BIDDING WITH BUDGETS: ALGORITHMIC AND DATA-DRIVEN BIDS IN DIGITAL ADVERTISING

By

Dirk Bergemann, Alessandro Bonatti and Nicholas Wu

March 2025

COWLES FOUNDATION DISCUSSION PAPER NO. 2429

COWLES FOUNDATION FOR RESEARCH IN ECONOMICS

YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281

http://cowles.yale.edu/

Bidding with Budgets: Algorithmic and Data-Driven Bids in Digital Advertising *

Dirk Bergemann[†] Alessandro Bonatti[‡] Nicholas Wu[§]

March 2, 2025

Abstract

In digital advertising, the allocation of sponsored search, sponsored product, or display advertisements is mediated by auctions. The generation of bids in these auctions for attention is increasingly supported by auto-bidding algorithms and platform-provided data. We analyze the equilibrium properties of a sequence of increasingly sophisticated auto-bidding algorithms. First, we consider the equilibrium bidding behavior of an individual advertiser who controls the auto-bidding algorithm through the choice of their budget. Second, we examine the interaction when all bidders use budget-controlled bidding algorithms. Finally, we derive the bidding algorithm that maximizes the platform's revenue while ensuring all advertisers continue to participate.

Keywords: Data, Advertising, Competition, Digital Platforms, Auctions, Automated Bidding, Managed Advertising Campaigns

JEL Classification: D44, D82, D83.

*We gratefully acknowledge financial support from NSF SES 1948336 and ONR MURI N00014-24-1-2742. We thank Santiago Balseiro and Negin Golrezaei for helpful conversations. We thank David Wambach for excellent research assistance.

[†]Yale University, dirk.bergemann@yale.edu

[‡]MIT Sloan School of Management, bonatti@mit.edu

[§]Yale University, nick.wu@yale.edu

1 Introduction

1.1 Motivation and Results

The allocation of digital advertisements is increasingly conducted through auctions with *automated bidding* agents. Advertisers provide high-level constraints (such as total spending or a return-on-investment target) to the digital platform that operates the auction, and the platform then determines specific bids for each impression and viewer via an automated (or algorithmic) bidding protocol. These approaches are now the dominant mode of advertiser interaction with major digital platforms like Google, Meta, and Amazon (Aggarwal et al. 2024). Most often, the auto-bidding systems are offered directly by the digital platforms that run the auctions, although they can also be provided as a service by third parties that manage bidding on advertisers' behalf. Advertisers embrace these protocols because they reduce the complexity of handling thousands or millions of keyword- or user-level bids. In turn, the platforms' bidding algorithms harness real-time consumer data and large-scale machine learning to optimize each advertiser's campaign.

Automated bidding protocols rely on the large datasets that digital platforms accumulate over time. These datasets consist of past impressions and advertiser–viewer interactions, which help identify and monetize the most valuable matches. The interplay between algorithmic and data-augmented bidding reveals two distinct gatekeeper roles for digital platforms. First, the platform acts as a competition gatekeeper among advertisers, with its auction algorithm determining which advertisers get listed and how they are ranked. Second, it serves as an information gatekeeper, using past information and current bids to allocate positions in a way that creates new matches and further refines the platform's data for future allocations. In essence, the proprietary, data-driven bidding system ties these two gatekeeping roles together. Trained on past impressions, the auto-bidding system creates new impressions in real time. From the advertiser's perspective, this setup produces a *managed advertising campaign* subject to the advertiser's stated constraints. By *managing* campaigns through bids and data, the platform effectively controls which products and viewers get matched.

The objective of the current paper is to develop the details of the auctions and auto-bidding protocols that provide the market-design foundations of a *managed advertising campaign*, thereby clarifying the separate role of the *auction engine* that powers the *matching engine*. We focus on algorithmic bidding protocols–also called auto-bidding agents–that let advertisers specify higher-level targets such as budgets or return-on-investment requirements and then delegate the actual bidding over many impressions and keywords to the auto-bidding agents. We characterize the properties of auto-bidding mechanisms that address the platform's design objectives. In particular, we show that such mechanisms preserve *standard auction logic*—for example, awarding an impression to the highest bidder under a second-price or first-price payment—while guaranteeing that each advertiser's total expenditure remains within its stated budget.

In earlier work, Bergemann & Bonatti (2024) and Bergemann, Bonatti & Wu (2025) studied how managed advertising campaigns monetize the match between viewer and product by selling impressions. We defined such campaigns as protocols that form matches based on past data and current bids, focusing primarily on the *matching design* and downplaying the specifics of the *auction design*. The central finding was that these centralized mechanisms enhance platform profits and push equilibrium product prices upward. However, advertisers' budgets appeared mainly as coarse, aggregate constraints that mapped directly to advertiser payoffs under the platform's chosen mechanism, without explicitly modeling the intermediate steps of bidding and bid-based ranking.

In Bergemann & Bonatti (2024), we focused on multi-product sellers, demonstrating that the optimal managed advertising campaign induces product steering that facilitates second-degree price discrimination. By leveraging data on past interactions, the platform can efficiently match products to viewers' preferences. Equipped with extensive consumer-level information, the platform tailors not only which advertiser is displayed to each user but also which product and price each viewer is offered.

In Bergemann, Bonatti & Wu (2025), our focus turned to single-product sellers, showing that the optimal managed advertising campaign delivers personalized product discounts—effectively a form of third-degree price discrimination. The principal goal of that work was to identify the managed advertising campaign maximizing the platform's revenue, subject to advertisers' participation constraints and, in particular, their ability to reach consumers on competing sales channels.

Here, in contrast, we develop an equilibrium model of these managed campaigns that incorporates both budgets and bidding strategies. Our framework explicitly ties each advertiser's budget to a formal bidding mechanism in which total spending is capped, yet each impression-level bid adjusts to user-specific factors. In other words, we connect (i) a *managed* *campaign*, which internally enforces budget constraints, with (ii) a visible *shadow-auction*, where an advertiser's effective bid is a function of its chosen budget. This formulation captures the core idea of *auto-bidding with pacing*: at each impression, the platform sets or scales the advertiser's bid to stay within the daily (or monthly) spend constraint while seeking to maximize that advertiser's return.

Formalizing the link between advertisers' budgets and real-time bidding explicitly is essential for two main reasons. First, nearly all major platforms implement budget-based auto-bidding, yet classical auction theory (without budgets) often struggles to predict advertiser behavior in modern, large-scale ad systems. Second, while earlier studies such as Bergemann, Bonatti & Wu (2025) largely treat the mapping from budgets to per-impression bids as a "black box," our mechanism models the intermediate *pacing* or *shadow-price* step directly. This fuller view shows precisely how an advertiser's total spending constraint translates into a continuum of value-dependent bids.

The perspective that we offer here departs from much of the recent work on auto-bidding agents in three main respects. First, we deliberately set aside the explicit dynamic dimension—how a budget is gradually spent over many auction events. Although budgets are stock variables and impressions arrive as a flow, we adopt a steady-state (implicitly static) view. The auto-bidding protocol must keep total expected spending below or equal to the declared budget over a given distribution of impressions, items, and payments. Second, we enlarge the standard notion of a bidding equilibrium to accommodate algorithmic bidding. The auto-bidding algorithm effectively lets each advertiser commit to a protocol mapping budgets into bids. In this context, the equilibrium notion must capture not only that the resulting bids are optimal given a budget, but also that advertisers choose their budget constraints optimally. Third, rather than focusing on a single exogenous bidding protocol, we investigate which auto-bidding protocol maximizes the digital platform's revenue, in contrast to much of the literature that analyzes protocols maximizing social surplus or advertisers' net surplus.

