Submitted to *Management Science*
manuscript MS-00425-2005.R1

Temporary and Permanent Buyout Prices in Online Auctions

Jérémie Gallien
MIT Sloan School of Management, Cambridge, MA 02142, jgallien@mit.edu,

Shobhit Gupta
MIT Operations Research Center, Cambridge MA 02139, shobhit@mit.edu,

Buyout options allow bidders to instantly purchase at a specified price an item listed for sale through an online auction. A *temporary* buyout option disappears once a regular bid is submitted, while a *permanent* option remains available until it is exercised or the auction ends; such buyout price may be *static* and remain constant throughout the auction, or *dynamic* and vary as the auction progresses. We formulate a game-theoretic model featuring time-sensitive bidders with independent private values and Poisson arrivals but endogenous bidding times to answer the following questions: How should a seller set the buyout price (if at all)? What are the implications of using a temporary buyout option relative to a permanent one? What is the potential benefit associated with using a dynamic buyout price? For all buyout option types we exhibit a Nash equilibrium in bidder strategies, argue that this equilibrium constitutes a plausible outcome prediction, and study the problem of maximizing the corresponding seller revenue. Our numerical experiments suggest that when any of the participants are time-sensitive, the seller may significantly increase his utility by introducing a buyout option, but that dynamic buyout prices may not provide a substantial advantage over static ones. Furthermore, while permanent buyout options yield higher predicted revenue than temporary options, they also provide additional incentives for late bidding and may therefore not be always more desirable.

Key words: buyout option, online auction, Nash equilibrium, dynamic pricing

1. Introduction

As they were initially conceived during the last decade of the previous century, online auctions were arguably suffering from two perceived drawbacks relative to posted price mechanisms: waiting time and price uncertainty. Many auction sites have since introduced a new feature known as a *buyout option*, which offers potential buyers the opportunity to instantaneously purchase at a specified price an item put for sale through an online auction. Augmented with this option, an online auction becomes a hybrid between a fixed-price catalogue and a traditional auction.

Buyout options are now widespread and have significant economic importance: in the fourth quarter of 2003 alone, fixed income trading (primarily from the buyout option "Buy It Now") contributed $2 billion or 28% of *eBay*'s gross annual merchandise sale[1]; other examples of buyout options include *Yahoo*'s "Buy Price", *Amazon*'s "Take-It" and *uBid*'s "uBuy it!". Remarkably, buyout options in these large auction sites currently differ in one important aspect: *eBay*'s "Buy

[1] Source: *http://investor.ebay.com/*, see also Reynolds and Wooders (2003).

It Now" option disappears as soon as a regular bid above the reserve price is submitted, so it is called *temporary*; in contrast *Yahoo*, *Amazon* and *uBid*'s options remain until they are exercised or the auction in which they are featured ends, so they are called *permanent* (Hidvégi et al. 2003). However, all auction sites just mentioned (and for that matter all auction sites we are aware of) use *static* buyout options, meaning that the buyout price is fixed at the outset and may not be modified during the auction.

These observations motivate in our view the following questions:

1. How should a seller using an online auction set the buyout price (if at all)?
2. What are the implications of using a temporary buyout option relative to a permanent one?
3. What is the potential benefit associated with using a *dynamic* buyout price that may vary as the auction progresses?

This paper contains the description and analysis of a game-theoretic model designed to answer these questions in a stylized setting. It is organized as follows: after a discussion of our contribution relative to the existing literature in §2, we present and discuss our model in §3. Section §4 contains an equilibrium analysis for the temporary buyout option (in §4.1), the permanent buyout option (in §4.2), and a study of the associated seller's optimization problem (in §4.3). A comparative discussion of the insights obtained for both types of buyout options, relying on numerical experiments and our theoretical results of the previous subsections, is then provided in §4.4. We next discuss dynamic buyout prices in §5, where §5.1 focuses on outcome prediction, §5.2 on the resulting optimization problem, and §5.3 on numerical experiments. Section §6 contains our concluding remarks, and all proofs can be found in a technical supplement available online at *http://web.mit.edu/jgallien/www/*.

2. Literature Review and Paper Contributions

While the literature on auction theory is large, existing research work on buyout prices is recent and relatively limited. Indeed, the comprehensive 1999 survey of the auction literature by Klemperer (1999) makes no mention of buyout prices, and while Lucking-Reiley (2000) observes the use of buyout prices in his 2000 survey of internet auction practices, he points out that he is "[...] not aware of any theoretical literature which examines the effect of such a buyout price in an auction." The theoretical papers written since on buyout prices are listed in Table 1, which shows some of their model features and which one(s) of the three motivational questions listed in the previous section they address.

A first set of papers relies on simple models to shed light on why the addition of a buyout option may increase the seller's revenue under various circumstances. Kirkegaard and Overgaard (2003) show that when two bidders with multi-unit demand face two sequential auctions of one

	number of bidders			auction participants		bidding times			buyout option		motivational question(s) addressed		
	2	*arbitrary*	*unknown*	*risk-averse*	*time-sensitive*	*not captured*	*exogenous*	*endogenous*	*temporary*	*permanent*	*1*	*2*	*3*
Kirkegaard and Overgaard (2003)	•					•			•		•		
Budish and Takeyama (2001)	•			•		•				•	•		
Hidvégi et al. (2003)		•		•		•				•	•		
Reynolds and Wooders (2003)		•		•		•			•	•		•	
Caldentey and Vulcano (2006)			•		•		•			•	•		
Mathews (2003a, 2003b, 2004)		•		•	•			•	•		•		
this paper			•		•			•	•	•	•	•	•

Table 1 Existing Theoretical Models of Buyout Options

item each, the seller may benefit from using a buyout option in the second auction because of the information revelation occurring in the first. Budish and Takeyama (2001) show that a seller facing two risk averse bidders may improve expected profit by using an optimal permanent buyout price, and Hidvégi et al. (2003) show that this last observation still holds in an extended model with an arbitrary number of bidders and continuous valuation distributions.

Among the papers focusing on the impact of risk aversion, Reynolds and Wooders (2003) deserves special mention, as it appears to be the only theoretical paper besides ours comparing temporary and permanent buyout options in the same unified framework, which seems an indispensable methodological requirement when addressing the second motivational question listed in §1. However, their analysis ignores the issues of time sensitivity and bid arrival times altogether. As a first consequence, it fails to identify the different implications of temporary and permanent buyout options on bidding concentration near the end of the auction (see also discussion of Caldentey and Vulcano 2004 below). Secondly, in the model of Reynolds and Wooders (2003) all participating bidders (who are not differentiated by their arrival time) may exercise both types of options at the outset, with success determined by tie probabilities. The first bidder's strategic option in practice to remove a temporary buyout option is thus not captured, and neither is the near certainty that a bidder will get the item when exercising either type of buyout option. This casts some doubts on the validity or at least generality of their analysis, since the key difference between temporary and permanent options is the set of bidders to whom the option is available. Finally, we point out that Reynolds and Wooders (2003) do not address question one and, like all other papers cited here, question three from §1.

More generally, the discussion forums of experienced online auction users suggest that time sensitivity of both sellers and buyers is an important reason for using a buyout option in practice (Gupta 2006). The importance of time sensitivity on buyout exercise is also confirmed by the survey of auction users reported in Wan et al. (2003), and eBay itself, in its user guidelines, states the reduction in waiting time as the very first reason why both sellers ("Sell your items fast.")

and buyers ("Buy items instantly!") would want to use their buy it now feature (eBay 2006). Recognizing this, a second and more recent subset of papers that includes ours investigates models of buyout options that capture explicitly the time-sensitivity of participants and the dynamic arrivals of bidders.

In particular, Caldentey and Vulcano (2004) use a model similar to ours in several respects (Poisson arrival of bidders, utility functions with exponential time-discounting) in order to study equilibrium behavior and optimal price for a "dual auction and list price channel" resembling an auction with a permanent buyout option. In Caldentey and Vulcano (2004) however, when making decisions during the bidding period buyers are assumed to ignore past bid values and current auction asking price, and even whether the item listed is still available or not. In our model (as is currently the case on all major auction sites), bidders are provided with that information, and their decisions may change accordingly. The information structure assumed in Caldentey and Vulcano (2004) is thus less realistic, but more importantly it could not support the comparative study between permanent and temporary buyout options that we undertake here. Indeed, the key differences between bidder strategies in the temporary and permanent cases stems from bidders' different strategic options and different assessment of the competition they face when confronted with given past bidding activity (see §4.4), the very information assumed away in that paper. As a result, the information model of Caldentey and Vulcano (2004) is too coarse to predict any difference in bidder behavior between temporary and permanent buyout option cases. Secondly, bidders in Caldentey and Vulcano (2004) are assumed to bid or exercise the buyout option instantly upon their arrival, whereas in practice bidding activity concentrates near the end of the auction (Roth and Ockenfels 2002). Because bidding times are endogenous and an integral part of the bidders' strategy in our model, we are able to show that using a permanent option likely increases last-minute bidding relative to a temporary one. The issue of last-minute bidding seems important to major auctioneers (Amazon.com offers a 10% discount to the first bidder in order to induce early bidding), so that these model differences are material. We point out however that Caldentey and Vulcano (2004) investigate the multi-unit case, whereas we only consider the case of a single item.

