University Library

MINERVA ACCESS

A gateway to Melbourne's research publications

Minerva Access is the Institutional Repository of The University of Melbourne

Author/s:
Lim, B;Shemesh, J

Title:
Good Bids Come to Those Who Wait: The Value of Late Bidding in Online Auctions

Date:
2019-06-01

Citation:
Lim, B. & Shemesh, J. (2019). Good Bids Come to Those Who Wait: The Value of Late Bidding in Online Auctions. Economic Record, 95 (309), pp.248-261. https://doi.org/10.1111/1475-4932.12456.

Persistent Link:
https://hdl.handle.net/11343/286846

Author Manuscript

This is the author manuscript accepted for publication and has undergone full peer review but has not been through the copyediting, typesetting, pagination and proofreading process, which may lead to differences between this version and the Version of Record. Please cite this article as doi: 10.1111/1475-4932.12456

This article is protected by copyright. All rights reserved

Author Manuscript

Good Bids Come to Those Who Wait: The Value of Late Bidding in Online Auctions

Bryan Lim and Joshua Shemesh*

January 19, 2019

Abstract

This paper proposes an intuitive rationale for late bidding in online venues. The expected surplus from bidding on subsequent auctions for equivalent items creates an option value to losing the current auction. This option is dynamic due to the stochastic arrival of new auctions and early bids on later-closing auctions. We demonstrate that late bidding can be optimal given the decentralized and heterogeneous nature of online auctions, in which the option value is exogenous to an individual bidder's actions. Late bidding precludes the bidder from being locked into a suboptimal bid as her opportunity set evolves. (*JEL* C73, D44, D83, L81)

Keywords: late bidding; bid timing; online auctions; eBay; sniping

*We thank Flavio Menezes (the editor), several anonymous referees, Spencer Martin, Yongli Zhang, and seminar participants at the University of Melbourne for their helpful comments.
Lim (Corresponding Author)*:* Senior Lecturer, Department of Finance, University of Melbourne, Australia. E-mail blim@unimelb.edu.au
Shemesh: Senior Lecturer, Department of Banking and Finance, Monash University, Australia. E-mail joshua.shemesh@monash.edu

This article is protected by copyright. All rights reserved

I Introduction

A distinctive feature of online auctions is the preponderance of bidding in the last moments of an auction. At its extreme, with bids placed in the final seconds, the phenomenon has both a name, "sniping", and secondary websites devoted strictly to facilitating it (e.g., esnipe.com, auctionsniper.com, justsnipe.com). Late bidding has been formally documented by Wilcox (2000), Roth and Ockenfels (2002), Bajari and Hortaçsu (2003), Schindler (2003), and Gonzalez, Hasker and Sickles (2009), among others. Hopenhayn and Saeedi (2016) document that as recently as June 2014, "more than 40% of all bids, and more than 60% of winning bids, are submitted in the last 10% of the duration of an auction."

While late bids can be readily explained in the context of common-valuation (CV) auctions,[1] the motivation in private-valuation (PV) auctions, the type which constitutes the overwhelming volume of online auctions, is less clear. Online sites like eBay typically run second-price auctions, with the winning bidder paying the second-highest bid plus some nominal increment. Ely and Hossain (2009) observe that "the prevalence of sniping in second-price auctions in a private-value setting is surprising as auction theory suggests that sniping would be, at best, no more profitable than bidding early if rival bidders follow undominated strategies." To explain late bids in PV settings, some theorists have argued that they are optimal if rivals condition their bids on existing bids.[2] Roth and Ockenfels (2002), for example, posit that late bidding represents best response to naïve rivals who bid as though the auction were a standard (first-price) English auction, an argument that requires at least some bidders to be sufficiently irrational so as to employ dominated bidding tactics.

[1]For common valuation (CV) auctions, sniping is optimal to prevent free riding, as argued by Schindler (2003), Bajari and Hortaçsu (2003), and Hossain (2008). Hossain (2008), for example, conjectures that some bidders do not know how much they value an item but only know whether they value it more than the current price. Such bidders learn their valuation through the bidding process, which creates incentive for better informed agents to withhold their bids until the last second.

[2]Ockenfels and Roth (2006) theorize that last-second bidding represents tacit collusion by experienced bidders, who keep the final price low in the process. They note that very late bids have a positive probability of not being successfully submitted, and this opens a way for bidders to implicitly collude and avoid bidding wars. Barbara and Bracht (2006) argue that sniping is rational since early bidders can retract bids prior to the auction's close. Though not commented upon by the authors themselves, Engleberg and Williams' (2009) "bid-and-discover" algorithm for shill bidding on eBay can be circumvented by sniping.

This article is protected by copyright. All rights reserved

Author Manuscript

Empirical studies, however, have failed to identify significant differences in the final sale prices of auctions with and without late bids. Wintr (2008) finds no significant differences in sale prices across a range of commonly listed computer equipment in auctions with late bids versus those without. Ely and Hossain (2009) explicitly test whether late bids generate differential surplus to early bids. Participating in a series of auctions on eBay, they document an economically minute advantage in average surplus from late bidding and conclude that whatever respective advantages of either early or late bids roughly cancel each other out.[3]

The motivation for late bids offered by the literature no longer holds if a bidder cannot influence the other interested bidders in a given auction. What's more, in a world where other buyers are not conditioning their bids based on existing bids for the auction in question, there are reasons to bid early rather than late. Hossain (2008) suggests that bidders place multiple bids to learn about their own valuations. Bapna, Goes, and Gupta (2003) suggest that early bidders usually aim to evaluate or explore the auctions. Ariely and Simonson (2003) note that prices – i.e., the current highest bid – tend to be very attractive in the early stages of an auction, making the decision to enter the auction appear easy and risk-free. As of this writing, the eBay system explicitly encourages early bids with the prompt: "Bid now so you can get a deal!" Having identified a suitable auction on which to bid, a bidder may observe that since she cannot affect her rivals' bids, a reasonable strategy is simply to bid immediately.

Our paper revisits the bid-timing question but takes as given that bidders understand the structure of online auction markets along three key dimensions. First, they understand the second-price nature, such that they are not pursuing bidding wars. Second, they understand that new auctions on substitutable goods arrive stochastically. Finally, bidders in a given auction understand that they are unlikely to overlap in future auctions given the heterogeneous nature of online auctions. By "heterogeneous", we mean that even on a consolidated

[3]Ely and Hossain (2009) posit two possible effects of early bids. The first is the "competition effect", whereby an early bid deters entry into the auction by announcing one's presence, lowering the expectation of the winning bid. The second is the "escalation effect", whereby an early bid induces a bidding war from a naive bidder who behaves as though he is involved in a first-price auction.

This article is protected by copyright. All rights reserved

Author Manuscript

market like eBay, participation in auctions requires search across items which are nearly always differentiated in some form, like brand, condition, or proximity.

In this setting, we offer an intuitive rationale for late bidding in online PV auctions, one which doesn't require the existence of unsophisticated bidders. As Zeithammer (2006) observed, the online bidder's goal is not to win a particular auction but to win any auction from a set of auctions for substitutable goods. This implies an option value to losing an auction, a value which evolves dynamically as market conditions change. We posit two sources of exogenous change: the arrival of new auction listings and the arrival of early bids by other bidders on later-closing auctions. As new auctions open, the option value increases; as later-closing auctions get bid up, the option value decreases. We show that even when a bidder cannot affect others' bids through her own bids, late bidding may be optimal as it allows the bidder to maximize her information regarding the value of the option. If a bidder bids early, her preferred bid just prior to the auction's close may be below what she has already submitted. In this case, she is locked into a higher bid than she would like since bids can only be revised upward. If the cost to revisit the market is sufficiently low, she should wait as late as possible to collect information about her alternative buying opportunities. Late bidding precludes her from being locked into a suboptimal bid in the future.