1.2 Related Literature

We study the role of bidding algorithms, or automated bidding strategies, in matching advertisers with viewers on digital platforms. The rising prevalence of auto-bidding in online advertising has spurred substantial research, particularly in theoretical computer science and operations research (see Aggarwal et al. (2024) for a survey). The typical approach is to design algorithms that generate bids subject to advertiser-defined constraints that apply collectively over numerous auction events. For instance, one may seek to maximize the advertiser's total value of matches given a budget constraint or a target cost per acquisition. The resulting algorithm then issues a sequence of bids across different times, websites, impressions, and users, striving to meet the global constraint. Because the arrival of impressions and viewers is random, the algorithm must react in real time with incomplete knowledge of future events (see, e.g., Balseiro et al. 2015, Balseiro & Gur 2019). Often, the best practical solutions involve approximate methods or performance bounds rather than an exact ex-post optimum.

Given the prevalence of bidding agents, a natural question is how best to design auctions when auto-bidders are present, as explored by Golrezaei et al. (2021) and Balseiro et al. (2021). Unlike those papers, which treat the algorithm's design as exogenous, we examine how to optimize the auction design when the seller (e.g., a platform like Google or Meta) controls both the auction rules and the bidding algorithms.

Although algorithmic approaches to auctions remain underexplored, algorithmic pricing in strategic settings has generated considerable discussion (Johnson et al. 2023, Calvano et al. 2020, Brown & MacKay 2023, Asker et al. 2024). Johnson et al. (2023) show how AI-driven pricing strategies on platforms can reward price cuts with increased exposure, boosting consumer surplus and undermining algorithmic collusion. Calvano et al. (2020) provide experimental evidence that repeated oligopoly settings with AI pricing can foster tacit collusion, which emerges from self-enforcing punishment-and-reward schemes. Brown & MacKay (2023) demonstrate that high-frequency pricing algorithms help online retailers sustain higher prices and amplify merger effects. Finally, Asker et al. (2024) distinguish between asynchronous learning (which may converge to monopoly prices) and synchronous learning (which is more likely to produce competitive pricing). They find that asynchronous learning, which updates only based on results of past actions, leads to pricing close to monopoly levels, whereas synchronous learning, which accounts for counterfactuals, results in more competitive pricing.

By contrast, our focus here as well as in the earlier contributions, Bergemann & Bonatti (2024) and Bergemann, Bonatti & Wu (2025), is on algorithmic tools provided by a central player, rather than separate algorithms independently introduced by multiple parties. This centralization potentially heightens the risk of collusion through coordination, even if each advertiser's algorithm nominally responds only to that advertiser's own data and inputs. Another key difference from much of the auto-bidding literature (and from studies of collusion by algorithm) is that we explicitly allow for competition through other sales channels outside the platform. Each advertiser thus retains an outside option, shaping how it assesses the platform's inside option and responds strategically to the central auto-bidding algorithm.

Our analysis also complements the literature on auto-bidding and auctions with budget constraints. Pai & Vohra (2014), for instance, demonstrate that standard auction formats often fail to maximize revenue under budget limitations, and they propose a revenue-maximizing mechanism that pools bidders strategically and distorts allocations to heighten competition. Crucially, the top bidder may not always be awarded the item outright, and standard formats (first-price or second-price) are generally suboptimal when buyers' budget constraints are private information.

Recent work continues to investigate how algorithmic decision-making and platform-driven interventions shape digital markets. Musolff (2024) studies algorithmic pricing in e-commerce and finds that although repricing tools may initially drive prices down, sellers eventually engage in strategic resets that raise prices, enabling tacit collusion. Likewise, Lee & Musolff (2021) examine Amazon's recommendation algorithms, showing that platform-guided search boosts both price elasticity and competition, even as it modestly favors Amazon's own products. Their results suggest that self-preferencing need not undermine consumer welfare; in fact, it can increase overall surplus by directing consumers to more suitable offers.

2 Model

We consider a single digital platform and J advertisers, indexed by $i = 1, \ldots, J$. A unit continuum of consumers visit the platform. Each consumer is characterized by a value profile

$$v \;=\; \big(v_1, v_2, \ldots, v_J\big) \;\in\; \mathbb{R}_+^J,$$

where v_i represents the consumer's willingness to pay for the product or service of advertiser i. Each value v_i is distributed identically and independently over an interval in $\mathbb{R}_+$ according to $F_i\left(v_i\right)$ and thus the joint distribution of v is given by

$$F\left(v\right) = \prod_i F_i(v_i).$$

The associated density function is given by $f_i(v_i)$.

Advertisers and Budgets Each advertiser i chooses a nonnegative budget $T_i \in \mathbb{R}_+$ to allocate on the platform. An advertiser's *total spending* cannot exceed this budget. After the advertisers set budgets $T = (T_1, \ldots, T_J)$, the platform determines how to bid for each consumer type v on behalf of each advertiser.

Mechanism There is a single sponsored slot available for each impression of a viewer. The platform operates a mechanism that translates each value and budgets profile (v, T) into a shadow bid $b_i(v, T)$ for advertiser i. Conditional on these bids, the platform awards the sponsored placement to the highest-bidding advertiser and charges the winner a payment. At this level of generality, we allow the bid of advertiser to be informed by the entire vector of values v. In the analysis, we start in Section 3 with a bidding algorithm that generates the bids for advertiser i with information about v_i alone. We denote by $p_i(v)$ the payment that advertiser i makes if it wins consumer v.

We require the platform's mechanism (i.e., bidding and payment functions) to be such that each advertiser i's payment does not exceed their submitted budget T_i. Thus, letting

$$S_i(T) \triangleq \{v : b_i(v, T) = \max_k b_k(v, T)\}$$

denote the set of consumers won by advertiser i, we require

$$\int_{S_i(T)} p_i(v)\, dF(v) \;\leq\; T_i. \tag{1}$$

Our leading example is that of a second price auction, where bids are nonetheless constrained by the budget and thus do not necessarily follow the standard logic of bidding equal to the value. In this case, the payment of the winning bidder is

$$p_i(v) = \max_{k \neq i} \{b_k(v, T)\}.$$

In another example, the platform chooses a bidding strategy for each advertiser to maximize their expected profits given the realized bid distribution of the other advertisers.

We abstract from the downstream pricing mechanisms and assume for now that the advertiser is able to extract the consumer's full willingness to pay. Hence, the ex-post payoff for advertiser i given consumer type v is

$$v_i \; - \; p_i(v, T),$$

if i is assigned consumer v, and zero otherwise.

The expected profit of advertiser i under the budget profile T and bid strategies b (which determine the market segments S), is then given by

$$\Pi_i(T_i, T_{-i}) \; = \; \int_{S_i(T)} \big(v_i - p_i(v, T)\big) \, dF(v).$$

Discussion of Assumptions and Real-World Auto-bidding Our modeling approach posits that an *auto-bidding* system runs the detailed auction-level decisions for each advertiser. In practice, platforms like Google, Meta, and Amazon offer advertisers a simplified interface where the advertiser specifies a *budget*, and the platform then determines individualized bids across a heterogeneous stream of potential impressions. This structure is consistent with our mechanism that translates (v, T_i) into an effective bid $b_i(v; T_i)$. Empirical evidence shows that most advertisers adopt such delegated solutions because it eliminates the need to micromanage bids across thousands or millions of auctions (see also the discussion in Aggarwal et al. 2024).

