reassuring in showing that the optimal fluid mechanism is near optimal in the discrete model. Hence, the continuous time fluid model can provide a basis to construct an easily computable mechanism for an otherwise seemingly intractable problem. In the next section, we also provide a theoretical foundation for this claim.

5.3. Provable Near Optimality for Single-Buyer Case

In this section, we aim to provide analytical support for the fluid model by proving that the optimal fluid dynamic mechanism is approximately incentive compatible for a single buyer and approximately sequentially rational for the seller in the discrete stochastic model when $\delta \downarrow 0$. Note that for the single-buyer case, the mechanism $\hat{M}[x,t] = (\hat{P}[x,t], \hat{Z}[x,t])$ takes the simplified form

$$\begin{split} \hat{P}[x,t](v) &= \mathbf{1}\{v > r(x,t)\},\\ \hat{Z}[x,t](v) &= \min\left\{\frac{x}{\delta}, \frac{r(x,t)}{1+\mu(x,t)}\right\} \mathbf{1}\{v > r(x,t)\} \end{split}$$

Our first result shows that the mechanism \hat{M} is $O(\delta^{1/2})$ -incentive compatible for the buyer in the

discrete stochastic model.¹³ We prove the result under the assumption that the initial budget is not too small (in an appropriately defined sense). To this end, we define $\bar{x}(s; x, t)$ as the budget remaining at time $s \in [0, t]$ in the fluid model when the initial state is $(x, t) \in \mathcal{G}$ and $\hat{\mathcal{G}}^{\delta} = \{(x, t) \in \mathcal{G} : \bar{x}(0; y, t) > \delta \bar{v} \ \forall y \in [x - \delta \bar{v}, x]\}$ as the initial states of the fluid model for which the budget remaining at time zero is greater then or equal to $\delta \bar{v}$ even if the initial budget is reduced by an amount of at most $\delta \bar{v}$.¹⁴ When the initial state lies in the set $\hat{\mathcal{G}}^{\delta}$, the mechanism along the equilibrium path is guaranteed to coincide with M^* because $\frac{r(x,t)}{1+\mu(x,t)} \leq \bar{v}$, and the budget decreases monotonically under the fluid mechanism.

Theorem 3 (Approximate Buyer Incentive Compatibility). Suppose that the state $(x, n) \in \mathcal{F}^{\delta}$ satisfies $(x, t_n) \in \hat{\mathcal{F}}^{\delta}$. Then stage mechanism $\hat{M}[x, n]$ is $O(\delta^{1/2})$ -incentive compatible for the buyer at state (x, n) with respect to dynamic mechanism \hat{M} .

There are two main challenges associated with showing this result. The first challenge stems from the fact that the prescribed dynamic mechanism (adapted from the fluid continuous time model) is a closed-loop mechanism designed to stay on the equilibrium path of the fluid model. In the discrete model, the presence of stochastic fluctuations will take the state off the fluid equilibrium path, and the discrete model path may slowly diverge from the original continuous time equilibrium path. A first key part of the proof revolves around establishing that these paths stay appropriately close; this is done by combining concentration inequality–type arguments with dynamical system analysis arguments.

Although the preceding point guarantees that the mechanism offered in the discrete model is "close" to the one offered along the equilibrium path of the fluid model, it does not necessarily imply that the mechanism is approximately DIC. The second key part of the proof establishes that dynamic incentives in both the discrete and fluid models are approximately "aligned" by guaranteeing that the utility per period and the marginal utility of an additional unit of budget in both systems are close. The first condition follows directly from the fact that the realized mechanisms in both systems are close. For the second condition, we need to show that the marginal utility of an additional unit of budget for the buyer in the discrete model under dynamic mechanism \hat{M} is close to the marginal utility of budget for the buyer in the fluid model $\mu(x, t) = U_x(x, t)$. In particular, we show the following.