As a simple example, consider someone looking to purchase a bike for up to \$100 within the next week. Presently, there is only one suitable bike auction to bid on, and it closes in three days. She could bid immediately or she could wait. Regardless of her bid, one of two possibilities can occur over the next three days: additional bike auctions ending within her deadline arrive or they do not. Her bid on the current auction will be considerably less than \$100 if new auctions arrive but will be exactly \$100 otherwise. If she bids some amount $x \leq \$100$ immediately and new auctions subsequently arrive, she may be locked into a sub-optimal bid, since bids cannot be revised downward.

Formally, we model an auction market as a sequence of open-bid, second-price auctions, similar though not exclusive to the structure on eBay. We show that in a market that can

This article is protected by copyright. All rights reserved

Author Manuscript

have multiple auctions on equivalent goods which are potentially overlapping in time, late bidding is weakly dominant when monitoring is costless: a bidder is always at least as well off submitting a bid late as opposed to early. When monitoring is costly, the decision to bid late depends on whether the gain in expected surplus exceeds the cost of monitoring. Rather than solving for the distribution of bids that results from optimal bidding behavior and finding a fixed point, we take the distribution of rival bids as given and find the bidder's best response, a modeling decision we believe best captures the bidder's problem. The model yields a closed-form solution for the option value of losing, in which the bidder's expected bid increases on each successive auction.

A bidder may face uncertainty about her ability to revisit the market in the future. In this case, she can contract out the bidding to a sniping service or software, which will submit a last-second bid on her behalf. In the earlier example, if the bidder is uncertain whether she will be able to revisit the market just prior to the auction's close, she can schedule a snipe bid of $100, an amount which she can revise if she is able to revisit the market and observe whether new auctions have arrived. Automated bidding may be extended to schedule last-second bids on a sequence of auctions for substitutable goods, for similar reasons as a scheduled bid on a single auction.

We contribute to the literature by demonstrating that late bidding does not require rivals to be conditioning their bids and can simply be the result of stochastic changes to the bidder's opportunity set. Crucially, our explanation does not require rivals to bid suboptimally (by engaging in bidding wars, for example). Despite its apparent intuitiveness, this point has not yet been made in the literature. Although our theoretical results match many empirical facts, our story is ultimately normative, not positive. Rather than try to explain observed late bidding, we illustrate why it can nevertheless be part of a surplus-maximizing strategy.

Our thesis, that late bidding is an optimal tactic with the stochastic arrival of new auctions and bids on later closing auctions, is subtly but critically differentiated from Roth and Ockenfels' (2002) observation that a non-strategic reason to bid late is the "desire to

This article is protected by copyright. All rights reserved

Author Manuscript

retain flexibility to bid on other auctions offering the same item". A desire to retain flexibility is not equivalent to surplus maximization over a stochastically evolving opportunity set. Our paper is related to a contemporaneous paper by Hopenhayn and Saeedi (2016), who similarly examine the optimal bidding when bidders' information sets are dynamic. A key distinction between the two papers is the respective auction markets. Their paper concerns optimal bidding on a single auction, while ours concerns optimal bidding on a stochastically evolving sequence of auctions.

A distinguishing feature of our paper relative to much of the extant literature is that we do not rely on what Zeithammer and Adams (2010) term the "sealed-bid abstraction": assuming away the possibility of early bidding, a modeling decision typically supported by citing the empirical evidence. Zeithammer (2006), Said (2011), Backus and Lewis (2012), and Arora, Xu, Padman and Vogt (2004) have all analyzed sequential options taking as given that all bidding occurs at the last second. Consequently, these papers either explicitly or implicitly model sequences of *sealed*-bid second-price auctions. By contrast, our paper models sequences of *open*-bid second-price auctions. With sealed bids, the bid-timing intuition of our paper would be trivial. With open bids, however, the optimality of a late bid is not obvious, particularly without assuming rivals' playing potentially suboptimal strategies. Accordingly, our paper provides a robust justification for not only late bidding but also the sealed-bid abstraction.

The paper is arranged as follows. The generalized model is given in Section II. A solution for the option value is given in Section III. Section IV examines several additions to the basic model, including uncertain monitoring, fixed-price alternatives, tolerance for substitutes, and state-dependent beliefs. Section V concludes.

Author Manuscript

This article is protected by copyright. All rights reserved

II The Model

The model is based on the auction structure on eBay, though it more generically applies when the market for a substitutable good is not centralized on any single venue. The description of the auction market generally follows Ockenfels and Roth (2006), with the sequential nature of bidder's problem echoing Zeithammer (2006).

There is an auction market where sellers periodically list a single unit of a good and where buyers arrive, bid, and depart. The market is ongoing, with new auctions opening and existing auctions closing. At any given time, there is a (possibly empty) sequence of second-price auctions, each for one unit of the good. Time is discrete.

The auction market is characterized by the following features:

- The arrival of new auctions is exogenous, implying that sellers are not strategically timing when they list their auctions.
- An auction arriving at time t will close at $a \equiv t + D$, where $D \in \mathbb{N}^+$. An auction is uniquely identified by its closing time a.
- The starting bid is 0.
- At the start of each period, the history of bids up to and including the current second-highest bid $b \equiv b(a,t)$ is publicly known for each auction. The current high bid is known only to the bidder who made it. If there is only one bid on the auction, the only public information for that auction is that a bid has been submitted.
- A bid greater than b can be submitted on a at any time prior to its closing time.
- When the current time equals a, the highest bidder wins the auction at the second-highest bid.[4]

Figure 1 illustrates the lifetime of an auction arriving at t.

[4]In practice, the highest bidder wins at the second-highest bid plus some nominal increment. For expositional simplicity, we set that increment to zero, but our results are robust to it being strictly positive.

This article is protected by copyright. All rights reserved

Author Manuscript

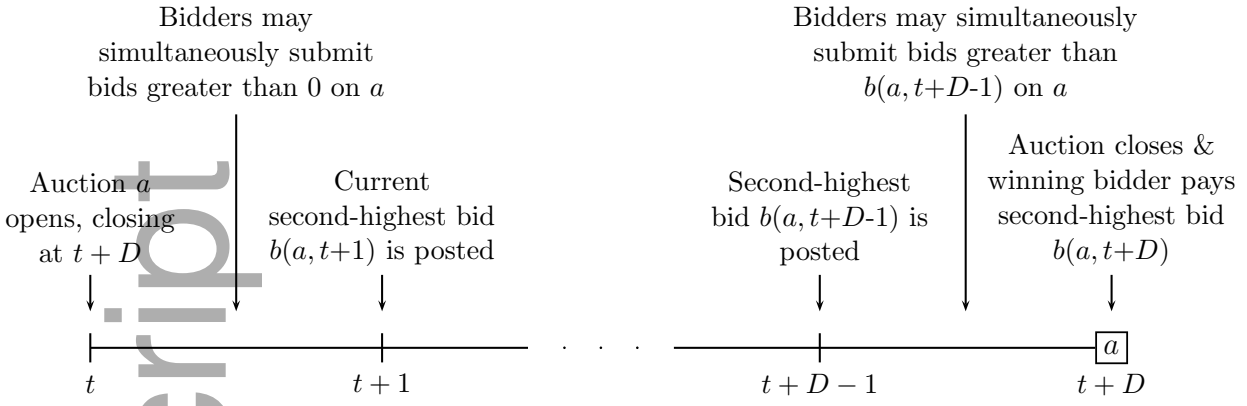


Figure 1: Lifetime of an Auction Arriving at t and Closing at $t + D$

(i) The Bidder's Problem

We consider the optimal strategy of a risk-neutral bidder who is currently in the market for a unit of a good from a set of acceptable substitutes. At the start of each period, a single auction for a unit of the good arrives with probability λ. The arrival probability is bidder-specific and measures the intersection of her willingness to accept substitutes and the exogenous arrival rates for auctions on those substitutes. The more specific her requirements are, the lower the probability is.[5] She has a private valuation $V \in [0, 1]$ for the good, regardless of the potential variation in the actual items she will bid on. (We discuss type-specific V in Section IV.)