In line with Bergemann, Bonatti & Wu (2025), we let the platform *fully observe* each consumer's value profile v. In reality, modern advertising platforms use extensive consumer data to predict relevant signals, such as click propensity or expected conversions, which proxy for willingness to pay. Our model interprets these advanced machine-learning predictions as the $\{v_i\}$. Consequently, we allow the platform to condition each advertiser's (auto-generated) bid on the entire vector v. As Bergemann, Bonatti & Wu (2025) underscore, the ability to finely segment users and tailor bids or prices is a central feature of digital advertising markets. In a model of competing heterogeneous advertisers (and products) as analyzed here, Bergemann, Brooks & Morris (2025) show how the shape of the information can influence the distribution of the social surplus even in the absence of an intermediary platform.

Finally, each advertiser's ex-ante choice of a budget $\{T_i\}$ arises naturally from institutional features. Many real-world platforms (e.g. Google Ads or Facebook Ads) prompt advertisers to enter a daily or monthly spend cap, or to adopt a particular return-on-investment target. The platform's auto-bidding system then implements the campaign with continuous pacing (Aggarwal et al. 2024). This reduced-form approach aligns with how leading digital platforms convert high-level constraints ("target spend" or "target CPA") into a bidding policy.

3 Bidding with Budgets and Data

We begin by analyzing a second-price auction with a basic bidding algorithm that converts each bidder's chosen budget into a distribution of bids. This automated algorithm ensures that an advertiser's budget translates into bids across a range of auctions. In particular, it (i) maximizes the advertiser's value subject to the chosen budget and (ii) respects the budget constraint by keeping the advertiser's total expected payment across all auctions below its budget. Thus, the budget provided by the advertiser effectively supports a managed advertising campaign over multiple auctions. Given this specific bidding algorithm, we determine the optimal budget for an individual advertiser. We then determine the equilibrium choice of all bidders, where each bidder takes the distribution of bids as given. We call this the "bid equilibrium."

We next extend the analysis by examining the budgeting behavior of bidders who all use the same bidding algorithm. Under this assumption, we can explore a more advanced equilibrium concept. Specifically, if the strategic decisions lie in choosing budgets rather than individual bids, and if everyone employs the same auto-bidding procedure, then we can ask how those budgets are determined in equilibrium. We refer to this as the "budget equilibrium."

Because advertisers vary widely—some bid manually, others employ third-party services, and still others rely on the platform's own auto-bidding tools—the literature on auto-bidding has mostly focused on bid equilibria. However, as platform-provided bidding algorithms become increasingly popular, we want to understand how widespread adoption of a single class of algorithm affects bidding outcomes. As we will see, the budget equilibrium framework allows bidders to engage in more nuanced strategic behavior. As the auto-bidding algorithm commits each side to a specific mapping from budgets to bids, any particular budget choice by bidder j influences how other bidders' budgets translate into their bids.

3.1 Pacing Algorithm

In the pacing algorithm, the auction platform transforms the budget of an individual bidder into a distribution of bids that are adapted to the value realization v_i of the bidder. More formally, the algorithm chooses a bid function as function of the budget T_i and the realized value v_i for each auction event:

$$b_i : V_i \times T_i \to \mathbb{R}_+. \tag{2}$$

In order to describe the budget constraint for bidder i explicitly, we introduce the bid distribution by the other bidders, $-i$. In fact, it suffices to consider the distribution of the highest bid among the remaining bidders $-i$, denoted by $G_{-i}(b_{-i})$ The distribution has a support given by

$$\left[\underline{b}_{-i}, \overline{b}_{-i}\right],$$

with $0 \leq \underline{b}_{-i} < \overline{b}_{-i} < \infty$. The bid b_{-i} is the bid that bidder i has to beat in order to win the auction. If the other bidders had bid their true values, then the support of the bids would of course coincide with the support of the values, but we do not assume truthful bidding from the outset. We assume that this first order statistic has finite expectation:

$$\int b_{-i} dG_{-i}(b_{-i}) \triangleq \overline{B}_i < \infty.$$

The expectation $\overline{B}_i$ of the highest competing bid against bidder i is the maximal expected payment of bidder i in a second price auction.

Lemma 1 (Maximal Spend Budget)
Given a distribution of competing bids $G_{-i}(b_{-i})$, there is a maximal budget that any bidding algorithm can spend for bidder i

$$T_i = \overline{B}_i.$$

Proof. Given the rules of the second price auction, the payment of the winning bidder i is the highest bid among the competing bidders, b_{-i}. By bidding sufficiently high, bidder i wins all impressions and pays the second-price each time; hence the maximum feasible budget to be fully spent is indeed the expectation of that highest competing bid, determined by the distribution $G_{-i}(b_{-i})$. ■

With a budget of $\overline{B}_i$ bidder i wins all auction events but the outcome may be less than desirable for bidder i since they may bid for an item more than the value v_i that they assign to the item. Now for any budget $T_i < \overline{B}_i$, a bidding algorithm determines what is the least expensive way to bid to satisfy the budget constraint while maximizing the *net utility* of bidder i:

$$\max_{b_i(\cdot,T_i)} \int_0^{\infty} \int_0^{b_i(v,T_i)} (v_i - b_{-i})\, g_{-i}(b_{-i})\, f_i(v_i)\, dv_i \tag{3}$$

subject to the budget constraint given by

$$\int_0^{\infty} \int_0^{b_i(v_i,T_i)} b_{-i} g_{-i}(b_{-i})\, f_i(v_i)\, dv_i \le T_i. \tag{4}$$

Lemma 2 (Optimal Pacing Algorithm)
Given a distribution of competing bids $G_{-i}(b_{-i})$, the bidding algorithm which maximizes the bidder's net utility is given by constant fraction of the value:

$$b_i(v_i, T_i) = k_i \cdot v_i.$$

Proof. This is a constrained optimization problem which can be solved pointwise for every v_i of the bidder identifying the optimal choice $b(v_i, T_i)$. The resulting first order condition is:

$$(v_i - b_i(v_i, T_i))\, g_{-i}(b_i(v_i, T_i)) - \lambda b_i(v_i, T_i)\, g_{-i}(b_i(v_i, T_i)) = 0\,,$$

where $\lambda \in \mathbb{R}_+$ is the Lagrange multiplier associated with the budget constraint. Thus we obtain that

$$b_i(v_i, T_i) = \frac{1}{1+\lambda} v_i, \quad \text{for all } v_i, \tag{5}$$

which establishes the result. ■

The boost or pacing parameter k_i is chosen such that the expected payment in the second price auction equals the committed budget T_i:

$$T_i = \int_0^{\infty} \left(\int_0^{k_i v_i} b_{-i} dG_{-i}(b_{-i}) \right) dF_i(v_i)\,. \tag{6}$$

Thus, the budget T_i allows bidder i to bid $k_i v_i$ for a item of value v_i. In consequence, bidder i wins as long as the bid $k_i v_i$ which defines the upper integration limit is the highest realized bid in the auction. Moreover, in the second price auction, the expected payment of such a winning bid is simply the expectation of the highest losing bid in the integral between 0 and $k_i v_i$.

The budget constraint is thus interpreted as a constraint in terms of an ex-ante expectation. The committed budget supports a bidding strategy

$$b_i(v_i, T_i) = k_i(T_i)\, v_i.$$

across a distribution of values v_i and v_{-i} such that in expectation the budget is sufficient to fund the payments in the winning auction. We refer to the realization of a particular auction with realized values v as an auction event. The budget is being spent over a notional time horizon over which the values of the advertisers are realized. Thus, the bidding strategy defines an interim bid, conditional on the value realization v_i that is consistent with an ex-ante expectation of auction payments.