Proposition 2. For all states $(x, n), (y, n) \in \mathcal{S}^{\delta}$ with $0 \le x - y \le \delta \overline{v}$ and $(x, t_n) \in \hat{\mathcal{S}}^{\delta}$, the buyer's expected utility under dynamic mechanism \hat{M} and truthful reporting in the discrete model satisfies

$$\left| U^{\hat{M}}(x,n) - U^{\hat{M}}(y,n) - \mu(x,t_n)(x-y) \right| = O(\delta^{3/2}).$$

It is important to highlight that although it is typical to show results on how well value functions are approximated by those in an appropriate fluid model, it is much more challenging to establish a tight connection between the finite differences of the value functions and the derivative of the value function in the fluid model. To illustrate this point, usually, one can approximate the value functions up to order $\delta^{1/2}$ invoking concentration inequalities. However, such control does not allow approximation of the finite differences with any satisfactory accuracy. Instead, we need to directly estimate finite differences by studying the joint evolution of two stochastic trajectories in the discrete model with close initial state. This is the core of the challenge to establish Proposition 2. This result and the techniques used to establish it might be of interest beyond the current application.

This result, together with the fact that utilities per period are close (because trajectories are close), allows us to translate the dynamic incentive compatibility condition of the fluid model given in Equation (6c) to its discrete counterpart given in Definition 2. Thus, if the derivatives of the buyer's value function are close to the corresponding finite differences (in the discrete model), reporting truthfully at any point in time is indeed approximately DIC.

We are now in a position to state our main result on the characterization of \hat{M} . The following result shows that at every stage $(x, n) \in \mathcal{G}^{\delta}$ the seller's incentive to deviate from $\hat{M}[x, n]$ to any other stage mechanism is small. We allow the seller to deviate to stage mechanisms that are $O(\delta^{1/2})$ -incentive compatible to include the proposed mechanism \hat{M} as a feasible deviation.

Theorem 4 (Approximate Seller Sequential Rationality). Suppose that the state $(x, n) \in \mathcal{F}^{\delta}$ satisfies $(x, t_n) \in \hat{\mathcal{F}}^{\delta}$. Then dynamic mechanism \hat{M} is $(O(\delta^{1/3}), O(\delta^{1/2}))$ -sequentially rational for the seller at state (x, n).

We prove the result in four steps. One of the challenges is similar to that of Theorem 3 in that the prescribed mechanism is a closed loop and is designed to stay on the equilibrium path of the fluid model. The proof of Theorem 3 shows that the fluid equilibrium path and the discrete model path stay appropriately close, and thus, the realized mechanisms in both systems are close.

Second, we note that as soon as one allows for approximate DIC (as opposed to DIC), one expands the set of mechanisms one could consider. In turn, we establish that any mechanism satisfying $O(\delta^{1/2})$ -incentive compatibility is $O(\delta^{1/2})$ -feasible for the one-shot fluid mechanism design problem in Equation (6). Here we use Proposition 2 again to show that for the buyer, the marginal utility of an additional unit of budget under mechanism \hat{M} in the discrete model is close to $\bar{U}_x(x, t)$, that is, the marginal utility of budget in the fluid model. Third, we prove the following general result

for one-stage mechanism design: relaxing incentive compatibility and individual rationality constraints by ϵ increases the seller's profit by $O(\epsilon^{2/3} + \epsilon)$. This result drives the difference in convergence rates between the buyer's approximate incentive compatibility and the seller's approximate individual rationality.

Finally, we conclude in a similar fashion as in the last step in the proof of the approximate incentive compatibility result. Here we leverage that the fluid mechanism is optimal in the HJB Equation (6) to show that the expected seller profit from deviating to stage mechanism $m \in \mathcal{M}$ is close to $\Pi^{\hat{M}}(x, n)$, the expected seller profit under \hat{M} and truthful reporting in the discrete model. To prove this result, we first show a parallel result to Proposition 2: the marginal profit of an additional unit of budget under mechanism \hat{M} in the discrete model is close to $\gamma(x, t) = \overline{\Pi}_x(x, t)$, that is, the marginal profit of budget in the fluid model.