The bidder is endowed with a fixed deadline T, which can be interpreted as the time at which she will exercise a zero-surplus option to purchase the item outside the auction market, say, at a retail outlet. She will exit the market upon having won a single auction or the time equaling T, whichever comes first. The deadline T measures her patience, with low realizations implying a need to obtain the item sooner rather than later. Note that the time when the bidder originally arrived to the market is irrelevant; only the interval until T

[5]If she is searching for a "bike", her λ should be considerably higher than if she is searching for a "Cannondale carbon fiber road bike" that is less than three years old and located within 20 miles of her.

Author Manuscript

This article is protected by copyright. All rights reserved

matters.[6]

The bidder's goal is to win a single auction at the lowest possible price below V at some time before or on T. Surveying the market at time t, she observes the set of auctions $\mathbf{a} \equiv \mathbf{a}(t) = [a_1\, a_2 \ldots a_N]$ which meet her search requirements, close before or on T, and have minimum bids $\mathbf{b} \equiv \mathbf{b}(t) = [b_1\, b_2 \ldots b_N]$ lower than V. The auctions are indexed chronologically, with a_1 closing the earliest and a_N the latest. Collectively, $\mathbf{a}$ and $\mathbf{b}$ represent the current market *state*.

Focusing on the current auction a_1, the optimal bidding strategy has two components: 1) when to bid, and 2) how much to bid. Putting aside the former for the moment, assume the bidder was going to submit a bid on a_1 immediately. Her expected surplus from a bid of x_1 is

$$\pi_{1,t}(x) \equiv \int_{b_1}^{x_1} (V - y_1) \frac{f(y_1)}{1 - F(b_1)} dy_1 + \frac{1 - F(x_1)}{1 - F(b_1)} \pi_{2,t} \tag{1}$$

where $f(y_1)$ is the distribution of the highest bid y_1 by all other bidders on a_1 and $\pi_{2,t}$ is the option value of bidding on auctions a_2 onward, calculated at t. The first component in (1) represents the expected surplus from winning the auction, and second component the expected surplus from bidding on all subsequent auctions for equivalent goods. Both $f(y_1)$ and $\pi_{2,t}$ represent subjective calculations from which she determines the optimal bid, the latter bounded between 0 and V. For dates $t' > t$, $\pi_{2,t'}$ is a random variable that can be either greater or less than $\pi_{2,t}$, conditional on whether new auctions arrive or later-closing auctions get bid up.

[6]The specification of a finite deadline represents a departure from most prior models. Both Zeithammer (2006) and Backus and Lewis (2012), for example, assume losing bidders exit the market randomly. In the context of the typical consumer goods for which PV applies, it is not obvious that one approach bears more fidelity to reality than the other. Under our specification, bidders face a fixed deadline, while under the alternative specification, bidders have an expectation of when they expect to leave, but may randomly leave earlier or (potentially) much later. Recent work by Coey, Larsen, and Platt (2016) employs a bidder deadline similar to ours.

Author Manuscript

This article is protected by copyright. All rights reserved

(ii) Bidder Beliefs

The motivation of this paper is to determine the optimal strategy when the bidder cannot affect the distribution of the highest rival bid:

$$df(y_n)/dx_1 \quad = \quad 0 \, \forall n \tag{2}$$

The underlying beliefs in Equation (2) are motivated by the following observations[7]:

- Online auctions are second price in format. Vickrey (1961) established the dominant strategy in a single, sealed-bid, second-price auction is to bid one's private valuation. The intuition holds for *open*-bid second-price auctions as well, so rivals who are bidding only on a_1 optimally should not condition on others' bids.

- Online auctions rarely traffic in truly homogeneous goods. A sizable proportion of auctions are for used goods, and even new goods are often differentiated by brand, quality, or proximity. The heterogeneity in supply virtually necessitates heterogeneity in demand, as bidders must identify suitable auctions through searches. The intersection of both suitable (substitutable) goods and their searchable descriptions will vary across bidders who are ostensibly looking for the same item.[8] Given two bidders in a particular auction, one may be more flexible than the other in terms of her willingness either to tolerate substitutes or to conduct her search across multiple venues, like retail or classifieds. Accordingly, the likelihood that both bidders will overlap in any future auctions for their respective desired good is attenuated, and any information gleaned about a rival by a bidder is likely to be of limited predictive value.

[7] We do not make the case that the beliefs in Equation (2) are optimal. Optimality requires, among other things, ad hoc assumptions about the behavior (e.g., higher-order beliefs) of rivals, assumptions which may not apply to actual bidders in online auctions.

[8] This heterogeneity in both supply and demand is fundamentally different from what Englebrecht-Wiggans (1994) terms "stochastically equivalent" goods. The latter refers to a single set of goods for a fixed group of bidders, while the former refers to different bidders bidding on different sets of goods.

Author Manuscript

This article is protected by copyright. All rights reserved

In addition to these structural reasons for Equation (2), the empirical evidence documents that winning bids for auctions on identical items with sufficiently frequent listings tend to exhibit little volatility, and the timing of a bid does not generate econonomically differential surplus (Wintr, 2008; Ely and Hossain, 2009). A bidder may observe these tendencies and conclude that she cannot game her rivals into submitting lower bids or that whatever marginal surplus she might hope to extract is not worth the effort required to do so.

Equation (2) implies that a given bid generates the same expected surplus regardless of whether it was submitted early or late. Relaxing this condition requires one of two alternatives to replace it: relative to a late bid, an early bid either increases or decreases her expected surplus. But neither alternative is supported by the evidence. Nekipelov (2007) and Ely and Hossain (2009) speculate that an early bid announces the bidder's presence in an auction and deters potential rivals from bidding. While Ely and Hossain indeed empirically observe that an early bid tends to draw fewer rival bidders, they simultaneously find the average surplus from this strategy is no greater than from submitting the same bid late. Wang (2006) argues that an early bid announces the bidder's presence to her rivals, which in turn lowers their respective option values of bidding in subsequent auctions due to the expected competition from the bidder if she loses the current auction. Accordingly, an early bid would raise rivals' bids on the current auction. As the intuition requires the assumption that the pool of bidders remains fixed from auction to auction, we should observe more late bidding in auctions for more specialized items. Wintr's (2008) results, however, run counter to this. The more generic goods (e.g., laptops and monitors) exhibit considerably more skewness toward late bidding than the more specialized goods (e.g., stamps and coins). Moreover, to the extent that Wang's intuition is correct, we should observe higher average surplus to late bids, but again, this is not supported by the existing evidence.