Lemma 3 (Bid Monotonicity)
The bidding boost $k_i(T_i)$ is monotone increasing in the budget T_i for all $T_i \leq \overline{B}_i$.

Proof. By the budget constraint equation (6) a larger T_i can support a larger upper bound of integration, which is given by $k_i v_i$. ■

We assume that submitted budgets T_i larger than $\overline{B}_i$ are executed by the platform but returning funds to the bidder in case the platform fails to spend the budget. Given the bid distribution $G_{-i}(b_{-i})$ there is a budget that allows bidder i to win if and only if his bid is less than his value, in other words a paced bidding strategy that sets $k_i = 1$:

$$T_i(G_{-i}) = \int_0^{\infty} \left(\int_0^{v_i} b_{-i} dG_{-i}(b_{-i}) \right) dF_i(v_i). \tag{7}$$

Each bidder i who avails themselves of the auto-bidder is maximizing the net utility from their budget commitment. The net utility can be written as

$$\int_0^{\infty} v_i \left(\int_0^{b_i(v_i, T_i)} dG_{-i}(b_{-i}) \right) dF_i(v_i) - T_i,$$

where by Lemma 1 it is sufficient to consider budgets $T \leq \overline{B}_i$. The problem for bidder i is then to find the budget that maximizes the above expression, thus

$$\max_{T_i \in [0, \overline{B}_i]} \left\{ \int_0^\infty v_i \left(\int_0^{b(v_i, T_i)} dG_{-i}(b_{-i}) \right) dF(v_i) - T_i \right\}. \tag{8}$$

We can now identify the optimal choice of budget T_i for advertiser i when the choice of budget T_i is only constrained by their objective of maximizing the net utility given by (8).

Proposition 1 (Optimal Budget Strategy)
The optimal budget for bidder i is to set $T_i^* = T_i(G_{-i})$ and the optimal budget supports truthful bidding

$$b_i(v_i, T_i^*) = v_i \Leftrightarrow k_i^* = 1.$$

Proof. We remove the constraint on the budget in the choice of the optimal bidding strategy in Lemma 2. This amounts to reducing the value of the Lagrange multiplier $\lambda = 0$ which delivers the result. ■

We can now formally define the bidding equilibrium.

Definition 1 (Bidding Equilibrium)
A bidding equilibrium is a profile of budgets $T = (T_1, ..., T_J)$ and associated bid distributions $G_i(b_{-i})$ such that every bidder i chooses their net-utility maximizing budget T_i and each generated bid b_i is a best response against the realized bids b_{-i} of the other bidders given the budget constraints.

The definition of a bidding equilibrium applies even if only some bidders adopt budget-based strategies, while others submit manual bids or use different algorithms. Specifically, the equilibrium concept only requires that those bidders who use auto-bidding choose their budgets optimally. All remaining bidders are simply assumed to satisfy ex-post optimality of their spending. This flexibility in defining equilibrium notions likely explains why such frameworks (with small variations) dominate the study of auto-bidding algorithms in the literature.

In digital auctions, the participants typically bid with a diversity of bidding strategies. Some bidders deploy their proprietary algorithms, other employ auto-bidding algorithms by the auction platform or third parties. Even if all bidders deploy auto-bidding algorithms from the same platform or service provider, their bidding strategy may differ as the define their budget and their requirements over different objects, time horizons and constraints. Nonetheless, we can ask what the equilibrium might look in an environment in which all bidders avail themselves to the same budget pacing algorithm.

Proposition 2 (Bidding Equilibrium)
The unique symmetric bidding equilibrium is given by:

$$b_i\left(v_i, T_i^*\right) = v_i \Leftrightarrow k_i^* = 1,$$

with

$$T_i^* = \int_0^\infty \left(\int_0^{v_i} s dF^{J-1}(s) \right) dF_i(v_i).$$

Proof. The benefit of bidder i with a budget T_i against a bid distribution by the remaining bidders given by $G_{-i}(b_{-i})$ in a second price auction is:

$$\int_0^\infty \left[\int_0^{b_i(v_i,T_i)} (v_i - b_{-i}) \, dG_{-i}(b_{-i}) \right] dF_i(v_i). \tag{9}$$

For every v_i, the optimal bidding strategy is therefore to choose T_i in such a manner that $k_i = 1$. This allows the bidder to maximize the expected return against the competing bidders for every v_i:

$$\left[\int_0^{v_i} (v_i - b_{-i}) \, dG_{-i}(b_{-i}) \right].$$

Now, if the other bidders are bidding truthfully, meaning that $b_j(v_j) = v_j$ for all v_j and j, then the budget by bidder i to support $k_i = 1$ is given by:

$$\begin{aligned} T_i &= \int_0^\infty \left(F^{J-1}(v) \frac{\int_0^v s dF^{J-1}(s)}{\int_0^v dF^{J-1}(s)} \right) dF_i(v_i) \\ &= \int_0^\infty \left(\int_0^v s dF^{J-1}(s) \right) dF_i(v_i) \end{aligned} \tag{10}$$

Namely, the expected payment of bidder i to support bidding his value is simply the expectation of his payments in the second price auction. ■

Note that in a symmetric equilibrium, the discount k_i is the same for all bidders, thus $k_i = k$, and therefore the resulting allocation is efficient.

Now if the competing bidders bid truthfully, or $b_j(v_j) = v_j$, then the individual bidder needs to provide a budget large enough to displace the bids of the competing bidders. If the displacement budget is supposed to attain the efficient placement of bidder i then the budget needs to pay to displace the less efficient bidders, or

$$T_i = \int_0^\infty \left(\int_0^v s dF^{J-1}(s) \right) dF_i(v_i). \tag{11}$$

We refer to this as the efficient displacement budget.

3.2 Budget Equilibrium

In the pacing equilibrium, each bidder submitted a budget T_i which generated bids b_i as a best response to (i.e., holding fixed) the realized bids b_{-i} of the other bidders.

We now extend our notion of competition in budgets by allowing the platform to condition each advertiser's bid on the budgets submitted by all bidders. In particular, we maintain our focus on the uniform pacing environment, where $b_i(v, T) = k_i v_i$, but we assume that all $k_i(T)$ are chosen by the platform as a function of all the submitted budgets so that each bidder's aggregate payment is exactly T_i.

Formally, we let $\{k_i(T)\}_{i=1,..,I}$ denote the solution to the system of equations

$$\int_{S_i(T)} \max_{j \neq i} k_j v_j dF(v) = T_i \quad \text{for all } i, \tag{12}$$

where

$$S_i(T) \triangleq \{v : k_i v_i = \max_j k_j v_j\}.$$

In other words, the algorithm sets the boosts for all bids in such a way that each bidder exhausts the budget they submitted.

Under this mechanism, the direct effect of bidder i raising their submitted budget T_i is typically to increase their boost k_i. It will also have an ambiguous effect on the boost of the other bidders. On the one hand, bidder i needs to spend more, so k_j may increase in order to raise the losing bids. On the other hand, all budgets T_j are unchanged, and if k_i increases, so will the payments of all other bidders in the auctions they still win. Thus, k_j may have to decrease.