Finally, Mathews (2003b), Mathews (2004) and Mathews (2003a) investigate variations of a model of a temporary buyout option featuring time-sensitive participants, dynamic bidder arrivals and endogenous bidding times. This model also assumes uniform valuations, and that the total number of bidders is fixed at the outset and known to all. As a result, these papers fail to capture that the availability of the temporary buyout option conveys some important information to an incoming bidder about the amount of competition he is likely to face in the auction, which obviously affects the attractiveness of exercising the buyout option relative to placing a normal bid: in Mathews' model, the first bidder does not modify his assessment of the likely number of competing bidders, regardless of whether he arrives at the very beginning or at the very end of the bidding period. This results in bidder equilibrium strategies that are qualitatively different from those we derive: when bidders' time sensitivity becomes very small in Mathews (2003a), the valuation threshold triggering the buyout option exercise becomes independent of the bidder's arrival time (this is easily seen in Mathews (2003b) and Mathews (2004), which consider time-insensitive bidders). In contrast, the buyout threshold function in our model does depend on (specifically, increases with) the first bidder's arrival time, even in the limit where bidders are time-insensitive. Furthermore, Mathews' model of bidder arrivals is not specific enough to support a comparative study between temporary and permanent buyout options: a bidder finding the buyout option still available likely infers that he is the first to arrive in the temporary case, but only that previously arrived bidders did not decide to exercise the buyout option in the permanent case. As a result, a bidder arriving to an auction with a given buyout price and no past bidding activity is more likely to exercise a permanent option than a temporary one. Because bidders' belief about their competitors' arrival process is not specified in Mathews (2003a, 2003b, 2004), this important issue (see our Figure 1 and following comments) could not be captured without a more complete information structure of the type we propose. Finally, although Matthews (2004) exhibits as we do an equilibrium whereby the first bidder bids or exercises the temporary buyout option immediately, his paper does not provide any argument why this specific equilibrium constitutes a better outcome prediction than any of the many other ones. It should be noted however that Mathews (2003b) explores the issue of bidder welfare, which we do not address here.

We highlight in closing that this paper is to the best of our knowledge the first to:

- present an optimization study generating qualitative insights on whether sellers should use a posted-price, a pure auction or an auction with a buyout price when confronted with various time-sensitivities and bidder arrival rates;
- compare temporary and permanent buyout options using a unified modeling framework capturing the impact of participants' time-sensitivity;
- analyze dynamic buyout prices. Note that our model feature of endogenous bidding times is also key for that purpose, as for example bidders knowing of a future decrease of the buyout price could not realistically be assumed to bid or buyout upon their arrival regardless.

3. Model

In this section, we first describe our game-theoretic model, focusing on the market environment in §3.1 and the auction mechanism in §3.2. We then discuss its realism in §3.3.

3.1. Market Environment. We consider a monopolistic seller opening at time 0 a market for one item. From that point on, he faces an arrival stream of potential buyers (or bidders) which is non-observable per se, but is correctly believed by all participants to follow a Poisson process with a known, exogenous and constant rate λ. Bidders valuations (or the prices at which they are indifferent between purchasing the item and not participating in the market) are assumed to follow an independent private values model – see Klemperer (1999) for background. Specifically, each bidder has a privately known valuation, and all other participants initially share the correct belief that this valuation has been drawn independently from a distribution with cdf F (which is assumed to be Lipschitz continuous) and compact support $[\underline{v}, \bar{v}]$ (define $m = \bar{v} - \underline{v}$).

All participants are risk-neutral and time-sensitive. In particular, the utility of the seller when earning revenue R at time τ is assumed to be $U^S(R, \tau) \triangleq e^{-\alpha\tau} R$, where $\alpha > 0$ denotes his time discounting factor. Likewise, a bidder arriving at time $t > 0$ with valuation $v \in [\underline{v}, \bar{v}]$ who purchases the item at time $\tau \geq t$ for a payment of x gets utility $U(v, t, \tau) \triangleq e^{-\beta(\tau - t)}(v - x)$, where $\beta > 0$ denotes his time discounting factor, assumed to be the same for all bidders. A losing bidder is assumed to derive zero utility from the market.

3.2. Auction Mechanism. The basic market mechanism we consider is a second-price auction with a time-limited bidding period $[0, T]$. That is, any bidder arriving at time $t \in [0, T]$ may submit a bid at any time in $[t, T]$, provided it is larger than any other he may already have submitted (i.e. bidders are not allowed to renege on their purchasing offers). At time T, the item is sold to the highest bidder who pays then a price equal to the second highest bid; if only one bidder has submitted a bid by T the item is sold to him for a price of $\underline{v}$, and if there are no bids the item is not sold. Note that the lower bound of the distribution support $\underline{v}$ thus effectively corresponds to a publicly advertised minimum required bid (any bids lower than $\underline{v}$ are ignored).

In addition to all the other information described previously, every bidder is assumed to know at every time τ subsequent to his arrival the value of I_τ, defined as the payment that would be made by the winning bidder if the auction were instead terminated at τ. That is, I_τ is equal to (i) the second highest bid submitted over $[0, \tau]$ if there are at least two such bids; (ii) $\underline{v}$ if there is only one; and (iii) 0 if there is none. As is the case on all auction websites we are aware of, we assume that I_t is truthfully revealed to any arriving bidder.

The basic auction mechanism just defined is investigated for example in Gallien (2006). The critical extension that we study in the present paper is the addition by the seller at the outset of a buyout price p, either temporary or permanent. Any bidder may exercise that buyout option at any time between his arrival and the end of the auction T, provided the option is still open then; this amounts to purchasing the item instantaneously at a price of p, effectively terminating the auction. A temporary buyout option remains open from the beginning until its exercise or the first time that a regular bid is submitted by any bidder, while a permanent buyout option remains open until its exercise or the end of the auction. That is, submissions of regular bids do not prevent bidders from subsequently exercising a permanent buyout option, but they do terminate a temporary buyout option. In line with observed practice, we assume that all participants know at any point in time whether the buyout option is still open.

While we assume in §4 that the buyout price p remains constant throughout the auction, we study dynamic buyout prices in §5. In the dynamic extension we consider then, the seller commits upfront to a function of time $[p(t)]_{t\in[0,T]}$ describing the evolution of the buyout price (either temporary or permanent) over time, and that function is known to all bidders.

3.3. Model Discussion. We first comment on our allocation mechanism. Online auction sites now typically feature "proxy bidding" systems, allowing bidders to enter the maximum amount they are willing to pay for the item. The system then submits bids on behalf of the bidder, increasing his outstanding bid whenever necessary and by as little as possible to maintain his position as the highest bidder, up until the maximum amount stated is reached[2]. As observed by Lucking-Reiley (2000), an online auction with a proxy bidding system effectively amounts to a second-price auction, the payment mechanism we assume.

For the closing rule, we assume a hard bidding expiration deadline similar to the one used on *eBay*, whereas some other sites such as *Amazon* use instead a floating deadline that automatically extends (within some limits) whenever a new bid close to the current deadline is submitted. As pointed out in Roth and Ockenfels (2002), this difference is material and *eBay*-like hard bidding deadlines account for a demonstrably higher concentration of bids near the end of the auction. In principle, our model allows to predict such surge of bids shortly before the end, because while we assume exogenous bidder arrival times, their bidding times are endogenous. In fact, our analysis in §4 confirms the intuition that last-minute bids seem more likely with a permanent buyout option than with a temporary one. However, our model does not capture some of the other important reasons why last-minute bidding does occur: presence of inexperienced bidders engaging in irrational

[2] See http://pages.ebay.com/help/buy/proxy-bidding.html for a detailed description.

bidding wars; informational value of bids when the item being sold has a common value component; reluctance to bring an auction to the attention of competing bidders performing searches on individual users's bidding activity; possibility that late bids may not reach the auction site due to network transmission delays... while we refer the reader to Roth and Ockenfels (2002) for an excellent discussion and empirical study of this phenomenon, we argue that factors such as the loss of last minute bids due to network transmission capacity and the presence of inexperienced bidders may not remain as prevalent in the long run, partly justifying these modeling choices (otherwise primarily motivated by tractability considerations). Consequently, in the model we assume for an online auction without a buyout option (or after a temporary buyout option has been removed), any sequence of bids culminating in the submission of one's true valuation before the bidding deadline T forms a weakly dominant strategy.

Another feature of the market mechanism we consider is the possible presence of a publicly announced minimum required bid, effectively captured in our model by the lower bound $\underline{v}$ of the valuation distribution support. Note that this is distinct from what some auction sites (such as *eBay*) call a "reserve price", which is likewise set by the seller as a minimum selling price for the item but, in contrast with the minimum required bid we use, is not publicly announced – when used by the seller, bidders are typically only informed that a reserve price has been set for the auction, and whether or not it has already been met by any of the existing bids. We assume that the seller does not use such concealed reserve price, in part because this would entail some inference of its value by the bidders, and may lead to further strategic interactions in the form of post auction negotiations between the winning bidder and the seller.

Several limitations of our analysis also stem from the market environment we consider. Our assumption that bidder arrivals follow a Poisson process seems more realistic than assuming that the number of bidders is known to all with certainty (as in nearly all other papers discussed in §2), and is partly justified by the classical Palm limit theorem on the superposition of counting processes. Nevertheless, the assumption that its arrival rate is constant and known to all participants (common to all other auction models assuming Poisson bidder arrivals that we are aware of) is still a strong one. In practice, the arrival rate of potential bidders to an auction could not only be variable but also endogenous, and depend for example on the bidding activity it has generated to date. In practice, the arrival rate of bidders to a specific auction is also influenced by factors such as advertising, the presence of a reserve price, the seller's feedback ratings, the presence and quality of photographs describing the item, etc. Our assumption of a constant known arrival rate saliently implies that bidders, including those arriving early in the auction when only little bidding history is available, correctly synthesize the impact of these factors when estimating how many competing bidders they are likely to face.

The structure assumed here for the utility functions of the seller and the bidders (time-discounted quasi-linear incentives) is also used for example in Caldentey and Vulcano (2004) and Gallien (2006), and reflects the proposed time sensitivity of participants. It saliently implies that bidders in our model do not have any fixed bidding or waiting costs. This is an important limitation, since bidders arriving to an auction may in practice decide to balk if their estimated transactional utility does not make up for these fixed costs. Our model, which only reflects that this utility is discounted by the transaction time, may thus significantly underestimate this balking behavior and its impact. Also, while all the results in the paper have been derived for auctions with risk neutral participants, some of them generalize to the case of risk averse bidders – see §6 for a detailed discussion.

In summary, while our model does capture some of the key features of an online auction, there are some others that it does not reproduce as faithfully. We point out that an actual online auction is a complex and random process involving multiple heterogeneous participants with various incentives and rationality levels interacting in a dynamic manner. As such, any tractable analytical model designed to predict its outcome (including ours and every other one described in the literature) must necessarily rely on fairly restrictive assumptions. Given one of our main research objective is to understand the differential impact of temporary and permanent buyout prices, we observe that several of these assumptions (e.g. common beliefs, bidder arrival process) may not specifically impact our model predictions when one type of buyout option is used as opposed to the other. From that perspective, we find it reassuring that our results rationalize some of the actual practices of auction sites using buyout options (see §6), and that some of our model predictions can be verified through a statistical analysis of real auction data (see §A.7 in the online supplement).