Author Manuscript

This article is protected by copyright. All rights reserved

(iii) Optimal Bid Strategy

If the bidder is committed to submitting a bid immediately at t, the optimal bid is given by the first-order condition from (1):

$$x^*_{1,t} = V - \pi_{2,t}$$

$x^*_{1,t}$ represents a once-and-for-all bid on a_1, after which she will not revise her bid should additional auctions arrive. Of course, the bidder is under no obligation to make a once-and-for-all bid. There are then two broad possibilities to consider. The first is that a_1 closes immediately: $t = a_1 - 1$. In this case, she cannot postpone the decision of whether or not to bid on a_1. If she doesn't bid, she expects a surplus of $\pi_{2,t}$ from bidding on the remaining auctions. If she does bid, she knows bidding $x^*_{1,t}$ yields a surplus at least as large as $\pi_{2,t}$. Hence, she should submit $x^*_{1,t}$ if $V - \pi_{2,t} > b_1$ and not bid otherwise. This is a *late bid*.

The second and more interesting case is that a_1 does not close immediately: $t < a_1 - 1$. With each passing unit of time, the option value may change. It may increase (decrease) because new auctions open (do not open). It may decrease because later-closing auctions get bid up. The interval of time between t and $a_1 - 1$ reveals information to the bidder and, in doing so, updates her valuation of losing a_1.

She will incur a monitoring cost c if she revisits the market at $\bar{t}$ and observe the updated state variables. Even on a centralized platform like eBay, finding suitable auctions requires searching potentially across multiple sets of keywords, identifying potential candidates from the limited information (e.g., title and image) provided by the search results, and finally verifying suitable auctions after reading individual item descriptions. Depending on the nature of the bidder's requirements – including but not limited to her flexibility with respect to substitutes, geographic distance, and price – the time required to learn the updated state variables can vary. The cost c is her opportunity cost of this time.[9]

[9]The willingness to accept substitutes (λ) and the monitoring cost (c) are not necessarily correlated. The monitoring cost is the product of the cost to verify an individual result times the total number of results. If a bidder is more willing to accept substitutes, the per-item verification cost may be low, but the number items to verify will be high, so ex-ante, their product is unclear. The verification cost is itself a function

This article is protected by copyright. All rights reserved

Author Manuscript

We define $\bar{t} \equiv a_1 - 1$ as the last period the bidder can submit a bid on the current auction. For now, we limit her timing decision to bidding either immediately or at $\bar{t}$ or both. The value of the option at $\bar{t}$ is denoted by $\pi_{2,\bar{t}}$, and at t, it is unknown and exogenously determined. The bidder estimates its distribution based on λ, $\mathbf{a}$, $\mathbf{b}$, T, and $f(y_n)$, among other things. Whatever the realization of $\pi_{2,\bar{t}}$ ultimately is, the corresponding optimal bid will be $x^*_{1,\bar{t}} = V - \pi_{2,\bar{t}}$. Since she cannot affect $\pi_{2,\bar{t}}$, her bidding decision is based on her time-t expectation of the expected surplus from bidding optimally at $\bar{t}$, conditional on the distribution of $\pi_{2,\bar{t}}$.

Let $\pi_{1,\tau}(x)$ equal the expected surplus from bidding x at a generic time τ – i.e., the value of Equation (1) given a bid x made at τ. The bidder's decision involves weighing the cost to revisit the market against the increase in expected surplus from bidding at $\bar{t}$: c versus $E_t[\pi_{1,\bar{t}}(x^*_{1,\bar{t}})] - \pi_{1,t}(x^*_{1,t})$.

Proposition 1. *Assuming $V - \pi_{2,t} > b_1$, the optimal bid is $x^*_{1,t} = V - \pi_{2.t}$. When $t = \bar{t}$ or when $t < \bar{t}$ and $c > E_t[\pi_{1,\bar{t}}(x^*_{1,\bar{t}})] - \pi_{1,t}(x^*_{1,t})$, the weakly dominant strategy is to bid $x^*_{1,t}$. Otherwise, the weakly dominant strategy is to revisit the market at $\bar{t}$.*

Proof. When $t = \bar{t}$, bidding $x^*_{1,t}$ is trivially weakly dominant, since the expected surplus from doing so is at least as large as from not submitting a bid. When $t < \bar{t}$, there are three pairwise comparisons to consider:

1. Single bid at $\bar{t}$ versus a single bid at t. At t, the expected surplus from bidding $x^*_{1,\bar{t}} = V - \pi_{2,\bar{t}}$ at $\bar{t}$ is $E_t[\pi_{1,\bar{t}}(x^*_{1,\bar{t}})]$, while the expected surplus from bidding $x^*_{1,t}$ immediately is $\pi_{1,t}(x^*_{1,t})$. For any non-degenerate distribution of $\pi_{2,\bar{t}}$, the former is at least as large as the latter – they are equal only when $x^*_{1,t} = x^*_{1,\bar{t}}$ – but incurs the cost c to implement. Hence, only bid at t if the increase in expected surplus is greater than c.

2. Single bid at $\bar{t}$ versus bids at both t and $\bar{t}$. Let $x_{1,t}$ denote a generic bid made at t.

of both a willingness to accept substitutes and the opportunity cost of the bidder's time. In this way, an unemployed lottery winner searching for a specific item can have a lower monitoring cost than a hedge fund manager searching for a generic item. We take as given that from the bidder's point of view, both λ and c are exogenously determined.

This article is protected by copyright. All rights reserved

Author Manuscript

The comparison in this case is $E_t[\pi_{1,\bar{t}}(x^*_{1,\bar{t}})]$ versus $E_t[\pi_{1,\bar{t}}(\max(x_{1,t}, x^*_{1,\bar{t}}))]$. If $x_{1,t} > x^*_{1,\bar{t}}$, the expected surplus from bidding $x_{1,t}$ will be less than from bidding $x^*_{1,\bar{t}}$. In this case, a late bid is always at least as good as an earlier bid, since both strategies incur c.

3. Single bid at t versus bids at both t and $\bar{t}$. Since a bid at $\bar{t}$ weakly dominates bidding at both t and $\bar{t}$, there will not be a case when bidding at both t and $\bar{t}$ dominates a single bid at either t or $\bar{t}$. □

An implication of Proposition 1 is that whatever strategy at t that weakly dominates bidding at $\bar{t}$ also weakly dominates bidding at a future time t' prior to $\bar{t}$. If it is too costly to revisit the market at $\bar{t}$, it is similary too costly to revisit at t'. If is optimal to revisit the market at $\bar{t}$, then it is suboptimal to revisit earlier at t'. Bidding at either t' or $\bar{t}$ incurs the same monitoring cost, but the expected surplus from the former is smaller than from the latter, for the same reason that bidding immediately yields lower expected surplus than bidding at $\bar{t}$.

Abstracting from monitoring costs, one would bid early if doing so would increase either the expected surplus from winning a_1 or the expected surplus from bidding on all subsequent auctions. Neither is the case. An early bid does not affect the distribution of the highest rival bid, so the expected surplus from winning the current auction is not affected. Similarly, an early bid does not affect the realization of the option value at $\bar{t}$, so the expected surplus from all subsequent auctions is also not affected.

Given strictly positive monitoring costs, if monitoring costs are sufficiently low, she should wait until the last moment to bid on a_1. Conversely, when the cost is sufficiently high relative to the gain in expected surplus from waiting, the bidder submits a bid immediately. In this case, waiting is not worth her time. This tradeoff between waiting and monitoring becomes especially acute if the current time is close to the closing time.[10]

[10]In a practical sense, the interval between the last moment to bid ($\bar{t}$) and the closing time will be bidder-specific, as it corresponds to the time required to observe the updated state variables. For some bidders, this may be seconds, while for others, it may be minutes or even hours, depending on the characteristics of their search.