In the remainder of this section, we specialize to the case of two bidders ($J = 2$) and a symmetric equilibrium. We begin by rewriting the expected payment (which is equal to the budget) of each bidder as

$$T_i = \int_0^\infty \int_0^{\frac{k_i v_i}{k_j}} k_j v_j dF(v_j)\, dF(v_i),$$

and the expected payoff of each bidder as

$$V_i(T) = \int_0^\infty \int_0^{\frac{k_i(T)v_i}{k_j(T)}} dF(v_j)\, v_i\, dF(v_i) - T_i,$$

where $k_i(T)$ and $k_j(T)$ solve the budget spending constraint above.

We characterize the symmetric budget equilibrium in Proposition 3, which is proved in the Appendix.

Proposition 3 (Budget Equilibrium) *There exists a unique symmetric budgets equilibrium. In this equilibrium, each firm i submits a budget*

$$T_i = T^* := \frac{\int_0^\infty v_i^2 f(v_i)^2\, dv_i \int_0^\infty v_i(1-F(v_i))f(v_i)\, dv_i}{2\int_0^\infty v_i^2 f(v_i)^2\, dv_i - \int_0^\infty v_i(1-F(v_i))f(v_i)\, dv_i} > 0$$

and their bids receive a boost

$$k_i = k^* := \frac{T^*}{\int_0^\infty v_i(1-F(v_i))f(v_i)\, dv_i}.$$

From this statement, it is immediate to see that the submitted budget exceeds the total payment in a second-price auction if and only if $k^* > 1$. In turn, the equilibrium boost factor exceeds one if and only if $\partial k_j/\partial T_i < 0$, i.e., if one firm raising its budget above its rival's level induces a *less aggressive* response. We then have the following corollary of our equilibrium characterization.

Corollary 1 (Revenue Comparison) *The revenue in the two-bidder symmetric budget equilibrium exceeds the revenue in the manual-bidding second-price auction if and only if $\partial k_j(T,T)/\partial T_i < 0$, i.e.,*

$$\int_0^\infty (1 - F(v_i) - v_i f(v_i)) v_i f(v_i)\, dv_i \geq 0.$$

This condition suggests that fat-tailed distributions yield an equilibrium in budgets that is more profitable than manual bidding. In fact, it is easy to show that $k^* = 1$ for any exponential distribution and that $k^* < 1$ for any power distribution on $[0,1]$. Furthermore, numerical results suggest that for a lognormal distribution with parameters (μ, σ), we have $k^* < 1$ for σ low enough and $k^* > 1$ for σ high enough.

Figure 1 illustrates this finding and displays values v, equilibrium bids $b(v) = k^*v$ and expected payments for type v conditional on winning, i.e.,

$$p(v) = \frac{\int_0^v k^*v\, dF(v)}{F(v)}.$$

In panel 1a, the winning bids (and therefore the expected payments) are below the winner's valuation. Conversely, in panel 1b, the winning bid is always larger than the winner's valuation and, for low enough v, so is the expected conditional payment. In other words, the platform spends the advertisers' budgets in a way that induces losses on low types and cross subsidizes them with gains on higher types.

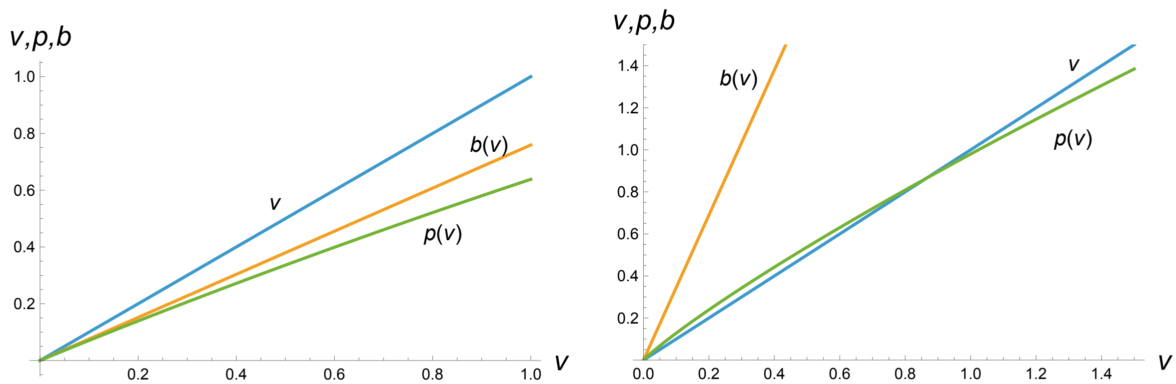

(a) $\mu = 1, \sigma = 1/2, k^* \approx 0.76$ (b) $\mu = 1, \sigma = 3, k^* \approx 3.45$

Figure 1: Values, bids, and expected payments–Lognormal distribution.

4 Optimal Mechanisms

In the previous sections, the platform was limited to running a second price auction with bids generated from budgets according to a simple class of (pacing) rules.

In this section, we maintain the second price format, but we now consider mechanisms that allow for arbitrary bidding functions $b_i : V_i \times T_i \to \mathbb{R}_+$. We first characterize the optimal bidding function in our baseline setting and then extend our results to a setting with parallel sales channels for each advertiser.

4.1 Vertical Integration Surplus

In this subsection, we show that when not restricted to pacing mechanisms, the platform can select a bidding rule that extracts the vertical integration surplus from the bidders. To that end, define the vertical-integration surplus as

$$\Pi_V \triangleq \int_0^\infty v dF^J(v). \tag{13}$$

We begin with a result showing how some bidding function induces an outcome where the platform earns all the surplus, even if the allocation rule is continuous and must always induce nonnegative profits for any budget offered.

Proposition 4 (Vertical Integration)

1. *There exists a bid function that implements the vertical integration solution and extracts all the bidders' surplus, induces continuous outcomes in the budget, and offers nonnegative profits regardless of budget submitted. Furthermore, any such bid function induces identical profits for each bidder and each budget submitted.*

2. *If a bidder chooses a lower budget than Π_V/J, the platform-optimal bid function submits the equilibrium bid with a probability proportional to the budget submitted and a zero bid otherwise.*

The complete proof of this result is in the Appendix. To establish the first claim, we show that there exists a bid function that grants all advertisers a payoff that is non-negative and continuous in their budget, and which generates revenue JT^* for the platform in equilibrium. The idea is to identify a bidding strategy for each bidder $b_i(v)$ that inflates the advertisers' bids so that, in a second-price auction, the winning firm pays their full value, $\arg\max_j v_j$. The following is an example of such a bid function

$$b_i(v) = \begin{cases} v_i + \epsilon & \{i\} = \arg\max_k v_k \\ v_i & i \in \arg\max_k v_k \text{ and } \exists j \neq i,\ j \in \arg\max_k v_k \\ v_j & j \in \arg\max_k v_k \text{ and } i \notin \arg\max_k v_k. \end{cases}$$

Under this bidding mechanism, if an advertiser submits the maximal budget T^*, their total spend is

$$T^* = \Pi_V/J = \int vF^{J-1}(v)dF(v).$$

Thus, the advertiser obtains zero net profits. If they submit any smaller budget, the platform will randomly and uniformly drop some of their bids down to zero, so that the entire budget is spent and the advertiser's net payoff remains nil.

Figure 2 illustrates the realized bids if two bidders submit budgets $T_1 < T^*$ and $T_2 = T^*$.