4. Static Buyout Prices

This section includes an equilibrium analysis for our model of an auction with a static buyout price, both temporary (in §4.1) and permanent (in §4.2), followed by a study of the seller's optimization problem (in §4.3). Numerical experiments are then described in §4.4.

4.1. Equilibrium Analysis of the Temporary Buyout Option. We now assume that the seller uses a fixed temporary buyout price p which disappears if a bid above the reserve price is placed in the auction. For any bidder arriving at time t with valuation v, consider the following family $\mathcal{T}[\cdot]$ of threshold strategies:

$$\mathcal{T}[\nu](v,t):\begin{cases}\text{Buyout at } p \text{ immediately} & \text{if buyout option available and } v>\nu(t,I_t)\\ \text{Bid } v \text{ immediately} & \text{if buyout option available and } v\le\nu(t,I_t)\ ,\\ \text{Bid } v \text{ at any time in } [t,T] & \text{otherwise}\end{cases}\qquad(1)$$

where $\nu : [0,T] \times [0,\bar{v}] \to [\underline{v},\bar{v}]$ is a threshold valuation function depending a priori on both the arrival time t and the second highest bid I_t defined in §3.2. Note however that a temporary option is only available when no bid has yet been placed or $I_t = 0$, so that ignoring the dependence of ν on I_t as we will do in the following entails no loss of generality for the temporary option case. Observe also that the only components of an increasing sequence of bids submitted by a participant that have strategic implications in our model of a temporary option are the time at which the first bid is submitted (because this may remove the buyout option), and the value of the highest bid submitted (because it may affect the auction outcome). Consequently, we will not distinguish between two bidding strategies that are equivalent modulo those two components. For example, even though the second possible action stated in the definition of $\mathcal{T}[\nu]$ in (1) is "Bid v immediately", a strategy whereby a bidder in the same case would place any bid v' in $[\underline{v}, v]$ immediately, then submit any sequence of bids with highest value v before the end of the auction would result in the exact same payoff and action space for himself, the other bidders and the seller. All the results to be stated about $\mathcal{T}[\nu]$ will thus also hold for any strategy or profile of strategies equivalent to it in the sense just defined. The concept of equilibrium uniqueness (see statement of Theorem 2) is also to be understood in this context, meaning that there does not exist any other equilibrium which is not equivalent to the one exhibited. Likewise, by a symmetric strategy profile we mean a set of strategies played by all players that are all equivalent. Finally, we will use the same notation for a strategy and the symmetric strategy profile obtained when every bidder plays that strategy, since no ambiguity arises from the present context.

The following theorem establishes the existence of a threshold function ν_{tmp} such that $\mathcal{T}[\nu_{tmp}]$ forms a Bayesian Nash equilibrium, and also provides a characterization of that function.

THEOREM 1. *Define function ν_{tmp} as $\nu_{tmp}(t) = \min\left(\hat{v}(t), \bar{v}\right)$ where $\hat{v}(t)$ is the unique solution on $[\underline{v}, +\infty)$ of the equation*

$$\hat{v}(t) - p = e^{-(\lambda+\beta)(T-t)} \int_{\underline{v}}^{\hat{v}(t)} e^{\lambda(T-t)F(x)} dx. \tag{2}$$

Then the symmetric strategy profile $\mathcal{T}[\nu_{tmp}]$ is a Bayesian Nash equilibrium for the online auction game with a temporary buyout price p.

The proof of Theorem 1 consists of deriving, for an arbitrary threshold function ν, a best response strategy to profile $\mathcal{T}[\nu]$, that is a strategy maximizing the utility of a bidder entering an auction where every other bidder uses strategy $\mathcal{T}[\nu]$. Specifically, denoting $\mathcal{R}(\mathcal{T}[\nu])$ the set of these best response strategies, we characterize a threshold function ν_{tmp} such that $\mathcal{T}[\nu_{tmp}] \in \mathcal{R}(\mathcal{T}[\nu])$. We further show that $\mathcal{T}[\nu_{tmp}] \in \mathcal{R}(\mathcal{T}[\nu_{tmp}])$, establishing that the profile $\mathcal{T}[\nu_{tmp}]$ constitutes indeed a Nash equilibrium; as in all our other proofs most of the difficulty stems from the bi-dimensional action space allowing for bidders to wait. Before discussing the intuition behind Theorem 1 and some qualitative implications, we state a proposition providing a closed-form expression for ν_{tmp} in the special (but widely assumed) case of uniformly distributed valuations:

PROPOSITION 1. *When bidder valuations are uniformly distributed on $[\underline{v}, \bar{v}]$, the threshold function ν_{tmp} characterizing the Bayesian Nash equilibrium described in Theorem 1 is*

$$\nu_{tmp}(t) = \min\left(p - \frac{m}{\lambda(T-t)}\left(W\left(-e^{-e^{-(\lambda+\beta)(T-t)} + \frac{(p-\underline{v})\lambda(T-t)}{m} - (\lambda+\beta)(T-t)}\right) + e^{-(\lambda+\beta)(T-t)}\right), \bar{v}\right), \quad (3)$$

where W is Lambert's W or omega function, i.e. the inverse of $W \mapsto We^W$.

In the equilibrium characterized by Theorem 1, the first incoming bidder compares upon his arrival the relative attractiveness of the buyout option and that of a regular bid, accounting for the likely competition resulting from the specific auction time remaining then; the dynamic threshold ν_{tmp} valuation characterized in (2) corresponds to the valuation of a bidder who at that time would be indifferent between the two options. Accordingly the threshold function ν_{tmp} is non-decreasing over time (this is easily established formally by inspection of (2)): the continuous buyout option availability over time indicates a reduced likely level of competition for the auction if it should take place, and therefore progressively makes the buyout option less attractive relative to submitting a regular bid.

Note also that strategy $\mathcal{T}[\nu_{tmp}]$ and the associated equilibrium result just stated do not provide a prediction of when the second and subsequent bidders will submit their bid. That is, the timing of bid submissions for these bidders does not have any strategic implication within the strict boundaries of our model definition. In practice however, it could be affected in various ways by features not captured by our model; for example a high cost of monitoring the auction could hasten bid submissions, while common value signaling could delay them – see §3.3 for a more complete discussion and related references.

An important observation is that the equilibrium $\mathcal{T}[\nu]$ specified in Theorem 1 is not unique. Indeed, for any $w > 0$ one may choose a threshold function $\nu : [0, T] \to [\underline{v}, \bar{v}]$ such that the strategy $\mathcal{T}^{(w)}[\nu]$ generalizing $\mathcal{T}[\nu]$ and defined by replacing the action associated with the second case in (1) with "Bid v after $\min(w, T-t)$ time units" also constitutes an equilibrium. That is, in the equilibria $\mathcal{T}^{(w)}[\nu]$ with $w > 0$, a bidder finding the buyout option still available when he arrives may wait for some time before submitting a bid. This is because, provided the threshold function is non-decreasing, this first bidder would lose in the auction anyway to any second bidder exercising the buyout option while the first bidder is still waiting. In the remainder of this subsection, we argue that, in contrast to $\mathcal{T}[\nu]$, such equilibrium does not survive some perturbations of our hypotheses, and therefore does not provide a robust outcome prediction.

Let G denote the online auction game with a temporary buyout option described in §3.2 and §3.1. We use the classical methodology of payoff perturbations (see van Damme (1987)) in order to refine our equilibrium analysis, and define $G^{(\epsilon)}$ as a game identical to G except that with a small probability $\epsilon > 0$ an arriving bidder is *desperate*, meaning that his utility from the auction with a type (v,t) is described instead by

$$U_D(v,t) = \begin{cases} +M & \text{if he obtains the item at } t; \\ -M & \text{if he bids in the auction;} \\ 0 & \text{otherwise,} \end{cases} \qquad \text{where } M \gg 0. \tag{4}$$

In words, desperate bidders greatly value the item auctioned, have an outside alternative with negligible value, and cannot wait under any circumstances; the dominant strategy for a desperate bidder is obviously to exercise the buyout option if it is available and to not participate at all otherwise. This specific perturbation seems appealing, because it may reveal the limiting impact of irrational bidders or bidders with different time sensitivities that our model otherwise assumes away (see §3.3). We prove the following result:

Theorem 2. *For any $\epsilon > 0$, the game $G^{(\epsilon)}$ does not have any Bayesian Nash equilibrium where a non-desperate bidder, who arrives when the buyout option is present, waits before bidding (e.g. plays $\mathcal{T}^{(w)}[\cdot]$ with $w > 0$). In addition, there exists a threshold function $\nu^{(\epsilon)}_{tmp} : [0,T] \to [\underline{v}, \bar{v}]$ such that for non-desperate bidders the strategy profile $\mathcal{T}[\nu^{(\epsilon)}_{tmp}]$ is a unique Bayesian Nash equilibrium of the game $G^{(\epsilon)}$, and $\lim_{\epsilon \to 0} \nu^{(\epsilon)}_{tmp} = \nu_{tmp}$ where ν_{tmp} is defined in Theorem 1.*

The intuitive explanation for the first statement in Theorem 2 is that when the first bidder decides to bid in the auction he is strictly better off bidding immediately and remove the buyout option then, because this prevents any subsequent desperate bidders from participating. It is clear however that the utility function of desperate bidders has been precisely defined to achieve that effect, and may thus appear ad-hoc or arbitrary. But from the perspective of refining our outcome prediction, what is striking about Theorem 2 is not that the introduction of desperate bidders per se preserves the equilibrium $\mathcal{T}[\nu]$ we propose and eliminates all other equilibria. The meaningful part is that this selection of equilibria occurs regardless of how small the introduction probability ϵ of these desperate bidders is. Indeed, an equilibrium which would not survive an arbitrarily small perturbation of the model payoff structure (whatever that perturbation) could hardly be considered robust. Theorem 2 actually establishes that the equilibrium $\mathcal{T}[\nu_{tmp}]$ characterized in Theorem 1 is the only one to survive the specific perturbation defined above, however small its probability.