Proposition 3. For all states $(x, n), (y, n) \in \mathcal{P}^{\delta}$ with $0 \leq x - y \leq \delta \bar{v}$ and $(x, t_n) \in \hat{\mathcal{P}}^{\delta}$, the seller's expected profit under dynamic mechanism \hat{M} and truthful reporting in the discrete model satisfies

$$\left|\Pi^{\tilde{M}}(x,n) - \Pi^{\tilde{M}}(y,n) - \gamma(x,t_n)(x-y)\right| = O(\delta^{3/2}).$$

Remark 1. As a corollary of this result, one has that $U^{\hat{M}}(x, n)$, the expected buyer utility under \hat{M} and truthful reporting in the discrete model, converges to $\overline{U}(x, t_n)$, the buyer value function in the fluid model, as $\delta \downarrow 0$. This result follows because stochastic and fluid paths are close, together with the fact that the expected utility per period is Lipschitz continuous close to the fluid path. Additionally, one has that $\Pi^{\hat{M}}(x, n)$, the seller profit under \hat{M} and truthful reporting in the discrete model, converges to $\overline{\Pi}(x, t_n)$, the seller value function in the fluid model.

Remark 2. We note that the current proofs for the case of a single buyer heavily leverage the structural properties of the dynamic optimal mechanism given in Section 4.1.1. As a result, extending these approximation results to multiple buyers seems challenging because it is not clear whether such results can be established given the multidimensional nature of the problem. We suspect, however, that similar arguments would still go through if these properties are assumed true.

6. Conclusions

In this paper we studied dynamic mechanism design when selling a sequence of items with limited commitment to buyers who face a cumulative budget constraint. We showed that an envelope approach can be applied in a corresponding fluid continuous time model in such settings. We then used this approach to characterize the dynamic optimal mechanism, highlighting novel incentive issues at play. We also provide justification for using the prescription of the fluid model in the original discrete time model.

We believe there are several interesting directions for future research. First, extending the approximation result of Section 5.3 to multiple buyers would provide an even stronger justification for the fluid model. We conjecture this involves significant technical challenges though. Second, the model we introduced was the simplest possible extension of a classic setting so that we could highlight the effect of budgets on optimal mechanism design. There are several interesting extensions of our model that may be worth considering in the future, such as incorporating valuations with a common component, assuming that budgets are buyers' private information, considering nonstationary value distributions, and studying models in which campaigns do not simultaneously start and end at the same time. Similarly, to simplify the formulation and focus on the link with the fluid model, we anchored the model on one particular alternative to deal with buyers' nonparticipation decisions. Understanding alternative formulations with a richer class of dynamic mechanisms that may be contingent on the subset of participant bidders in each auction is also an interesting avenue.

Finally, we consider mechanisms under a strong form of limited commitment: the seller commits to the rules of the current auction, but she cannot commit to those of future auctions. On the other extreme, with full commitment and the ability to charge participation fees upfront, the seller can sometimes achieve full surplus extraction. In between there is a spectrum representing different levels of the seller's commitment power. It would be enlightening to explore the auction design problem under different practical constraints and how these relate to various levels of commitment.

Acknowledgments

The authors thank Dirk Bergemann, Qingmin Liu, Pino Lopomo, Alessandro Pavan, Saša Pekeč, Andrea Prat, and Andy Skrzypacz for helpful insights, as well as seminar participants at Columbia, Duke, Google, Informs, and the Revenue Management Conference, and Huseyin Gurkan for his help in preparing the manuscript.

Endnotes

¹We note that this regime is particularly pertinent for display advertising in which an advertiser spends a tiny fraction of her budget in each auction.