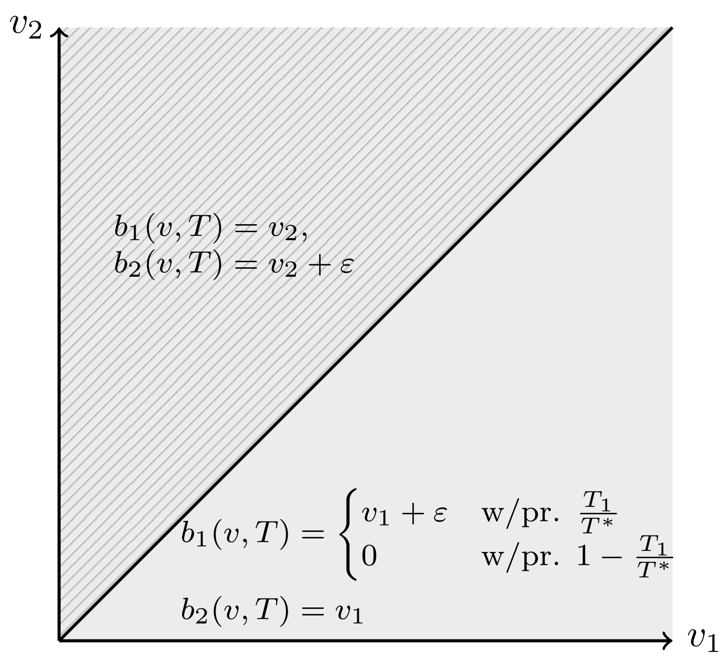

Figure 2: Realized Bids for Budgets $T_1 < T^*$ and $T_2 = T^*$.

4.2 On- and Off-Platform Markets

We now extend Proposition 4 to a model where advertisers also have access to an off-platform market. Following the setup of Bergemann, Bonatti & Wu (2025), suppose there is a unit mass of consumers which is divided into markets. Each advertiser controls a random "captive" set of consumers that only they can serve at a posted price, with total measure $(1-\lambda)/J$ per advertiser. Meanwhile, the on-platform market of measure λ is served via the auction. Advertisers must set a budget for the on-platform auctions and a posted price for their own off-platform market, knowing consumers can always choose the off-platform option if it is more attractive.

Thus, each advertiser selects a posted price $\bar{P}_i$ for their product in their off-platform market. On the platform, the platform sets prices as a function of the vector of the consumer's value, the budgets submitted, and off-platform prices. We suppose consumers on the platform can shop off the platform if they wish, so we additionally have the showrooming constraint, which requires that the platform never offers i's product at a price higher than their off-platform price $\bar{P}_i$.

The extensive form game is then

1. The platform proposes a mechanism (b, P) to all firms. The first component $b : V^J \times \mathbb{R}_+^J \times \mathbb{R}_+^J \to \mathbb{R}_+$ is a *bidding strategy*, and it maps consumer value, budgets, and off-platform prices into a bid. The second component $P : V^J \times \mathbb{R}_+^J \times \mathbb{R}_+^J \to \mathbb{R}_+$ is a *pricing policy*, which maps consumer value, budgets, and off-platform prices into a price to be offered on the platform.

2. The firms simultaneously decide how much budget to offer and what prices to post off the platform.

3. Advertising auctions are run; bids are chosen according to b and the submitted budgets, and the prices of products are set according to the policy P.

Bergemann, Bonatti & Wu (2025) show that the platform can attain its maximum profit by using a forcing contract, or discontinuous allocation at the threshold required budget. Under this forcing contract, the joint profit of the platform and the advertiser is maximized, but each advertiser is held to their outside option (i.e., the profit they could attain by only selling off-platform), and the platform appropriates all of the producer surplus it creates. More precisely, the pricing problem of the platform-advertisers integrated firm is given by

$$\Pi_C(P) \triangleq \left\{(1-\lambda)P(1-F(P)) + \lambda \int_0^\infty \min(v, P)\ dF^J(v)\right\}.$$

Let Π^* denote the maximized value; $\Pi^* \triangleq \max_P \Pi_C(P)$, and so Π^* is the vertical integration surplus in this environment. Bergemann, Bonatti & Wu (2025) show that with the forcing contract, the platform obtains $\Pi^* - J\Pi_O$, where Π_O is the outside option of a given advertiser:

$$\Pi_O = \max_P \left\{\frac{1-\lambda}{J} P(1-F(P)) + \lambda \int_P^\infty P F^{J-1}(v_j - P)\ dF(v_j)\right\}. \tag{14}$$

Note that (14) defines the price-setting problem where the firm can only sell via its posted price. In the next result, we show that the forcing contract of Bergemann, Bonatti & Wu

(2025) is not necessary; rather, the platform can appropriate the vertical-integration surplus up to the outside options of the advertisers with an appropriate bidding strategy and pricing policy, even if the allocation rule must be continuous and induce non-negative profits for any budget offered.

Proposition 5 (Full Extraction with Off-Platform Market)

1. *The platform can attain the optimal vertical integration surplus by having the bidders select budgets.*

2. *If bidder i does not submit the required budget, the platform-optimal bid function submits the equilibrium bid with probability chosen as a function of the budget submitted, and a zero bid otherwise.*

3. *The probability of a nonzero bid is not proportional to the budget but nevertheless induces a linear best-response profit function $\Pi_i(T_i)$ for the bidders.*

Here, the value of winning an auction is now disciplined by the pricing decisions made by the advertisers. However, the platform's optimal policy looks very similar to that of Proposition 4—the platform inflates bids so that the bidder pays their effective value (which depends on the off-platform prices).

Figure 3 below illustrates the expected bids made in equilibrium and the expected prices charged to consumers for an example with two bidders, uniformly distributed values, and off-platform prices $\bar{P}_1 = \bar{P}_2 = 1/3$. For each advertiser i, the expected prices charged in equilibrium to a consumer with valuation v_i are given by $p(v_i) = \min\{v_i, 1/3\}$, which also equals the advertiser's expected payment conditional on winning the auction. The expected bids exceed the expected payment for sufficiently low-value consumers, as with high probability, these bids serve the goal of raising the competitor's payment. In this example, they are given by $b(v_i) = \min\{1/3, (9v_i^2 + 5)/18\}$.

Similarly, in order to induce the advertisers to submit the optimal budgets, the platform makes zero bids with some probability whenever an advertiser submits a smaller budget, and designs this probability such that an advertiser cannot profit by submitting a smaller budget. More precisely, if the platform submits a nonzero bid with probability $q(t)$ whenever the budget submitted is t, the revenue of an advertiser who submits t and sets price P is

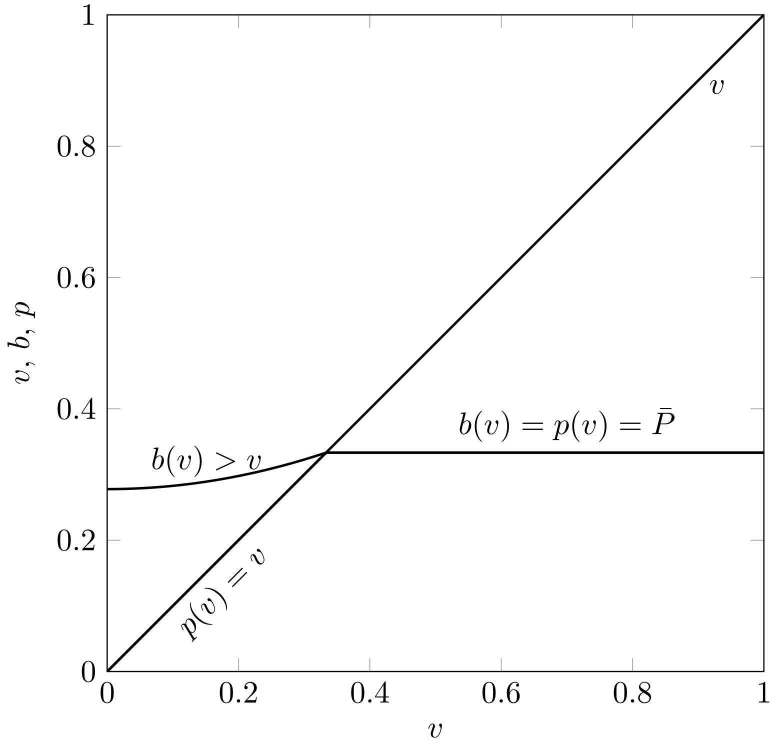

Figure 3: Values, Expected Bids, and Expected Payments ($F(v) = v$, $\bar{P} = 1/3$).