Another standard robustness test (or equilibrium refinement technique) for outcome prediction is to use the solution concept of trembling-hand perfect equilibrium instead of the less discriminating Bayesian Nash equilibrium (Fudenberg and Tirole 1991). While we omit that analysis here due to length restrictions, it can also be shown that $\mathcal{T}[\nu_{tmp}]$ is the unique trembling-hand perfect equilibrium of G (Gupta 2006).

Finally, we describe in Section A.7 of our online supplement a simple empirical study that we have conducted in order to validate our model predictions. That is, we collected bidding data from a number of actual auctions of similar items (iPod music players and accessories) on the site eBay (which features a temporary buyout option), and focused on testing an implication of our analysis on the bidding times of auction participants. Specifically, the equilibrium analysis in our model does not generate any prediction for the first bidding time (or any other bidding time for that matter) in an online auction without a buyout option, while Theorems 1 and 2 do imply that the first bidder will act (bid or buyout) immediately upon his arrival. Consequently, in auctions featuring a buyout option, the first activity (bid or buyout) should occur earlier than in an auction without a buyout option. The data set we constructed was clearly imperfect, because it did not capture many factors besides the buyout option which could also conceivably explain differences in bidder behavior: quality of items auctioned, presence and size of accompanying photographs, how prominently the auction is listed by eBay's search engine, etc. Another issue is that a buyout price set excessively high is very unlikely to generate any modification of bidding behavior, as it will effectively be discarded by the buyers. We attempted to control for these factors and others by proxy using the value of the buyout price, the final selling price and the ratio of the buyout price to the final selling price, and our statistical analysis relying on a two sample t-test lead to accepting the hypothesis of earlier first activity with a p-value of 2×10^{-27}.

Taken together, these observations support in our view the use of equilibrium $\mathcal{T}[\nu_{tmp}]$ in the remainder of this paper as a predictor for the outcome of an online auction with a temporary buyout price.

4.2. Equilibrium Analysis of the Permanent Buyout Option. We assume now that the seller uses a fixed permanent buyout price p that remains available until it is exercised or the auction ends. For any bidder with valuation v arriving at time t and observing then a current second-highest bid I_t (see §3.2), consider the following family $\mathcal{P}[.]$ of threshold strategies:

$$\mathcal{P}[\nu](v,t,I_t):\begin{cases}\text{Buyout at } p \text{ immediately} & \text{if } v > \nu(t,I_t)\\ \text{Bid } v \text{ at time } T & \text{if } v \leq \nu(t,I_t)\end{cases}, \tag{5}$$

where $\nu : [0,T] \times [\underline{v}, \bar{v}] \cup \{0\} \to [\underline{v}, \bar{v}]$ is a continuous function. Note that the action of bidding at time T in the definition of $\mathcal{P}[.]$ is clearly a theoretical limit, and would correspond in practice to submitting a first bid as close as possible to the end of the auction, with the goal of denying other bidders the opportunity to respond. Observe also that the definition of $\mathcal{P}[.]$ in (5) corresponds to a single strategy, whereas in the temporary option case (1) defines a class of equivalent strategies, as discussed in §4.1. That is, with a permanent buyout option it is no longer only the timing of the first bid and the highest bid submitted which hold strategic implications for a bidder. The time and value of every bid from a sequence submitted by a participant is now material, because that information affects how I_t evolves over time, which in turn is relevant to whether competitors decide to exercise the buyout option.

The following theorem establishes the existence of a threshold function ν_{prm} such that the symmetric strategy profile $\mathcal{P}[\nu_{prm}]$ constitutes indeed a Bayesian Nash equilibrium, and also provides a characterization.

Theorem 3. *Consider a maximal solution $\tilde{v}(.)$ of the following functional equation on $[0,T] \to [\underline{v}, +\infty)$:*

$$\tilde{v}(t) - p = \mathbf{E}_t \left[e^{-\beta(T-t)} \left(\int_{\underline{v}}^{\tilde{v}(t)} \frac{\prod_{i=1}^{N(t)} F\big(\min(\tilde{v}(t_i), x)\big)}{\prod_{i=1}^{N(t)} F\big(\tilde{v}(t_i)\big)} F(x)^{N(t,T)} dx \right) \right], \tag{6}$$

where the expectation $\mathbf{E}_t$ is with respect to the number $N(t)$ and epochs $t_1, ..., t_{N(t)}$ of arrivals in $[0,t)$ of a non-homogeneous Poisson process with rate $\lambda F\big(\tilde{v}(\tau)\big)$ for $\tau \in [0,t)$, and number $N(t,T)$ of arrivals in $(t,T]$ of a Poisson process with rate λ. Let $\tilde{v}(t,I)$ be a continuous extension of $\tilde{v}(.)$ to $[0,T] \times [\underline{v}, \bar{v}] \cup \{0\}$ such that $\tilde{v}(t,0) = \tilde{v}(t)$ and $\tilde{v}(t,I)$ is non-increasing in I for all t, non-decreasing in t for all I, and define $\nu_{prm}(t,I) = \min\big(\tilde{v}(t,I), \bar{v}\big)$. The symmetric strategy profile $\mathcal{P}[\nu_{prm}]$ is a Bayesian Nash equilibrium for the online auction game with a permanent buyout price p.

Denoting by $\mathcal{R}(\mathcal{P}[\nu])$ the set of best response strategies to the profile where every other player follows strategy $\mathcal{P}[\nu]$, the proof of Theorem 3 first establishes that $\mathcal{P}[\nu_{prm}] \in \mathcal{R}(\mathcal{P}[\nu_{prm}])$ if and only if $\nu_{prm}(t,0)$ satisfies (6). The most challenging part of the proof then consists of proving the existence of a solution to (6); to do so we establish that a generalization of Schauder's fixed point theorem applies to an appropriately defined functional space and continuous mapping on that space. Before discussing the intuition behind Theorem 3 and its qualitative implications, we show how the threshold function characterization (6) specializes to a first-order nonlinear differential equation in the case where valuations follow a uniform distribution:

Proposition 2. *When bidder valuations follow a uniform distribution with cdf F on $[\underline{v}, \bar{v}]$, the threshold function ν_{prm} characterizing the Bayesian Nash equilibrium described in Theorem*

3 satisfies $\nu_{prm}(t,0) = \min\left(\tilde{v}(t), \bar{v}\right)$ *where* $\tilde{v}(t)$ *is the unique solution on* $[0,T]$ *of the differential equation*

$$\frac{d\tilde{v}(t)}{dt} = \frac{\left(\beta + \lambda\left(1 - F\left(\tilde{v}(t)\right)\right)\right)\left(\tilde{v}(t) - p\right)}{1 - e^{-\left(\beta+\lambda\left(1-F\left(\tilde{v}(t)\right)\right)\right)\left(T-t\right)}} \tag{7}$$

with initial value $\nu_{tmp}(0)$ *as defined in (3).*

In the statement of Theorem 3, the requirement that $\nu_{prm}(t,I)$ be non-decreasing for I fixed is intuitive: from the auction still running at time t it can be inferred that $v_i \leq \nu(t_i, I_{t_i})$ for all bidders i with type (v_i, t_i) observing a current second-highest bid $I_{t_i} \leq I$ upon their arrival in $(0,t)$. Consequently as t increases with I fixed, the expected final second highest valuation among all bidders decreases, thus increasing the expected utility from bidding in the auction relative to exercising the buyout option; this effect is compounded with the reduced relative discounting of the utility from bidding as t increases. The requirement that $\nu(t,I)$ be non-increasing in I for every t is likewise easily interpreted: holding t fixed, a higher value of I implies that the expected second highest bid in the auction is higher, which lowers the expected utility from bidding relative to exercising the buyout option. Note that Theorem 3 only provides a stringent characterization of the equilibrium threshold function value $\nu_{prm}(t,I)$ for $I = 0$. This is because when all bidders follow strategy $\mathcal{P}[\nu]$ then on the equilibrium path $I_t = 0$ for all t in $[0,T)$, since all bidders not exercising the buyout option only bid then at time T. Indeed, equation (6) specifies quantitatively the valuation for which an incoming bidder should be indifferent between exercising the option and submitting a regular bid, accounting for the information about the valuations of potential competing bidders provided by the presence of an open buyout option. Other values of $\nu_{prm}(t,I)$ correspond to off-equilibrium path behavior, and are only required to satisfy the monotonicity properties discussed above. While in a strict game-theoretic sense Theorem 3 thus defines multiple equilibria, all of them result in the same equilibrium path and therefore yield the same utility for the bidders and the seller.

As in the temporary case however, there also exists equilibria for the permanent buyout option game other than the one(s) characterized above. Indeed for a threshold function ν satisfying the conditions of Theorem 3 and such that $\nu(t,I) \geq p$ for all (t,I), consider the strategy

$$\mathcal{P}'[\nu](v,t,I_t) \begin{cases} \text{Buyout at time } \tau \geq t & \text{if } v > \nu(\tau, I_\tau) \\ \text{Bid } v \text{ at time } T & \text{if } p < v \leq \nu(\tau, I_\tau) \text{ for all } t \leq \tau \leq T \\ \text{Bid } v \text{ at any time in } [t,T] & \text{if } v \leq p \end{cases}.$$

Note that when following $\mathcal{P}'[\nu]$ a bidder with valuation $v \leq p$ bids at any time, whereas such bidder only bids at T when following $\mathcal{P}[\nu]$. It can be shown however that $\mathcal{P}'[\nu]$ also induces an equilibrium for some function ν. This is because any bidder with a valuation $v \le p$ has no incentive in our model for delaying his bid in order to prevent the published second highest bid I_t from increasing, even if this would potentially trigger the exercise of the buyout option by another bidder: the valuation v' of any bidder exercising the buyout option would be larger than his since buyout exercise at τ implies $v' > \nu(\tau, I_\tau) \ge p \ge v$, therefore that other bidder would win in the auction anyway.