² When δ is small, the scaling is simply meant to highlight that each auction can only consume a small fraction of the budget. Note that the total number of items is equal to T/δ . Hence, as δ decreases, the number of items increases, and by imposing the scaling above, the total value gathered by buyers over the time horizon, as well as their total expenditure, remains bounded.

³For example, in online advertising, opportunity costs are common because publishers sell their inventory across multiple channels. We also comment on the case c = 0 in later sections.

⁴ In the context of online advertising, this could be motivated by the fact that within a time frame, the order of visits of unique users to a website is random.

⁵We restrict attention to deterministic payments. We note that this restriction is without loss of optimality in the fluid model that we introduce in Section 3 because there budgets are depleted deterministically according to the expected payments.

⁶We also considered other alternatives to discipline seller's off-theequilibrium-path behavior, such as imposing stronger forms of individual rationality or a formulation in which mechanisms are conditioned on the subset of participant bidders. However, these involve significant additional complexities that would obfuscate the analysis of the fluid model.

⁷ Technically, to apply the revelation principle, one needs to formalize the extensive form game between the seller and the buyers. One alternative is to formulate this game as a sequence of stage games in which in each stage the buyers first learn their values, then the seller chooses a mechanism in \mathcal{M} , and finally, the buyers play their actions. For the solution concept to be well defined, we need to restrict the seller's mechanisms to guarantee that the continuation bidding game between buyers always admits a pure equilibrium, even off the equilibrium path. We also note that the argument presented by Bester and Strausz (2001) regarding the inapplicability of the standard revelation principle in settings with limited commitment does not apply in our setting because private information is not persistent.

⁸We note that because the fluid model represents a limit in which bidders spend an infinitesimal amount of their budgets in each auction, budgets will not bind within each period whenever they are positive. Therefore, the features of the optimal mechanism in a oneshot auction with budgets, as in Pai and Vohra (2014), do not appear in our analysis.

⁹Given a matrix $\mathbf{A} \in \mathbb{R}^{K \times K}$ and a subset $\mathcal{H} \subseteq \{1, \dots, K\}$, we denote by $\mathbf{A}|_{\mathcal{H}} \in \mathbb{R}^{|\mathcal{H}| \times |\mathcal{H}|}$ the submatrix restricted to indices in \mathcal{H} .

¹⁰ In particular, one can show that the mechanism consisting of waiting until the time that implementing the Myerson auction would deplete the budget exactly at the end of the horizon cannot be fluid optimal.

¹¹ Recall that we assumed c > 0 to derive the dynamic optimal mechanism of Corollary 1.

¹²We also run instances with other valuation distributions, for example, truncated exponential, and obtained similar results that we omit for brevity.

¹³ We say $f(\delta)$ is big O of $g(\delta)$ or $f(\delta) = O(g(\delta))$ if and only if there exists C > 0 and a real number δ_0 such that $|f(\delta)| \le C|g(\delta)|$ for all $0 \le \delta \le \delta_0$. ¹⁴ The set $\hat{\mathcal{P}}^{\delta}$ can be readily characterized by noting that $\bar{x}(0; x, t) = t_0(x, t)(\rho_0 - e_0)$, where we define $t_0(x, t)$, following Proposition 1, as the first time that the shadow prices are zero in the fluid model when the starting state is $(x, t) \in \mathcal{P}$. Although a closed-form expression for $t_0(x, t)$ is not available in general, we can obtain a subset of $\hat{\mathcal{P}}^{\delta}$ by using the lower bound for $t_0(x, t)$ provided in Proposition 1. Additionally, the previous bound implies that $\hat{\mathcal{P}}^{\delta}$ converges to the set of positive states $\{(x, t) \in \mathcal{P} : x > 0, t > 0\}$ as $\delta \to 0$.

References

Akan M, Ata B, Dana JD Jr (2015) Revenue management by sequential screening. J. Econom. Theory 159(September, Part B):728–774.