This article is protected by copyright. All rights reserved
Author Manuscript

Proposition 1 demonstrates that even if a bidder cannot affect her expected surplus on a_1 with the timing of her bid, there exists a reason to submit a late bid. As time elapses, new auctions may arrive and later-closing auctions may get bid up. If she cannot favorably affect the distributions of these events, the option value is orthogonal to her bidding on the current auction, so she is always at least as well off submitting a late bid. Waiting to observe if new auctions arrive provides her with the most complete information set about the option value of losing the current auction. If she submitted a bid $x_{1,t}$ prior to $\bar{t}$, should this bid turn out to be more than $x^*_{1,\bar{t}}$, she will be locked into a suboptimal bid. Letting $\mathbf{\Pi}_{2,\bar{t}}$ represent the support of $\pi_{2,\bar{t}}$, the only undominated bids placed at $t < \bar{t}_1$ are those less than $\min_{\pi_{2,\bar{t}} \in \mathbf{\Pi}_{2,\bar{t}}} V - \pi_{2,\bar{t}}$.[11]

One can frame Proposition 1 in terms of the example in the introduction. A bidder identifies a set of bike auctions on which to bid, with the possibility of that set increasing in size as time passes and new auctions open. She must decide whether to bid on the current auction and how much to bid. If the auction does not close immediately and her monitoring cost is sufficiently low, there is no benefit to bidding now. Between now and the current auction's closing time, new auctions may arrive. If there are later-closing auctions open, new bids from rivals may arrive on these. Her estimate of the value of losing the current auction will be updated in the future, prior to the auction's close. Whatever bid she might place initially is unlikely to be the bid she would place in the future, given the additional information she will have. Moreover, if she cannot affect the bidding strategy of her marginal competitor for that auction, she cannot affect the expected surplus from winning the auction with a bid of x_1 by bidding early.

Remark. When $t < \bar{t}$, the gain in expected surplus from revisiting the market at $\bar{t}$ is increasing in the volatility of $\pi_{2,\bar{t}}$.

Since $\pi_{1,\bar{t}}(x_1)$ is concave in x_1 and maximized at $x_1 = V - \pi_{2,\bar{t}}$, the greater the variation in $\pi_{2,\bar{t}}$, the greater the probability-weighted average deviation between the state-dependent

[11] Though she could bid $\min_{\pi_{2,\bar{t}} \in \mathbf{\Pi}_{2,\bar{t}}} V - \pi_{2,\bar{t}}$ early, this strategy potentially involves repeated and ultimately superfluous bidding that would not generate any additional surplus relative to late bidding.

This article is protected by copyright. All rights reserved

Author Manuscript

Figure 2: Optimal Bid Timing for $D = 4$**.** Panels (a), (b), and (c) illustrate optimal bid timing for three different scenarios, all with $D = 4$. The solid black triangle in Panel (a) indicates the bidder should optimally revisit the market at that time, while the outlined white triangles in all three panels indicate the bidder should revisit the market at the indicated time only if she does not win a_1.

optimal bid and a fixed early bid $x_{1,t}$. It follows that the loss in expected surplus from an early bid is increasing in the volatility of $\pi_{2,\bar{t}}$. The more varied the potential option values can be, the more varied the corresponding optimal late bids will be. On average, then, a fixed early bid $x_{1,t}$ will be further from optimal when the option value is more variable. For goods with more option-value variability – e.g., due to lower arrival rates – the incentive to bid late increases, given a fixed cost c.

When the criteria for a late bid are met, the bidder needs to revisit the market only at $a_1 - 1$ to potentially bid. Assuming she is notified directly if she has won (i.e., she does not have to visit the market to verify whether she has won), should she not win a_1 she need only revisit at $a_2 - 1$ if a_2 existed at $t = a_1 - 1$ or at $a_1 + D - 1$ otherwise. Moreover, the bidder should submit an early bid if there is no possibility of an additional eligible auction opening ($t \geq T - D$).

Three examples with $D = 4$ are illustrated in Figure 2. In Panel (a), the current auction is closing in two periods, so the bidder should revisit the market in one period at $t + 1$ to potentially bid. If she does not win, she should revisit the market at $t + 3$. In Panel (b), since the current auction closes next period, the bidder should bid on a_1 immediately. If she does not win, she should revisit the market at $t + 2$. In Panel (c), again, the bidder should bid immediately. If she does not win, she should revisit the market at $t + 4$, since if $D = 4$ the earliest a subsequent auction could close is $t + 5$.

This article is protected by copyright. All rights reserved

Author Manuscript

III Solving for the Option Value

We have to this point taken the option value $\pi_{2,t}$ as given, but fundamentally it represents a transformation of state variables into a bidder-specific expected surplus. In this section, we elaborate on the basic mechanics of that calculation.[12] The dual specifications of a fixed deadline T and the independence of bid timing and rivals' bids allow us to solve for $\pi_{2,t}$ as a relatively straightforward dynamic programming problem.[13]

Dropping the subscript t and denoting the proposed solution as π_2^*, the calculation requires that the bidder has a strategy of how she will bid on future auctions and forecasts of how she expects her rivals to bid on those auctions. A forward-looking strategy can be deduced iteratively. Insofar as Proposition 1 yields the weakly dominant strategy for a_1, one can look forward and infer that it will similarly yield the weakly dominant strategy for a_2, a_3, and so forth. For analytical clarity, the cost c to revisit the market is set to 0 in the analysis, such that late bids are optimal for the current and all future auctions. When c is strictly positive, the analysis remains largely intact except the bidder must additionally integrate over the set of states in which future bids on later-closing auctions should optimally be placed early. For example, in the event that bidding at $\bar{a}_1$ is weakly dominant, the bidder must account for future states at $\bar{a}_1$ in which the weakly dominant strategy is not to bid on a_1 and in which the increase from waiting to make a late bid on a_2 exceeds the cost of revisiting.

Having established a set of beliefs regarding future auctions, one can now go about calculating π_2^*. If she loses a_1, on how many more auctions can she bid? When $t \geq T - D$, there is no uncertainty about this number. There are N auctions currently open, so she can potentially bid on $N - 1$ of them subsequent to a_1. In this case, the calculation of π_2^* is

[12] As a practical exercise, calculating $\pi_{2,t}$ is complex, involving but not limited to integrating over potential future states driven by estimated auction arrival probabilities and distributions of winning bids on future auctions. This calculation is almost surely beyond the scope of the representative bidder in an online PV auction, given the typical value of the good being sold relative to the cost – psychic or opportunity – involved in that calculation. Our solution is intended as the basis for whatever heuristic a bidder would employ.

[13] An examination of bidding without these specifications can be found in Backus and Lewis (2012).

This article is protected by copyright. All rights reserved

Author Manuscript

straightforward. Let π_n^* represent the option value of losing the $n-1^{th}$ auction. If there are N auctions, it must be the case that $\pi_{N+1}^* = 0$. It follows that $x_N^* = V - \pi_{N+1}^* = V$.[14] If $V - \pi_{n+1} > b_n$,

$$\begin{aligned} x_n^* &= V - \pi_{n+1}^* \\ \pi_n^* &= \int_{b_n}^{x_n^*} \frac{f(y_n)}{1-F(b_n)}(V-y_n)dy_n + \frac{1-F(x_n^*)}{1-F(b_n)}\pi_{n+1}^* \end{aligned}$$

Otherwise $\pi_n^* = \pi_{n+1}^*$, and no bid will be made on a_n.