given by $\Pi_i(t, P)$ which we can write explicitly as:

$$\frac{1-\lambda}{J} P(1-F(P)) + \lambda \left((1-q(t)) \int_P^\infty P F^{J-1}(v-P)\ dF(v) + q(t) \int_0^\infty \min(v, P) F^{J-1}(v) dF(v) \right).$$

In the environment of Proposition 4, the revenue from submitting a smaller budget equaled the budget spent; with the off-platform market, the revenue *maximized over off-platform pricing choices* must equal the budget spent; hence, the probability $q(t)$ is nonlinear in t.

5 Conclusion

In this paper, we examined the implications of budget-constrained auto-bidding in digital advertising auctions. We study the implications of the pacing algorithms used by a platform, their implications on the equilibrium budgets provided by advertisers, and characterize the extent to which platforms can extract surplus in equilibrium by optimizing auto-bidding mechanisms.

An important avenue for further exploration is the role of informational asymmetries between advertisers and platforms. Platforms have extensive data on user behavior and market conditions, which they use to shape bidding strategies and pricing mechanisms. Thus, understanding how a platform's informational role shapes budget-setting behavior could yield further insights into the equilibrium consequences of platform behavior. Additionally, we did not model competition among platforms; it would be interesting for future work to explore how competitive pressures between platforms shape auto-bidding algorithms and advertising spending; if advertisers allocate budgets across multiple platforms, how do competing platforms adjust their auto-bidding mechanisms to attract more spending?

By relating budget-constrained bidding to equilibrium platform outcomes, our work provides a foundation for understanding the strategic consequences of auto-bidding in digital advertising markets. We hope these insights inspire further research on the economics of algorithmic market design, platform competition, and the evolving structure of digital advertising.

Appendix

Proof of Proposition 3. We first fully differentiate V_i with respect to T_i and we impose $k_i = k_j = k > 0$. We then obtain

$$\frac{\partial V(T,T)}{\partial T_i} = \frac{A - 2Ak + Bk}{2Ak - Bk},$$

where

$$\begin{aligned} A &= \int_0^\infty v_i^2 f(v_i)^2 \, dv_i \\ B &= \int_0^\infty v_i(1 - F(v_i)) f(v_i) \, dv_i \\ k &= T/B. \end{aligned}$$

Solving for k yields the result. Finally, to see that the equilibrium budget T^* is strictly positive, we show that the denominator of T^* is positive. Consider the expression:

$$\int_0^\infty (1 - F(v)) v f(v) \; dv.$$

Note that

$$\frac{d}{dv}\left[-\frac{1}{2}(1 - F(v))^2\right] = (1 - F(v)) f(v),$$

and so we can integrate by parts:

$$\int_0^\infty (1 - F(v)) v f(v) \; dv = \left[-\frac{v}{2}(1 - F(v))^2\right]_0^\infty + \int_0^\infty \frac{1}{2}(1 - F(v))^2 \; dv = \frac{1}{2}\int_0^\infty (1 - F(v))^2 \; dv. \tag{15}$$

Additionally, we can also apply the Cauchy-Schwarz inequality to get

$$\begin{aligned} \int_0^\infty (1 - F(v)) v f(v) \; dv &\leq \left(\int_0^\infty (1 - F(v))^2 \; dv\right)^{1/2} \left(\int_0^\infty (v f(v))^2 \; dv\right)^{1/2}, \\ &= \left(\frac{1}{2}\int_0^\infty (1 - F(v))^2 \; dv\right)^{1/2} \left(\int_0^\infty 2(v f(v))^2 \; dv\right)^{1/2}. \end{aligned}$$

By (15), the first term on the RHS can be substituted, so this becomes

$$\int_0^\infty (1 - F(v)) v f(v) \; dv \leq \left(\int_0^\infty (1 - F(v)) v f(v) \; dv\right)^{1/2} \left(\int_0^\infty 2(v f(v))^2 \; dv\right)^{1/2}.$$

Then dividing out the term $\left(\int_0^\infty (1-F(v))vf(v)\;dv\right)^{1/2}$ from both sides (which is positive) we get

$$\left(\int_0^\infty (1-F(v))vf(v)\;dv\right)^{1/2} \leq \left(\int_0^\infty 2(vf(v))^2\;dv\right)^{1/2}.$$

Squaring both sides and rearranging, we get

$$2\int_0^\infty v^2 f(v)^2\,dv - \int_0^\infty v(1-F(v))f(v)\,dv \geq 0.$$

and hence $T^* > 0$. ■

Proof of Proposition 4. Let $T^* = \int vF^{J-1}(v)dF(v)$ denote the value a bidder gets from winning the consumers who have the highest value for that bidder. The vertical integration surplus is equal to JT^*. We first show that there exists a bid function that grants bidders a nonnegative payoff that is continuous in their budget and generates equilibrium revenues JT^* for the platform. Let $b_j(v)$ be a bidding strategy that inflates bids such that the winning bidder always pays their value. The following is an example of such a bid function:

$$b_i(v) = \begin{cases} v_i + \epsilon & \{i\} = \arg\max_k v_k \\ v_i & i \in \arg\max_k v_k \text{ and } \exists j \neq i,\; j \in \arg\max_k v_k \\ v_j & j \in \arg\max_k v_k \text{ and } i \notin \arg\max_k v_k \end{cases}.$$

Note that such a function need not be unique; in the above example, any $\epsilon > 0$ suffices. Observe that the constructed b_i has the feature that when all bidders bid according to b_i, bidder i will always end up winning and paying v_i for every consumer such that $i \in \arg\max_k v_k$; since every bidder wins exactly the set of consumers that prefer them most and pays their value, if $b_i(v)$ is implemented, the total auction payments by all bidders generated is precisely JT^*.

It remains to show that such a b_i can be obtained in equilibrium when the bidders select budgets and the platform turns them into bids. Consider then the following strategy: given a budget t, the platform bids 0 with probability $1 - \max(1, t/T^*)$ and bids according to $b(\cdot)$ with probability $\max(1, t/T^*)$. Note that by bidding T^*, the bidder's induced profit is T^*; the bidder earns exactly T^* in value from the auctions. If a bidder bids $t < T^*$, the bidder's induced profit is t; they earn T^* if they are randomly chosen to bid "correctly" with probability t/T^*, and 0 otherwise. Hence, the bidder's profit function is

$$\Pi_i(t) = \begin{cases} t & t \in [0, T^*) \\ T^* & t \geq T^* \end{cases} \tag{16}$$

It is clear to see that regardless of the budget choice, the bidder gets net zero; by offering any $t < T^*$, the bidder spends their entire budget but earns exactly their budget: $\Pi_i(t) = t$. Similarly, by offering any $t \geq T^*$, the bidder spends T^* and makes profit exactly T^*, and hence earns zero. Thus, it is clear that the bidder's profit is nonnegative and continuous in their budget offered, weakly maximized at T^*.