- Alós-Ferrer C, Ritzberger K (2008) Trees and extensive forms. J. Econom. Theory 143(1):216–250.
- Ashenfelter O (1989) How auctions work for wine and art. J. Econom. Perspect. 3(3):23–36.

- Başar T, Olsder G (1999) *Dynamic Noncooperative Game Theory*, 2nd ed. (SIAM, Philadelphia).
- Balseiro SR, Besbes O, Weintraub GY (2015) Repeated auctions with budgets in ad exchanges: Approximations and design. *Man-agement Sci.* 61(4):864–884.
- Benoît J-P, Krishna V (2001) Multiple-object auctions with budget constrained bidders. *Rev. Econom. Stud.* 68(1):155–179.
- Bergin J, MacLeod WB (1993) Continuous time repeated games. Internat. Econom. Rev. 34(1):21–37.
- Bester H, Strausz R (2001) Contracting with imperfect commitment and the revelation principle: The single agent case. *Econometrica* 69(4):1077–1098.
- Bhattacharya S, Goel G, Gollapudi S, Munagala K (2010) Budget constrained auctions with heterogeneous items. *Proc. 42nd ACM Sympos. Theory Comput.* (ACM, New York), 379–388.
- Board S, Skrzypacz A (2015) Revenue management with forwardlooking buyers. J. Political Econom. 124(4):1046–1087.
- Borgs C, Chayes J, Immorlica N, Mahdian M, Saberi A (2005) Multiunit auctions with budget-constrained bidders. *Proc. 6th ACM Conf. Electronic Commerce* (ACM, New York), 44–51.
- Bressan A, Shen W (2004) Small bv solutions of hyperbolic noncooperative differential games. SIAM J. Control Optim. 43(1):194–215.
- Brusco S, Lopomo G (2008) Budget constraints and demand reduction in simultaneous ascending-bid auctions. *J. Indust. Econom.* 56(1):113–142.
- Che YK, Gale I (1998) Standard auctions with financially constrained bidders. *Rev. Econom. Stud.* 65(1):1–21.
- Che YK, Gale I (2000) The optimal mechanism for selling to a budgetconstrained buyer. J. Econom. Theory 92(2):198–233.
- Deb R, Said M (2015) Dynamic screening with limited commitment. J. Econom. Theory 159(September, Part B): 891–928.
- Dilme F, Li F (2019) Revenue management without commitment: Dynamic pricing and periodic fire sales. *Rev. Econom. Stud.* Forthcoming.
- Dobrzynski JH (2012) How an acquisition fund burnishes reputations. New York Times (March 14), https://www.nytimes.com/ 2012/03/15/arts/artsspecial/a-fund-for-buying-art-burnishes collections-and-reputations.html.
- Evans LC (2010) Partial Differential Equations, Graduate Studies in Mathematics, 2nd ed. (American Mathematical Society, Providence, RI).
- Gallien J (2006) Dynamic mechanism design for online commerce. Oper. Res. 54(2):291–310.
- Gershkov A, Moldovanu B (2014) Dynamic Allocation and Pricing: A Mechanism Design Approach (MIT Press, Cambridge, MA).
- Gul F, Sonnenschein H, Wilson R (1986) Foundations of dynamic monopoly and the coase conjecture. J. Econom. Theory 39(1):155–190.
- Gummadi R, Key PB, Proutiere A (2012) Optimal bidding strategies and equilibria in dynamic auctions with budget constraints. *Proc. 49th Annual Allerton Conf. Comm. Control Comput.* (IEEE, New York), 588.
- Hendon E, Jacobsen HJ, Sloth B (1996) The one-shot-deviation principle for sequential rationality. *Games Econom. Behav.* 12(2):274–282.
- Hörner J, Samuelson L (2011) Managing strategic buyers. J. Political Econom. 119(3):379–425.
- Iyer K, Johari R, Sundararajan M (2014) Mean field equilibria of dynamic auctions with learning. *Management Sci.* 60(12): 2949–2970.
- Kakade S, Lobel I, Nazerzadeh H (2013) Optimal dynamic mechanism design and the virtual pivot mechanism. Oper. Res. 61(3): 837–854.
- Korula N, Mirrokni V, Nazerzadeh H (2015) Optimizing display advertising markets: Challenges and directions. *IEEE Internet Comput.* 20(1):28–35.