When $t < T - D$, the number of auctions on which she can potentially bid is random. Each unit of time between t and $T - D$ yields the possibility of a suitable auction arriving. The arrival or non-arrival on an auction at each unit of time in the future corresponds to a unique timeline of auctions. Abstracting from potential changes in the vector $\mathbf{b}$ of rival bids on those auctions, each timeline corresponds to a future state s. The arrival process governed by λ means the evolution of future states can be modeled as a binomial tree with 2^{T-D-t} terminal nodes, where each terminal state represents a unique sequence of auctions.

The binomial tree is path-dependent. An auction arrival followed by a non-arrival may generate a different strategy than a non-arrival followed by an arrival. Suppose, for example, that over a 3 period interval, a single auction arrives. If the bidder's deadline is in 7 periods, her strategy may be different if the auction arrived in the first period versus if it arrived in the third. Figure 3 illustrates the tree over four periods.

The 2^{T-D-t} states $s \in \mathbf{S}$ at $T-D-t$ each have a corresponding timeline of auctions. One can calculate $\pi_{2|s}^*$ as the option value conditional on state s occurring. π_2^* is the probability-weighted average over $\mathbf{S}$:

$$\pi_2^* = \sum_{s}^{s \in \mathbf{S}} pr(s)\pi_{2|s}^*$$

where $pr(s)$ is the probability of state s, as governed by λ.

[14] The reader may have anticipated this result. Bidding on a_N is conditional on losing all previous auctions and having no subsequent auctions on which to bid. With only one auction to bid on, the game collapses into the canonical game in Vickrey (1961). A similar result obtains if the arrival rate λ equals 0, in which case the option value is 0 as well.

This article is protected by copyright. All rights reserved

Author Manuscript

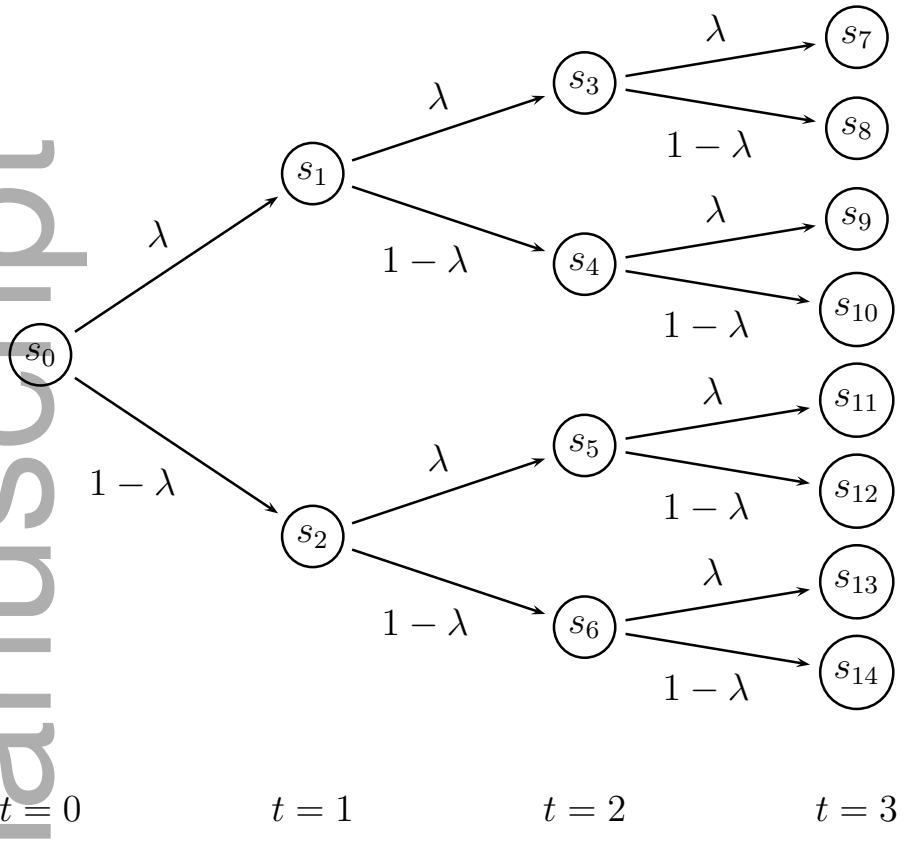


Figure 3: Binomial Tree when $t < T-D$. s_0 represents the current state, with the subsequent nodes corresponding to potential future states. Up nodes indicate an auction opening at t, while down nodes indicate an auction not opening at t. The state s_5, for example, is the state s_0 updated with an auction opening at $t = 2$ (and closing at $2 + D$)

Example 1. The bidder's deadline is in 5 periods, and newly-listed auctions close in 3 periods ($D = 3$). There are currently two eligible auctions, with the first (a_1) closing in one period and the second (a_2) closing in three periods. Next period, there are two possibilities: an auction opens or it does not. The former is denoted by s_1 and occurs with probability λ, while the latter is denoted by s_2 and occurs with probability $1 - \lambda$. Any auction which arrives after next period is not relevant since it closes after the bidder's deadline. To keep the intuition as simple as possible, we set all minimum bids b_n equal to 0, and $f(y_n)$ equal to 1 for all n (i.e. $y_n \sim U(0, 1)\ \forall n$). The evolution of possible timelines are illustrated in Figure 4.

Author Manuscript

This article is protected by copyright. All rights reserved

Author Manuscript

Figure 4: Timelines for Example 1, with $D = 3$ and $T = t + 5$

In this case, if the bidder were to lose a_1, she knows she can bid on a_2. Furthermore, she knows that she could wait to submit a bid on a_2 until $t + 2$. By waiting, she observes whether auctions arrive at $t + 1$ and $t + 2$, and she can condition her eventual bid on that information. One can express her option value of losing a_1 as

$$\begin{aligned} \pi_2^* &= \lambda \pi_{2|s_1}^* + (1-\lambda)\pi_{2|s_2}^* \\ &= \lambda^2 \pi_{2|s_3}^* + \lambda(1-\lambda)\pi_{2|s_4}^* + (1-\lambda)\lambda \pi_{2|s_5}^* + (1-\lambda)^2 \pi_{2|s_6}^* \end{aligned} \tag{3}$$

This article is protected by copyright. All rights reserved

The state-dependent optimal bids and surpluses are

$$
\begin{aligned}
x^*_{4|s_3} &= x^*_{3|s_4} = x^*_{3|s_5} = x^*_{2|s_6} = V \\
\pi^*_{4|s_3} &= \pi^*_{3|s_4} = \pi^*_{3|s_5} = \pi^*_{2|s_6} = V^2/2 \\
x^*_{3|s_3} &= x^*_{2|s_4} = x^*_{2|s_5} = V - V^2/2 \\
\pi^*_{3|s_3} &= \pi^*_{2|s_4} = \pi^*_{2|s_5} = V^2 - V^3/2 + V^4/8 \\
x^*_{2|s_3} &= V - V^2 + V^3/2 - V^4/8 \\
\pi^*_{2|s_3} &= (3/2)V^2 - (3/2)V^3 + (9/8)V^4 - (11/16)V^5 + (5/16)V^6 - (3/32)V^7 + (1/64)V^8
\end{aligned}
$$

One can solve for π^*_2 by substituting into (3). Since $t = a_1 - 1$, she should bid $x^*_1 = V - \pi^*_2$ immediately.