As in the temporary case, we now introduce a game perturbation to argue that $\mathcal{P}[\cdot]$ is a more robust outcome prediction than other equilibria such as $\mathcal{P}'[\cdot]$. Suppose that with probability ϵ each arriving bidder is a *common value* bidder, having type (v,t) and following strategy $\mathcal{P}[\nu_c]$ where $\nu_c(t,0) = \nu_{prm}(t,0)$ and $\nu_c(t,I) = \underline{v}$, $\forall I > 0, t$. That is, a common value bidder exercises the buyout option irrespective of its price as soon as any regular bid is placed in the auction, and if none is placed submits at the end a bid equal to his private valuation. Such a bidder can be rationalized as one whose valuation includes not only an independent private value component as assumed so far, but also a large common value component. That is, any bid placed before the auction end is perceived by that bidder as a signal drastically increasing the estimated value of the item being sold, thus triggering the exercise of the buyout option (see McAfee and McMillan 1987 for background). Again, we point out that while the perturbation just described is completely ad-hoc, we are only interested here in its impact when the perturbation probability ϵ is arbitrarily small. Denoting by $G^{(\epsilon)}$ the corresponding perturbed game, we specifically show the following result:

THEOREM 4. *For any $\epsilon > 0$, the game $G^{(\epsilon)}$ does not have any Bayesian Nash equilibrium where bidders bid at any time $\tau < T$ (e.g. play $\mathcal{P}'[\cdot]$). In addition, the only Bayesian Nash equilibria of the game $G^{(\epsilon)}$ are such that normal (non common value) bidders play strategy profile $\mathcal{P}[\nu_{prm}]$ defined in Theorem 3.*

The underlying intuition is that a normal bidder is strictly better off placing a bid near time T because bidding earlier may cause a common value bidder, who could have otherwise lost in the auction, to exercise the buyout option. Any equilibrium strategy where a normal bid is placed before time T is thus eliminated, however strategy profile $\mathcal{P}[\nu_{prm}]$ still constitutes an equilibrium of the perturbed game for normal bidders. Consequently, only the strategies characterized in Theorem 3 survive the above game perturbation, no matter how small.

Finally, we have conducted a simple empirical study in order to validate some of our model predictions in the permanent case as well, which is described more extensively in Section A.7 of our online supplement. In summary, we collected bidding data from actual auctions of the same category of items considered in the temporary case (iPod music players and accessories), this time from the site Yahoo! (which features a permanent buyout option), and likewise identified a testable implication of our analysis. Specifically, we observed that the equilibrium analysis in our model does not generate any prediction of bidding times in an online auction without a buyout option, while Theorems 3 and 4 do imply that with a permanent buyout option bidders not exercising the option will submit theirs bids as late as possible. Consequently, in auctions featuring a permanent buyout option that is not exercised, bids should be submitted later on average than in an auction without a buyout option. We applied the same proxy controls as described in §4.1 in order to mitigate the impact of our dataset flaws. While our statistical analysis lead to accepting the hypothesis of later average bidding times, the associated p-value was higher at 0.0963, so that our associated confidence level was much lower than for the test we conducted in the temporary case. This was hardly surprising however, as the dataset we were able to construct in the permanent case was significantly smaller (to date the site Yahoo! Auctions receives significantly less traffic than eBay), and also because auctions on Yahoo! feature an automatic bidding deadline extension mechanism (see §3.3), a deviation from our model known to impact the concentration of bids near the end of the auction (Roth and Ockenfels 2002).

Taken together, these observations support in our view the use of equilibrium $\mathcal{P}[\nu_{prm}]$ in the remainder of this paper as a predictor for the outcome of an online auction with a permanent buyout price.

4.3. Seller's Revenue Optimization Problem. We now consider the problem of finding the buyout price p maximizing the seller's expected discounted revenue from a temporary (resp. permanent) buyout price auction when all bidders follow the equilibrium strategy $\mathcal{T}[\nu_{tmp}]$ (resp. $\mathcal{P}[\nu_{prm}]$). Note that p is the only decision variable we consider here (see Vakrat and Seidmann 2001 and Gallien 2005 for optimization studies focusing on the variables T and $\bar{v}$).

4.3.1. Formulation and Numerical Solution. We first consider the temporary case. Making the dependence of the threshold function on p explicit from now on and conditioning on both the arrival time and the action of the first bidder, the problem can be stated mathematically as

$$\begin{aligned}\max_{p\in[\underline{v},\bar{v}]} \mathbf{E}[U^S_{tmp}(p)] = & \int_0^T e^{-\alpha T}\mathbf{E}_t[\max(\underline{v}, v^{(2)}_{N(t,T)+1})|v_1 \le \nu_{tmp}(p,t)]F(\nu_{tmp}(p,t))\lambda e^{-\lambda t}dt \\ + & \int_0^T e^{-\alpha t}p\big(1-F(\nu_{tmp}(p,t))\big)\lambda e^{-\lambda t}dt,\end{aligned} \tag{8}$$

where the expectation $\mathbf{E}_t$ in the first integrand is with respect to the number $N(t,T)$ of arrivals in interval $(t,T]$ of a Poisson process with rate λ and the second highest value $v^{(2)}_{N(t,T)+1}$ among $N(t,T)+1$ independent draws $v_1,...,v_{N(t,T)+1}$ from the valuation distribution with cdf F, where by convention $v_1^{(2)} = 0$ – note that the first and second integrals in (8) correspond respectively to the seller's expected revenue when the first bidder submits a regular bid upon his arrival and when he exercises the buyout option.

Turning next to the permanent buyout option, let $\nu_{prm}(p,t)$ denote the value of the threshold function on the equilibrium path (i.e. the variable $I_t = 0$ is omitted). In equilibrium, the arrivals of bidders who will exercise the buyout option follow a non-homogeneous Poisson process with instantaneous rate $\lambda\big(1 - F\big(\nu_{prm}(p,t)\big)\big)$, and we denote its counting measure by N_{buy}. Likewise, the arrivals of bidders who will wait until the end of the auction to submit a bid follow a non-homogeneous Poisson process with instantaneous rate $\lambda F\big(\nu_{prm}(p,t)\big)$, and we denote its counting measure by N_{bid}. As a result, the probability that the buyout option will not be exercised is $P(N_{buy}(T) = 0) = \exp(-\lambda \int_0^T \big(1 - F\big(\nu_{prm}(p,t)\big)\big)dt)$, and the problem can be stated as

$$\begin{aligned} \max_{p\in[\underline{v},\bar{v}]} \mathbf{E}[U^S_{prm}(p)] = & \int_0^T e^{-\alpha t} p\lambda\big(1 - F(\nu_{prm}(p,t))\big)e^{-\lambda\int_0^t \left(1-F\left(\nu_{prm}(p,\tau)\right)\right)d\tau}dt \\ + \quad & e^{-\lambda\int_0^T \left(1-F\left(\nu_{prm}(p,t)\right)\right)dt} e^{-\alpha T}\mathbf{E}[\mathbf{1}_{\{N_{bid}(T)>0\}}\max(\underline{v}, v^{(2)}_{N_{bid}(T)})|v_i \le \nu_{prm}(p,t_i)\ \forall i], \end{aligned} \tag{9}$$

where the expectation $\mathbf{E}$ is with respect to the number $N_{bid}(T)$ and epochs $t_1,...,t_{N_{bid}(T)}$ of arrivals in $[0,T]$ of the second Poisson process defined above, and second highest value $v^{(2)}_{N_{bid}(T)}$ among $v_1,...,v_{N_{bid}(T)}$ (by convention $v_0^{(2)} = v_1^{(2)} = 0$), where the i-th valuation v_i follows a distribution with cdf $F_i(v) = F(v)/F(\nu_{prm}(p,t_i))$. The first term in (9) is equal to the seller's expected discounted revenue from the option, while the second term is the expected discounted revenue from regular bidding, which only occurs if the buyout option is not exercised.

While solving analytically these optimization problems in the general case seems particularly challenging, computing through a line search over p a numerical solution to (8) and, in the special case of uniformly distributed valuations, to (9) is relatively straightforward: for each value of p, one may numerically solve (2) for $\nu_{tmp}(p,t)$ and (7) for $\nu_{prm}(p,t)$; the seller's expected utility can then be estimated through Monte-Carlo simulation by generating repeated random bidder arrival streams $\{(v_1,t_1),(v_2,t_2),...\}$. This is the method we implement to obtain the numerical results reported later in §4.4. The other method we have followed to study the difficult stochastic optimization problems (8) and (9) is an asymptotic analysis, which is discussed next.

4.3.2. Asymptotic Analysis. We were able to characterize analytically the limits of the solutions p^*_{tmp} and p^*_{prm} to (8) and (9) respectively for various asymptotic regimes of the bidders' arrival rate λ, seller's time sensitivity α, and bidders' time sensitivity β. While their somewhat lengthy and technical derivations are relegated to the online supplement, we provide here a summary of these results in Table 2, which uses the notations

$$\begin{cases} p_1 \triangleq \arg\max_{p\in[\underline{v},\bar{v}]} p(1-F(p)) \\ p_2 \triangleq \arg\max_{p\in[\underline{v},\bar{v}]} \big(p(1-F(p)) + \underline{v}F(p)\big) \\ p_3(\mu) \triangleq \arg\max_{p\in[\underline{v},\bar{v}]} \frac{p(1-F(p))}{\mu+1-F(p)} \text{ for } \mu\in[0,+\infty] \end{cases}. \tag{10}$$

Before interpreting p_1, p_2 and $p_3(\mu)$, we first observe that for these quantities to be uniquely defined we need to impose in the following some additional mild assumptions on the distribution function $F(\cdot)$ – a possible sufficient condition is for $F(\cdot)$ to be strictly increasing on $[\underline{v},\bar{v}]$, convex and continuously differentiable. It is then easy to prove (see Lemma 13 in online supplement) that $p_1 \le p_2$, $p_1 \le p_3(\mu)$ for any μ, $p_3(\mu)$ is decreasing in μ, and $\lim_{\mu\to+\infty} p_3(\mu) = p_1$, with the last statement justifying the notational extension $p_3(+\infty)$.