- Laffont JJ, Robert J (1996) Optimal auction with financially constrained buyers. *Econom. Lett.* 52(2):181–186.
- Lewis TR, Yildirim H (2002) Managing dynamic competition. Amer. Econom. Rev. 92(4):779–797.
- Liu Q, Mierendorff K, Shi X (2019) Auctions with limited commitment. Amer. Econom. Rev. 109(3):876–910.
- Mailath GJ, Postlewaite A, Samuelson L (2005) Contemporaneous perfect epsilon-equilibria. Games Econom. Behav. 53(1):126–140.
- Maskin ES (2000) Auctions, development, and privatization: Efficient auctions with liquidity-constrained buyers. *Eur. Econom. Rev.* 44(4–6):667–681.
- Milgrom P (2004) *Putting Auction Theory to Work* (Cambridge University Press, New York).
- Muthukrishnan S (2009) Ad exchanges: Research issues. Internet and Network Economics, Lecture Notes in Computer Science, vol. 3828 (Springer, New York), 1–12.
- Myerson R (1981) Optimal auction design. Math. Oper. Res. 6(1):58-73.
- Nazerzadeh H, Saberi A, Vohra R (2013) Dynamic cost-per-action mechanisms and applications to online advertising. *Oper. Res.* 6(1):98–111.
- Pai MM, Vohra R (2014) Optimal auctions with financially constrained bidders. J. Econom. Theory 150(0):383–425.
- Pavan A, Segal I, Toikka J (2014) Dynamic mechanism design: A myersonian approach. *Econometrica* 82(2):601–653.
- Pitchik C (2009) Budget-constrained sequential auctions with incomplete information. *Games Econom. Behav.* 66(2):928–949.
- Simon LK, Stinchcombe MB (1989) Extensive form games in continuous time: Pure strategies. *Econometrica* 57(5):1171–1214.
- Skreta V (2006) Sequentially optimal mechanisms. Rev. Econom. Stud. 73(4):1085–1111.
- Skreta V (2015) Optimal auction design under non-commitment. J. Econom. Theory 159(September, Part B):854–890.
- van den Berg GJ, van Ours JC, Pradhan MP (2001) The declining price anomaly in dutch dutch rose auctions. *Amer. Econom. Rev.* 91(4): 1055–1062.
- Vulcano G, van Ryzin G, Maglaras C (2002) Optimal dynamic auctions for revenue management Management Sci. 48(11):1388–1407.
- Yuan S, Wang J, Zhao X (2013) Real-time bidding for online advertising: Measurement and analysis. Proc. 7th Internat. Workshop Data Mining Online Advertising (ACM, New York), Article 3.

Santiago R. Balseiro is an associate professor in the Decision, Risk, and Operations Division at the Graduate School of Business, Columbia University. His primary research interests are in the area of dynamic optimization, stochastic systems, and game theory, with applications in revenue management and internet advertising.

Omar Besbes is an associate professor in the Decision, Risk, and Operations Division at the Graduate School of Business, Columbia University. His research focuses on the development of mathematical models and algorithms for dynamic and data-driven decision making, with applications in revenue management, operations, and service systems.

Gabriel Y. Weintraub is an associate professor in the Operations, Information, and Technology Division at the Stanford Graduate School of Business. His research covers several subjects that lie in the intersection between operations/management science and microeconomics. He is particularly interested in developing mathematical, computational, and econometric models to study the economics and the optimization of marketplaces.