(i) Comparative Statics

Proposition 2 outlines the comparative statics for π^*_2. The results are generally consistent with similar analyses in Zeithammer (2006) and Said (2011) despite considerably different modeling assumptions.

Proposition 2. *π^*_2 is characterized by the following:*

1. *From t to $t+1$, π^*_2 increases when an eligible auction opens and decreases otherwise*
2. *π^*_2 is decreasing in b_n for all $n > 1$*
3. *π^*_2 is increasing in $T - t$, and x^*_1 is decreasing in $T - t$*
4. *π^*_2 is increasing with respect to V*
5. *If $t < T - D$, π^*_2 is increasing with respect to λ*

Proof. Proofs are in the Appendix. □

Author Manuscript

This article is protected by copyright. All rights reserved

There is no conceptual inconsistency between Proposition 2.2 and the bidder's beliefs specified in Equation (2). The reason is chronological. Proposition 2.2 implies that bids on later auctions can affect bids on earlier auctions by way of the option value. In contrast, (2) states that a bid on an earlier auction cannot affect bids on later auctions.[15]

As noted earlier, the specification of a fixed deadline T is a departure from earlier models, which typically model exit as a Poisson process. The fixed deadline means that in expectation, the bidder should increase her bids over time, a result that matches both Englebrecht-Wiggans' (1994) discussion of the declining option value of later-closing auctions and Coey, Larsen, and Platt's (2016) empirical finding that losing bidders tend to increase their bids on subsequent auctions.[16]

The option value increases with the bidder's private valuation V, since a higher V increases both the probability of winning the current auction (by increasing the optimal bid x_1^*) and the expected surplus conditional on winning. Similarly, higher λ implies more future auctions on which to bid, which in turn raises the option value of losing earlier-closing auctions.

IV Alternative Specifications

For pedagogical purposes, we have attempted to keep the model as simple as possible. As these specifications may not necessarily apply to real-life auction markets, we now consider the effects of altering various features of the model. Rather than fully developing each alternative specification, our goal is to demonstrate that the reasoning of Proposition 1 persists in various extended frameworks.

[15]By symmetry, Proposition 2.2 implies that the bidder can potentially reduce her rivals' respective option values by bidding early on later auctions, but doing so would then push up their optimal bids on earlier auctions. Putting aside the operating assumptions that any two bidders are unlikely to overlap in their future auction sets and that each bidder only bids on one auction at a time, the only implication from Proposition 2.2 regarding Equation (2) is that the latter can potentially be amended to $df(y_n)/dx_m \geq 0 \, \forall n < m$: i.e., the bidder can only *adversely* affect the distribution of rival bids on earlier auctions through bidding on a later auction. Exclusion of such sub-optimal bidding yields the original specification of Equation (2).

[16]With stochastic exit rather than a fixed deadline, the bidder does not expect her future bids to be increasing.

This article is protected by copyright. All rights reserved

Author Manuscript

(i) Uncertain Monitoring

In the model, the bidder can revisit the market for a cost c. By incurring this cost, she accesses the market with probability 1. In practice, there may be exogenous constraints on her time, making her future access to the auction market uncertain.[17] We refer to this as *uncertain monitoring*. To accommodate uncertain monitoring, one can assign probability $\gamma(\tau)$ that she will have access to the market at time $\tau < a_1$ in the future. With uncertain monitoring, a value can be attached to the ability to submit a non-binding bid.

Suppose the bidder does not know today if she will be available to revisit the market prior to a_1 to determine if a new auction(s) has arrived. She would be willing to bid some relatively high amount $x^*_{1|H}$ assuming an auction doesn't arrive. As she is uncertain whether she will revisit the market, she may contract out the bidding to a service (or software) at a cost k, which will submit a late bid on her behalf. An early bid submitted directly on the auction effectively handcuffs her from recalibrating her strategy if a new auction opens. If both $\gamma(\tau)$ and k are sufficiently low, she can bid $x^*_{1|H}$ through the service, knowing that the scheduled bid is non-binding and can be amended if she is able to revisit the market prior to the auction closing and incorporate new information into an updated bid. Uncertain monitoring consequently provides a rationale for *last-second* (snipe) bids, as opposed to late bids more generally.

This reasoning extends beyond the first auction to $\mathbf{a}$, the set of all currently open auctions which fit her requirements. With uncertain monitoring, the bidder can schedule a series of bids which will get executed in sequence unless she is able to revisit the market and update her bids, conditioning on the increased information set. Indeed, the vast majority of sniping services and software provide for this functionality.

[17] In Ambrus, Ishii, and Burns (2016), for example, bidders have "periodic, random opportunities to check the auction's status".

Author Manuscript

This article is protected by copyright. All rights reserved

(ii) Fixed-Price Alternatives

Depending on the item sought, the market can potentially be very large, distributed across multiple online venues as well as non-auction markets. In particular, nothing prevents the bidder from conducting a dual search for fixed-price listings, either on an auction site(s) or via classified advertisements or even retail channels. Each of these channels – auction, classified, and retail – has a stochastic component to it, whether that be in arrival time and/or price. Some bidders may monitor multiple channels, all of which would get factored into $\pi_{2,t}$. The intuition of Proposition 1 is reinforced by the additional channels, as an early bid may prevent the bidder from purchasing the item more immediately from an alternate channel.

(iii) Tolerance for Substitutes

For simplicity, V has been normalized to [0, 1], though this specification is at odds with the possibility that the bidder is willing to accept one of many substitutes. While her willingness to accept substitutes informs the arrival probability λ of a suitable auction opening, to explicitly model differentiated substitutes we can specify type-specific V and λ, similar to Zeithammer (2006). In this case, a bidder has private valuation V_k for each of the $k \in K$ types she is willing to accept, and $\lambda \cdot \lambda_k$ is the probability that an auction for a k-type unit, arrives where $\sum_{k \in K} \lambda_k = 1$.

Differentiated values of V_k allow for the possibility that there are types of auctions for which, if one of those types arrives, the bidder has higher expected surplus than for other types. Sufficiently large variation in type-specific expected surplus reinforces the motivation to withhold bidding on the current auction until late.

Author Manuscript

This article is protected by copyright. All rights reserved

(iv) State-Dependent Beliefs

Even if subjective, $f_n(\cdot)$ may nevertheless be state-dependent, contrary to the specification in (2). Consider two states s' and s'' with the same state time ($\tau(s') = \tau(s'')$) such that either there are more auctions in state s' than in s'' or the minimum bids on all auctions are at least as low in s' as in s''. The bidder may speculate that $F(y|s') \geq F(y|s'')\,\forall y$, given the decreased competition in state s'.[18]

Since the state is exogenous, the solution for π_2^* remains intact. If one were to replace each $f(y)$ with $f_n(y_n|s)$ in the preceding section, the necessary refinement to the solution would be a redefinition of terminal states as those corresponding to a state-time of $T-1$ instead of $T-D$. With state-dependent beliefs about the distribution of y_n, the expected surplus is evolving with the state, even if the set of auctions on which to bid is fixed. The game tree increases by a factor of 2^{D-1}, but the mechanics of the solution remain intact.