To show uniqueness of this equilibrium, note that since the vertical integration surplus is at most JT^*, it is impossible for bidders to earn a positive surplus on expectation if they pay T^*. Hence, their profit function $\Pi_i(T^*) \leq T^*$. Further, bidding 0 earns 0, so $\Pi_i(0) = 0$. Since the profit function must extract all bidder surplus, each bidder must find it weakly optimal to submit T^*, so it must be the case that $\Pi_i(T^*) - T^* \geq \Pi_i(0)$, and so this is only possible if $\Pi_i(T^*) = T^*$. Since the profit function must induce a nonnegative payoff for any $t < T^*$, it follows that $\Pi_i(t) \geq t$ for any $t \in [0, T^*]$. But since the bidder must find it weakly profitable to submit T^*, $\Pi_i(T^*) - T^* \geq \Pi_i(t) - t$, so $t \geq \Pi_i(t)$. Combining these, together with the fact that the profit cannot exceed T^* in expectation, implies that (16) is the unique profit function that implements vertical integration solution and extracts all the bidder surplus. ■

Proof of Proposition 5. We claim that under an appropriate choice of bidding strategy and pricing policy, the platform can induce the advertisers to each offer budget $T^* = \Pi^*/J - \Pi_O$. Define the best-value price of firm j for some consumer v as follows, given off-platform prices $\bar{P}$:

$$p_j(v, \bar{P}) = \min(v_j, \bar{P}_j, \min_{k \neq j}(v_j - v_k + \bar{P}_k))_+.$$

Note that P_j is the maximum price that advertiser j could extract after winning consumer v. Then analogously to Proposition 4, define the bidding strategy:

$$b_i(v, \bar{P}) = \begin{cases} P_i(v, \bar{P}) + \epsilon & \{i\} = \arg\max_k v_k, \\ P_i(v, \bar{P}) & \{i\} = \arg\max_k v_k \text{ and } \exists j \neq i,\ j \in \arg\max_k v_k, \\ P_j(v, \bar{P}) & j \in \arg\max_k v_k \text{ and } i \notin \arg\max_k v_k. \end{cases}$$

Consider a platform policy that, given a budget T_i from i, randomizes between bidding $b_i(v, \bar{P})$ and setting the best-value price $P_j(v, \bar{P})$ with probability $q(T_i)$, and bidding zero with complementary probability. Note that the bidder's induced profit function depends on their submitted budget and their posted price. That is, the bidder's profit is

$$\Pi_i(t,P) = \tfrac{1-\lambda}{J} P(1-F(P)) + \\ \lambda\left((1-q(t))\textstyle\int_P^\infty PF^{J-1}(v-P)\,dF(v) + q(t)\int_0^\infty \min(v,P)F^{J-1}(v)dF(v)\right)$$

Define the maximized expression

$$\Pi_i^*(t) \triangleq \max_P \Pi_i(t,P).$$

If $q(t)=0$, then $\Pi_i^*(t)=\Pi_{O,j}$, as the profit function collapses to the same as the outside option in (14). Similarly, if $q(t)=1$, then $\Pi_i^*(t)=\Pi^*$, since the profit function collapses into the same objective as Π_C. Further, if Π_i^* is weakly concave and induces the advertisers to submit T^* on-path, it must be linear in the budget from 0 to T^*. We claim we can construct a q that satisfies this. By the envelope theorem, we must have

$$\frac{\partial}{\partial t}\Pi_i^* = \lambda q'(t)\left(\int_0^\infty \max(v,P^*(t))F^{J-1}(v)dF(v) - \int_{P^*(t)}^\infty P^*(t)F^{J-1}(v-P^*(t))\,dF(v)\right)$$

where $P^*(t)\in \arg\max_P \Pi_i(t,P)$. Observe that the parenthesized part is positive. For the profit function to be linear, this derivative must be constant, and equal to 1. Therefore, q must solve the differential equation:

$$q'(t) = \frac{1}{\lambda\left(\int_0^\infty \max(v,P^*(t))F^{J-1}(v)dF(v) - \int_{P^*(t)}^\infty P^*(t)F^{J-1}(v-P^*(t))\,dF(v)\right)}$$

with condition $q(0)=0$. Note that q does not appear on the right-hand side, and hence given the initial condition, we can integrate to obtain the unique solution q^*. By construction, it follows that under q^*, $\Pi_i^*(t)=\Pi_O+t$, and q^* is increasing since q' is positive. Further, at T^*, the total surplus attains the vertical integration benchmark, and hence we must have $q(T^*)=1$. Finally, q^* is a well-defined probability, because q^* is increasing, equal to 0 when no budget is submitted and equal to 1 when T^* is submitted. ■

References

Aggarwal, G., Badanidiyuru, A., Balseiro, S. R., Bhawalkar, K., Deng, Y., Feng, Z., Goel, G., Liaw, C., Lu, H., Mahdian, M. et al. (2024), 'Auto-bidding and auctions in online advertising: A survey', *ACM SIGecom Exchanges* **22**(1), 159–183.

Asker, J., Fershtman, C. & Pakes, A. (2024), 'The impact of artificial intelligence design on pricing', *Journal of Economics & Management Strategy* **33**(2), 276–304.

Balseiro, S., Besbes, O. & Weintraub, G. (2015), 'Repeated auctions with budgets in ad exchanges: Approximations and design', *Management Science* .

Balseiro, S., Deng, Y., Mao, J., Mirrokni, V. & Zuo, S. (2021), 'The landscape of auto-bidding auctions: Value versus utility maximization'.

Balseiro, S. & Gur, Y. (2019), 'Learning in repeated auctions with budgets: Regret minimization and equilibrium', *Management Science* **65**, 3952–3968.

Bergemann, D. & Bonatti, A. (2024), 'Data, competition, and digital platforms', *American Economic Review* **114**(8), 2553–2595.

Bergemann, D., Bonatti, A. & Wu, N. (2025), 'How do digital advertising auctions impact product prices?', *Review of Economic Studies* .

Bergemann, D., Brooks, B. & Morris, S. (2025), 'On the alignment of consumer surplus and total surplus under competitive price discrimination', *American Economic Journal: Microeconomics* **forthcoming**.

Brown, Z. Y. & MacKay, A. (2023), 'Competition in pricing algorithms', *American Economic Journal: Microeconomics* **15**(2), 109–56.

Calvano, E., Calzolari, G., Denicolò, V. & Pastorello, S. (2020), 'Artificial intelligence, algorithmic pricing, and collusion', *American Economic Review* **110**(10), 3267–97.

Golrezaei, N., Lobel, I. & Paes Leme, R. (2021), Auction design for roi-constrained buyers, *in* 'Proceedings of the Web Conference 2021', WWW 21, ACM.

Johnson, J. P., Rhodes, A. & Wildenbeest, M. (2023), 'Platform design when sellers use pricing algorithms', *Econometrica* **91**(5), 1841–1879.

Lee, K. H. & Musolff, L. (2021), 'Entry into two-sided markets shaped by platform-guided search', *Job Market Paper, Princeton University* .

Musolff, L. (2024), Algorithmic pricing, price wars and tacit collusion: Evidence from e-commerce, Technical report, Working Paper, Wharton School of the University of Pennsylvania.

Pai, M. M. & Vohra, R. (2014), 'Optimal auctions with financially constrained buyers', *Journal of Economic Theory* **150**, 383–425.