		Seller time sensitivity	
		Low ($\alpha\to 0$)	High ($\alpha\to\infty$)
Bidder time sensitivity	Low ($\beta\to 0$)	any p in $[\underline{v},\bar{v}]$	$p^*_{tmp}, p^*_{prm} \to \underline{v}$
	High ($\beta\to\infty$)	$p^*_{tmp}, p^*_{prm} \to p_2$	$p^*_{tmp}, p^*_{prm} \to p_1$

(a) Low demand rate limit ($\lambda\to 0$)

Seller time sensitivity	
Low ($\alpha\to 0$)	High ($\alpha\to\infty$)
$p^*_{tmp}, p^*_{prm} \to \bar{v}$	$p^*_{tmp} \to p_1$, $p^*_{prm} \to p_3(\mu)$ where $\mu = \lim_{\lambda\to\infty} \frac{\alpha}{\lambda}$

(b) High demand rate limit ($\lambda\to\infty$)

Table 2 Optimal buyout prices in asymptotic regimes

In Table 2 (a), each entry $(\lambda\to 0, \alpha\to A, \beta\to B)$ with $(A,B)\in\{0,\infty\}^2$ corresponds more precisely to the regime $\lambda\to 0$, $\alpha = f_1(\lambda)$ and $\beta = f_2(\lambda)$ where $f_i : [0,\infty)\to[0,\infty)$, $i\in\{1,2\}$ are any functions such that $\lim_{x\to 0} f_1(x) = A$ and $\lim_{x\to 0} f_2(x) = B$. A first interesting observation is that, in contrast with Table 2 (a), the limit statements in Table 2 (b) are independent of the bidders' time sensitivity parameter β. More precisely, each entry $(\lambda\to\infty, \alpha\to A)$ with $A\in\{0,\infty\}$ in Table 2 (b) corresponds to the asymptotic regime $\lambda\to\infty$, $\alpha = f_1(\lambda)$ and $\beta = h(\lambda)$, where f_1 is any function $[0,\infty)\to[0,\infty)$ such that $\lim_{x\to+\infty} f_1(x) = A$ and h is *any* non-negative function of λ. Our interpretation is that the seller's utility only depends on bidders' time sensitivity via the buyout threshold functions, where β discounts the utility from bidding relative to exercising the buyout option. However, the utility from bidding in an auction with a high bidder arrival rate is already made negligible by the very high associated level of competition, consequently the effect of β on the optimal buyout price vanishes then. That is, in this dynamic setting bidders' time sensitivity effectively acts as a negative adjustment to their market power, and thus looses leverage in the asymptotic regimes of Table 2 (b) where that market power is low.

Tables 2 (a) and (b) can be further interpreted as follows. The case $p^*_{tmp}, p^*_{prm} \to \underline{v}$ effectively amounts to using a fixed price mechanism, since no bidding activity will ever occur then; this is optimal for a very impatient seller facing very few time-insensitive bidders (i.e. $\lambda \to 0, \alpha \to \infty, \beta \to 0$). Indeed, a seemingly large number of auction listings on eBay now feature only a "Buy It Now" option and no "Place Bid" option, providing anecdotal evidence for the relevance of this case in practice. At the other extreme, the case $p^*_{tmp}, p^*_{prm} \to \bar{v}$ is equivalent to an auction without a buyout option since the buyout price is never exercised then. That is, a patient seller ($\alpha \to 0$) with high market power ($\lambda \to \infty$) finds it beneficial to not use any buyout option at all and only rely on a traditional bidding mechanism – there are clearly many examples of such sellers on auction sites as well. These results are thus reminiscent of those obtained by Harris and Raviv (1981), who study a mechanism design model in which the seller should use an auction when demand exceeds supply but a posted price otherwise (see also Gallien 2006). In our model, the relative values of the seller's and bidders' time sensitivity (α and β) and the expected number of bidders λ effectively capture the ratio between supply and demand and the seller's market power, and the hybrid mechanism relying on both bidding and posted price enabled by the buyout option makes for a continuous, smoother transition between those two mechanisms.

More specifically, consider first market environments where demand for the auctioned item is low but bidders are highly time-sensitive ($\lambda \to 0$, $\beta \to \infty$). Such bidders gain negligible utility from submitting a regular bid (which entails winning almost surely but waiting up to time T), and thus always exercise the buyout option provided the buyout price is no larger than their valuation (formally $\nu_{tmp}(p,t), \nu_{prm}(p,t) \to p$ for all t). A highly time-sensitive seller ($\alpha \to \infty$) also gets zero utility from selling the product at time T, and thus offers then the buyout price p_1 defined in (10), which maximizes his expected revenue from the event that the first (and most likely only) bidder exercises the buyout option. In the same environment however, a seller with low time-sensitivity ($\alpha \to 0$) can potentially wait up to the auction end to sell the product, and hence finds it optimal to offer a higher buyout price $p_2 \geq p_1$ maximizing his utility from the event that the first/only bidder either exercises the buyout option or bids, in which case the product sells for $\underline{v}$ at the end of the auction (see (10)).

Next, the regime ($\lambda, \alpha, \beta \to 0$) corresponds to the case when both the seller and the bidders have low time-sensitivity and the bidder arrival rate is small. Any incoming bidder is unlikely to face any competition then (as $\lambda \to 0$), and consequently a buyout option with price $p > \underline{v}$ is never exercised. Therefore in the limit any incoming bidder either exercises the option if the buyout price p is set to $\underline{v}$, or bids in the auction if $p > \underline{v}$ and gets the item at time T, still for $\underline{v}$ (since the probability of two or more bidder arrivals becomes negligible). In either case, the seller gains the same utility since he is not affected by the time of sale ($\alpha \to 0$), and is thus indifferent between any buyout price $p \in [\underline{v}, \bar{v}]$. More formally we show that for any $p' \in [\underline{v}, \bar{v}]$

$$\lim_{\substack{\lambda\to 0\\ \alpha=f_1(\lambda)\\ \beta=f_2(\lambda)}} \frac{\max_{p\in[\underline{v},\bar{v}]} \mathbf{E}[U^S_{tmp}(p)]}{\mathbf{E}[U^S_{tmp}(p')]} = 1 \text{ and } \lim_{\substack{\lambda\to 0\\ \alpha=f_1(\lambda)\\ \beta=f_2(\lambda)}} \frac{\max_{p\in[\underline{v},\bar{v}]} \mathbf{E}[U^S_{prm}(p)]}{\mathbf{E}[U^S_{prm}(p')]} = 1,$$

where $f_i : [0,\infty) \to [0,\infty)$, $i \in \{1,2\}$ are any functions such that $\lim_{x\to 0} f_i(x) = 0$. In words, the additional utility obtained by choosing the optimal buyout price relative to using any buyout price becomes asymptotically negligible in that regime.

Finally, the only regime where p^*_{tmp} and p^*_{prm} converge to different limits is the one in which a very impatient seller faces a high demand ($\lambda, \alpha \to \infty$). While such a seller could sell the product for $\bar{v}$ via an auction, that outcome would only occur at time T and would therefore give him negligible utility relative to a buyout option exercise. Hence, in the temporary case where the buyout option is only available to the first bidder, the seller offers the buyout price p_1 maximizing his expected utility from the event that the buyout option is exercised by that bidder. A key observation is that in the temporary case the buyout price affects only the probability, but not the time, of buyout exercise. In contrast, a permanent buyout option is available to all arriving bidders until exercised, and is therefore exercised with probability 1 in this limiting regime, provided $p < \bar{v}$. However, the buyout price in the permanent case does affect the time at which the option is exercised, confronting the seller with the tradeoff of selling time vs. selling price. The optimal balance in this tradeoff is dictated by the relative values of α and λ, explaining the impact of the ratio $\mu = \lim_{\lambda\to+\infty} \frac{\alpha}{\lambda}$ on the optimal permanent buyout price $p_3(\mu)$ shown in (10): as noted earlier $p_3(\mu)$ is decreasing with μ, which reflects that a seller more time sensitive or facing fewer bidders should reduce the permanent buyout price. The fact that $\lim_{\mu\to+\infty} p_3(\mu) = p_1$ reflects that in the extreme case where the seller's time sensitivity is very high relative to the bidders' arrival rate, the seller obtains negligible utility from waiting for a subsequent bidder beyond the first one and should then, as in the temporary case, maximize the revenue obtained from the first bidder alone. Finally, we point out that the arg max operand defining $p_3(\mu)$ in (10) is equal to the seller's expected discounted revenue $J(p)$ obtained by offering the item for a fixed posted price of p to an infinite Poisson arrival stream of bidders with valuation distributions given by F – this is easily seen by substituting terms and solving for $J(p)$ in

$$J(p) = E[e^{-\alpha X}]\left((1 - F(p))p + F(p)J(p)\right),$$

where X denotes an exponential random variable with mean $1/\lambda$. That is, in this limiting regime the impact of the finite time horizon T and regular bidding on the optimal buyout price are effectively obliviated by the high seller time sensitivity and bidders' arrival rate.

4.4. Numerical Experiments. In this section we compare the equilibrium behavior, optimal buyout price and seller's revenue associated with the temporary and permanent buyout options, drawing on both numerical experiments and our theoretical results from the previous subsections. While we only report here the results from a small number of experiments due to length restrictions, we found those to be representative of a larger set of scenarios.

A first insightful exercise is to compare the bidders' equilibrium buyout threshold functions ν_{tmp} and ν_{prm} (see statements of Theorems 1 and 3) corresponding to the same buyout price and market environment. For illustration purposes, Figure 1 shows a plot of these two functions for the specific case $p_{tmp} = p_{prm} = 350$, $\lambda = 0.25$, $T = 16$, $\beta = 0.03$; as in all other experiments to be discussed in this section we assume that valuations are uniformly distributed on $[50, 500]$.