V Conclusion

This paper examines a simple rationale for late bidding in online auctions: maximizing the information about the stochastically evolving option value of bidding on subsequent auctions for equivalent goods. We demonstrate that the optimality of late bidding does not require rivals to be conditioning their bids on existing bids. Taking both the arrival of new auctions and the bidding of other participants as exogenous, we demonstrate that late bidding is optimal for a bidder with sufficiently low monitoring costs. Losing an auction has an option value, as the bidder can potentially win a subsequent auction. She should postpone submitting a bid until as late as possible, as this will allow her to better estimate the option value of bidding on later-closing auctions. The contracting of sniping services is optimal if the bidder is uncertain whether or not she can revisit the market (just) prior to the current auction's close.

[18] In Example 1, for example, it is reasonable for the bidder to believe that $F_2(\cdot|s_3) > F_2(\cdot|s_6)$ given that at s_3 there are more auctions on which to bid than at s_6.

This article is protected by copyright. All rights reserved

Author Manuscript

A solution for the option value is provided. According to this solution, the bidder's preferred bids on each subsequent auction are increasing, with a maximum of her private valuation on the final auction. The more patient the bidder is, the less she prefers to bid on any auction. When her deadline extends to infinity, her optimal strategy is to consistently submit the lowest possible bid, as this will ensure the maximum surplus if and when she wins.

The contribution of this paper is providing a rationale for late bidding when buyer cannot affect her rivals' bids through her own bidding. We are motivated empirically by Wintr (2008) and Ely and Hossain (2009), both of whom document no economic difference in surplus from early bids versus late bids. Our theoretical motivation comes from the decentralized nature of online shopping – across not only auction markets but also potentially retail and classified markets – and the heterogeneity of online bidders, the combination of which reduces the likelihood that any pair of auctions has overlapping bidders. The random arrival of new auctions and early bids on later-closing auctions implies that late bidding is optimal for a bidder who can potentially bid on future auctions. The optimality is expressed with respect to an objective not to win a particular auction but rather to maximize one's expected surplus conditional on winning any auction prior to her deadline. While the narrative is couched in terms of online auctions, the underlying intuition applies more generally. One could, for example, drop the specification that auctions are second-price and use the key results of this paper for a housing search model, wherein a bidder may postpone placing an offer on a particular house for as long as possible while searching for substitutes on which to potentially bid. Indeed, the intuition may apply to any model of auctions involving search.

Author Manuscript

This article is protected by copyright. All rights reserved

References

[Ambrus, Ishii BurnsAmbrus .2013] ambrus2013gradualAmbrus, A., Ishii, Y. Burns, J. 2013. Gradual bidding in ebay-like auctions. Gradual bidding in ebay-like auctions.

[Arora, Xu, Padman VogtArora .2004] arora2004optimalArora, A., Xu, H., Padman, R. Vogt, W. 2004. Optimal bidding in sequential online auctions. Optimal bidding in sequential online auctions.

[Backus LewisBackus Lewis2012] backus2012demandBackus, M. Lewis, G. 2012. A demand system for a dynamic auction market with directed search. A demand system for a dynamic auction market with directed search.

[Bajari HortacsuBajari Hortacsu2003] bajari2003winnerBajari, P. Hortacsu, A. 2003. The winner's curse, reserve prices, and endogenous entry: Empirical insights from eBay auctions The winner's curse, reserve prices, and endogenous entry: Empirical insights from ebay auctions. RAND Journal of Economics329–355.

[Bapna, Goes GuptaBapna .2003] bapna2003replicatingBapna, R., Goes, P. Gupta, A. 2003. Replicating online Yankee auctions to analyze auctioneers' and bidders' strategies Replicating online yankee auctions to analyze auctioneers' and bidders' strategies. Information Systems Research143244–268.

[Coey, Larsen PlattCoey .2016] coey2016theoryCoey, D., Larsen, B. Platt, B. 2016. A theory of bidding dynamics and deadlines in online retail. A theory of bidding dynamics and deadlines in online retail.

[Ely HossainEly Hossain2009] ely2009snipingEly, JC. Hossain, T. 2009. Sniping and squatting in auction markets Sniping and squatting in auction markets. American Economic Journal: Microeconomics1268–94.

This article is protected by copyright. All rights reserved

Author Manuscript

[Engelbrecht-Wiggans .Engelbrecht-Wiggans .1994] engelbrecht1994sequentialEngelbrecht-Wiggans, R. . 1994. Sequential auctions of stochastically equivalent objects Sequential auctions of stochastically equivalent objects. Economics Letters441-287–90.

[Gonzalez, Hasker SicklesGonzalez .2009] gonzalez2009analysisGonzalez, R., Hasker, K. Sickles, RC. 2009. An analysis of strategic behavior in eBay auctions An analysis of strategic behavior in ebay auctions. The Singapore Economic Review5403441–472.

[Hopenhayn SaeediHopenhayn Saeedi2015] hopenhayn2015dynamicHopenhayn, H. Saeedi, M. 2015. Dynamic bidding in second price auction. Dynamic bidding in second price auction.

[HossainHossain2008] hossain2008learningHossain, T. 2008. Learning by bidding Learning by bidding. The RAND Journal of Economics392509–529.

[NekipelovNekipelov2007] nekipelov2007entryNekipelov, D. 2007. Entry deterrence and learning prevention on eBay. Entry deterrence and learning prevention on ebay.

[Ockenfels RothOckenfels Roth2006] ockenfels2006lateOckenfels, A. Roth, AE. 2006. Late and multiple bidding in second price Internet auctions: Theory and evidence concerning different rules for ending an auction Late and multiple bidding in second price internet auctions: Theory and evidence concerning different rules for ending an auction. Games and Economic Behavior552297–320.

[Roth OckenfelsRoth Ockenfels2002] roth2002lastRoth, AE. Ockenfels, A. 2002. Last-minute bidding and the rules for ending second-price auctions: Evidence from eBay and Amazon auctions on the Internet Last-minute bidding and the rules for ending second-price auctions: Evidence from ebay and amazon auctions on the internet. American Economic Review9241093–1103.

Author Manuscript

This article is protected by copyright. All rights reserved

[SaidSaid2011] said2011sequentialSaid, M. 2011. Sequential auctions with randomly arriving buyers Sequential auctions with randomly arriving buyers. Games and Economic Behavior731236–243.

[SchindlerSchindler2003] schindler2003lateSchindler, J. 2003. Late bidding on the Internet. Late bidding on the internet.

[VickreyVickrey1961] vickrey1961counterspeculationVickrey, W. 1961. Counterspeculation, auctions, and competitive sealed tenders Counterspeculation, auctions, and competitive sealed tenders. The Journal of Finance1618–37.

[WangWang2006] wang2006lastWang, JTY. 2006. Is last minute bidding bad. Is last minute bidding bad.

[WilcoxWilcox2000] wilcox2000expertsWilcox, RT. 2000. Experts and amateurs: The role of experience in Internet auctions Experts and amateurs: The role of experience in internet auctions. Marketing Letters114363–374.

[WintrWintr2008] wintr2008someWintr, L. 2008. Some evidence on late bidding in eBay auctions Some evidence on late bidding in ebay auctions. Economic Inquiry463369–379.

[ZeithammerZeithammer2006] zeithammer2006forwardZeithammer, R. 2006. Forward-looking bidding in online auctions Forward-looking bidding in online auctions. Journal of Marketing Research433462–476.

[Zeithammer AdamsZeithammer Adams2010] zeithammer2010sealedZeithammer, R. Adams, C. 2010. The sealed-bid abstraction in online auctions The sealed-bid abstraction in online auctions. Marketing Science296964–987.

Author Manuscript

This article is protected by copyright. All rights reserved