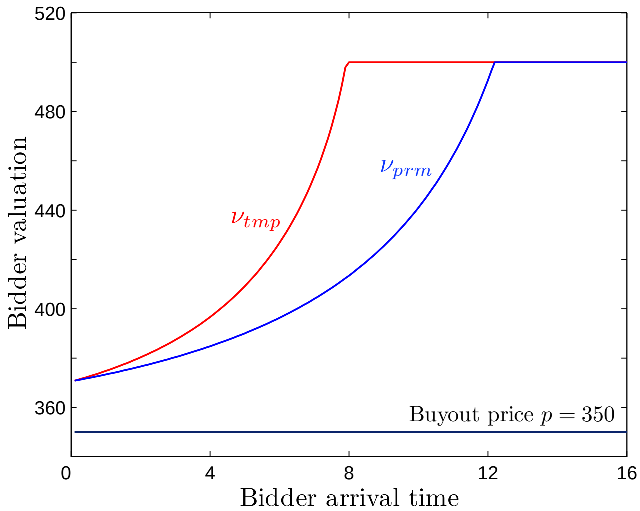

Figure 1 Equilibrium threshold valuation in temporary and permanent buyout price auction

A first observation is that both curves shown in Figure 1 are non-decreasing, which can be easily established for the general case from (2) and (6). That is, either type of buyout option remaining open as time goes by indicates reduced competition among bidders participating in the auction and therefore progressively makes the buyout option less attractive relative to submitting a regular bid, so that fewer bidders will decide to exercise it. The temporary threshold function ν_{tmp} does lie above the permanent threshold function ν_{prm} however, suggesting that the effect just described is less pronounced with a permanent option than with a temporary option. Indeed, when participants follow the equilibrium strategy $\mathcal{T}[\nu_{tmp}]$ described by (1) and Theorem 1, the fact that a temporary option is still open when a bidder arrives indicates to him that he is the first bidder and that the only competition he is likely to face should he submit a regular bid will come from bidders who are yet to arrive. On the other hand, under the strategy profile $\mathcal{P}[\nu_{prm}]$ described by (5) and Theorem 3, if a permanent option is still open when a bidder arrives he can only infer that all the bidders who have already arrived have valuations lower than the value of the threshold valuation at the time of their respective arrivals. Consequently, for such a bidder the decision to submit a regular bid appears less attractive relative to exercising the buyout option than it is for a bidder facing an open temporary option in circumstances that are otherwise the same. As a result, with identical buyout prices more bidders will tend to exercise a permanent option than a temporary one. Finally, note that the initial values $\nu_{tmp}(0)$ and $\nu_{prm}(0)$ shown in Figure 1 are identical, which is intuitive but can also be established analytically by calculating the right-hand sides of (2) and (6) for $t=0$.

Next we compare the optimal permanent and temporary buyout prices for the special case when participating bidders are very impatient, i.e. $\beta \to \infty$. The optimal temporary (resp. permanent) buyout price is obtained by solving numerically the optimization problem obtained when substituting the very impatient bidder condition in (8) (resp. (9)). These optimal buyout prices are plotted in Figure 2 for various values of the bidder arrival rate λ and seller time-sensitivity α.

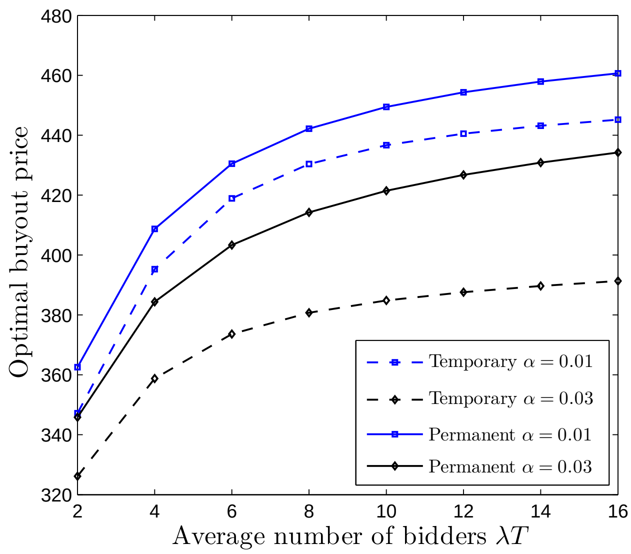

Figure 2 **Optimal temporary and permanent buyout prices with impatient bidders**

The graph in Figure 2 confirms the intuition that both optimal buyout prices should increase with the bidder arrival rate and decrease with seller time-sensitivity. Although not reported here, other experiments show that these prices also decrease with the bidders' time-sensitivity. Another observation is that the optimal buyout price is higher with a permanent option than with a temporary one; our explanation follows from examining the individual terms of the equation for the seller's total expected discounted revenue

$$\mathbf{E}[pe^{-\alpha\tau_{buy}}|\text{buyout}]P(\text{buyout}) + \mathbf{E}[e^{-\alpha T}\mathbf{1}_{\{N(T)>0\}}\max(\underline{v}, v^{(2)}_{N(T)})|\text{no buyout}]P(\text{no buyout}), \qquad (11)$$

where the first term is the expected discounted revenue from the buyout auction (τ_{buy} denotes the conditional buyout exercise time), while the second is the expected discounted revenue from regular bidding. For a given buyout price p, the permanent buyout option is exercised with higher probability and, conditional on its exercise, later on average than the temporary option (it may be exercised by other bidders besides the first one). This suggests that the price maximizing the first term alone in (11), which is a unimodal function of the buyout price, will be larger with a permanent option than with a temporary one. Figure 1 also indicates that for any given buyout price both the expectation and the probability forming the expected revenue from bidding (second term in (11)), which is increasing in the buyout price, will be smaller with a permanent option than with a temporary one. The buyout price value at which the marginal decrease in expected buyout revenue equals the marginal increase in expected bidding revenue in (11) should thus be higher with a permanent option than with a temporary option. Finally, note that the higher the seller time-sensitivity α, the larger the difference between the conditional buyout revenues $\mathbf{E}[pe^{-\alpha\tau_{buy}}|\text{buyout}]$ for permanent and temporary options, explaining the larger difference between optimal permanent and temporary buyout prices observed in Figure 2.

Our last set of experiments focuses on the seller's relative gain in utility from an auction with temporary and permanent buyout options over an auction with no buyout price, that is $(\mathbf{E}[U^S_{tmp}(p^*_{tmp})] - \mathbf{E}[U^S_{nb}])/\mathbf{E}[U^S_{nb}]$ or $(\mathbf{E}[U^S_{prm}(p^*_{prm})] - \mathbf{E}[U^S_{nb}])/\mathbf{E}[U^S_{nb}]$, where $\mathbf{E}[U^S_{tmp}(p^*_{tmp})]$ and $\mathbf{E}[U^S_{prm}(p^*_{prm})]$ denote the seller's expected utility from an auction with optimal temporary and permanent buyout options respectively, and $\mathbf{E}[U^S_{nb}]$ the seller's expected utility from the basic auction mechanism without a buyout price described in §3.2. As described in §4.3 the optimal buyout prices p^*_{tmp} and p^*_{prm} are obtained by performing a simulation-based line search; for all values estimated by simulation, the true value is within 1% of the estimate with 95% confidence. The results from these experiments are plotted in Figure 3 and 4, which show the seller's relative utility increase just defined for both option types in various environments. A first observation is that, as intuition suggests, the relative gain from both types of buyout option generally increases with both the seller's time sensitivity α and the bidders' time sensitivity β – the possibility of selling the item earlier is more valuable for a time-sensitive seller, and bidders with a high time-sensitivity are willing to pay

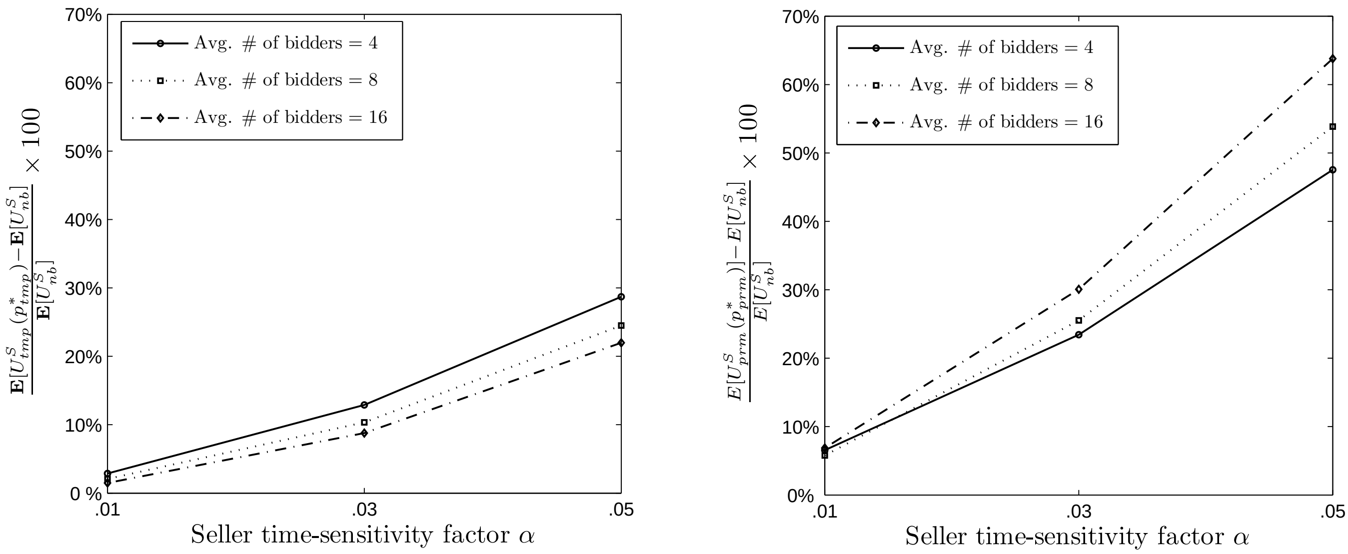

(a) Temporary buyout option (b) Permanent buyout option

Figure 3 Relative increase in seller's utility from a buyout option ($\beta = 0.03$)

more if they can get the product earlier. Figure 4 suggests however that the impact of the bidders' time sensitivity on the relative utility gain from a buyout option becomes insignificant when the expected number of bidders λT becomes moderately large. On the other hand, the expected utility gain from a buyout option always seems to increase substantially with the seller's time sensitivity, independently of the expected number of bidders. Our interpretation is that while the seller's time sensitivity directly impacts his utility, the effect of the bidders' time sensitivity is more indirect in that it only affects the bidders' relative preference between the buyout option and the regular online auction, without otherwise affecting the seller's discounted revenue from either alternative. Moreover, when the number of bidders is large, affecting the probability that a single one of them will exercise the buyout option for a given time-sensitivity β becomes relatively easier.

Another important finding is that the optimal seller's utility derived from a permanent buyout option is always larger than that obtained with a temporary buyout option, as can be seen from comparing the two vertical scales in Figure 3 and 4; although unable to show this analytically, we have more generally observed this in all the experiments we have conducted besides the ones reported here. Within the strict boundaries of our model definition, a permanent buyout option seems like a more powerful instrument than a temporary one, because it allows to leverage the time-sensitivity of all participating bidders as opposed to only the first one. This interpretation ignores some of the features of actual online auctions that our model does not capture however, and we come back to this issue in §6.

Finally, we observe that while the increase in seller's utility achieved by introducing a temporary buyout option (Figures 3(a) and 4(a)) is decreasing in the bidder arrival rate, with a permanent

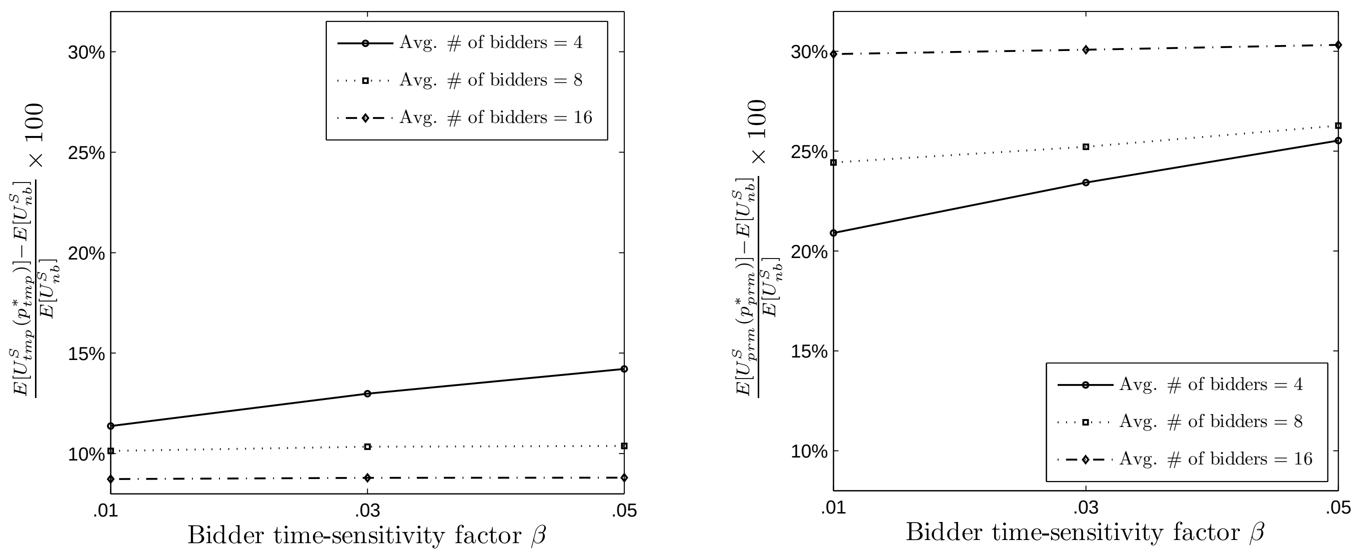

(a) Temporary buyout option (b) Permanent buyout option

Figure 4 Relative increase in seller's utility from a buyout option ($\alpha = 0.03$)

buyout option the exact opposite occurs (Figures 3(b) and 4(b)). Our interpretation is that since a temporary buyout option is only available to the first bidder, its relative impact diminishes in an environment with a high expected number of participants. On the other hand, a permanent option is potentially available to all arriving bidders and thus its relative impact does increase with the expected number of bidders.

5. Dynamic Buyout Prices

In this section we study the mechanism obtained when the buyout price, either temporary or permanent, is no longer constant but instead varies over time according to a pre-announced trajectory $[p(t)]_{t\in[0,T]}$. While we are not aware of any actual auction site currently implementing such a feature, our goal is to develop a theoretical analysis providing some prediction for what the outcome of such mechanism is likely to be (in §5.1), and bound the maximum expected revenue achievable by the seller when setting this buyout price trajectory optimally (in §5.2).

5.1. Outcome Prediction. In an auction with a temporary buyout price following a dynamic trajectory $[p(t)]_{t\in[0,T]}$, consider the extension of strategy $\mathcal{T}[\nu]$ obtained for any function $\nu:[0,T]\to[\underline{v},\bar{v}]$ by substituting $p(t)$ with p in the first line of (1); for notational simplicity we will still refer to the resulting strategy as $\mathcal{T}[\nu]$. The following result establishes that any non-decreasing continuous threshold function ν can be supported by some price trajectory in equilibrium:

Theorem 5. *For any non-decreasing continuous function $\nu:[0,T]\to[\underline{v},\bar{v}]$, define function $p:[0,T]\to[\underline{v},\bar{v}]$ as*

$$p(t) = \nu(t) - e^{-(\lambda+\beta)(T-t)} \int_{\underline{v}}^{\nu(t)} e^{\lambda(T-t)F(x)} dx. \tag{12}$$

The symmetric strategy profile $\mathcal{T}[\nu]$ is then a Bayesian Nash equilibrium for the auction with temporary buyout price trajectory $[p(t)]_{t\in[0,T]}$.

Likewise, in an auction with a permanent buyout price following trajectory $[p(t)]_{t\in[0,T]}$, for any function $\nu:[0,T]\times[\underline{v},\bar{v}]\cup\{0\}\to[\underline{v},\bar{v}]$ we can consider the extension of strategy $\mathcal{P}[\nu]$ obtained by substituting $p(t)$ with p in the first line of (5), and keep using the same notation. The following result is the exact analogue of Theorem 5 for the case of a permanent buyout option:

THEOREM 6. *For any continuous function $\nu:[0,T]\times[\underline{v},\bar{v}]\cup\{0\}\to[\underline{v},\bar{v}]$ such that $\nu(t,0)\triangleq\nu(t)$ is non-decreasing in t and $\nu(t,I)$ is decreasing in I for all t, define function $p:[0,T]\to[\underline{v},\bar{v}]$ as*

$$p(t)=\nu(t)-e^{-\beta(T-t)}\mathbf{E}_t\left[\int_{\underline{v}}^{\nu(t)}\frac{\prod_{i=1}^{N(t)}F\big(\min(\nu(t_i),x)\big)}{\prod_{i=1}^{N(t)}F\big(\nu(t_i)\big)}\big(F(x)\big)^{N(t,T)}dx\right] \tag{13}$$

where the expectation $\mathbf{E}_t$ is with respect to the number $N(t)$ and epochs $t_1,...,t_{N(t)}$ of arrivals in $[0,t)$ of a non-homogeneous Poisson process with rate $\lambda F\big(\nu(\tau)\big)$ with $\tau\in[0,t)$, and number $N(t,T)$ of arrivals in $(t,T]$ of a Poisson process with rate λ. The symmetric strategy profile $\mathcal{P}[\nu]$ is then a Bayesian Nash equilibrium for the auction with permanent buyout price trajectory $[p(t)]_{t\in[0,T]}$.

Theorems 5 and 6 have similar interpretations: for both the temporary and permanent case, any threshold function ν that is continuous and non-decreasing with time corresponds to a buyout price trajectory such that the strategy profile $\mathcal{T}[\nu]$ or $\mathcal{P}[\nu]$ forms an equilibrium. In fact, the negative of the second terms in the right-hand side of (12) and (13) both represent the expected utility that a bidder arriving at time t and having a valuation equal to the threshold would obtain by submitting a regular bid (as opposed to exercising the buyout option) in the corresponding game. Therefore, both (12) and (13) express that the buyout price $p(t)$ they define is such that a bidder arriving at time t with a valuation equal to the threshold $\nu(t)$ would be indifferent between submitting a regular bid and exercising the buyout option (provided it is still open) at that price. However, setting the buyout price $p(t)$ according to (12) or (13) is only a necessary condition in general, and would not eliminate alone the possibility that a bidder could benefit from waiting beyond his arrival before choosing between these two options – this could occur for example if the buyout price is known to substantially decrease in the future, and would give rise to a competitive optimal stopping situation in which neither strategy $\mathcal{T}[\nu]$ or $\mathcal{P}[\nu]$ would form an equilibrium. Theorem 5 and 6 actually establish in their respective settings that no rational bidder will ever find such wait to be more profitable a priori than acting immediately when the target valuation threshold is non-decreasing over time. Note that this does not imply that the buyout price itself is non-decreasing – in fact, for a constant valuation threshold $\nu(t)=\nu\in[\underline{v},\bar{v}]$, which satisfies the conditions of Theorem

5, the price trajectory defined by (12) is decreasing. Only, in the incoming bidders' assessment it does not decrease fast enough for the possible utility increase derived from waiting to strictly overcome time discounting and the risk associated with the arrival of another bidder while the option is still open.

Note that, as is the case with static buyout options, there may exist other equilibria for the temporary and permanent dynamic buyout price games besides those characterized here. In contrast with the static buyout case unfortunately, we have not been able to develop any formal robustness results rationalizing the use for outcome prediction of these specific equilibria among all possible ones. We do however make the observation that the following form of reciprocal holds for Theorems 5 and 6, as should be clear from their respective proofs: for every continuous valuation threshold curve ν that is strictly decreasing with time on some interval, there exist bidders whose best response to the symmetric profile $\mathcal{T}[\nu]$ (resp. $\mathcal{P}[\nu]$) will not be $\mathcal{T}[\nu]$ (resp. $\mathcal{P}[\nu]$). This suggests that any equilibrium we may be ignoring is likely to involve strategic and possibly risky waiting behavior relative to exercising the buyout option, which in practice may be unattractive to some bidders for reasons that our model does not capture (e.g. cost of auction monitoring efforts).

5.2. Seller's Optimization Problem. In this subsection we study the maximum expected discounted revenue achievable by the seller through the choice of a temporary or permanent buyout price trajectory $[p(t)]_{t\in[0,T]}$, using the equilibria characterized in Theorems 5 and 6 as a prediction of the relevant game outcome.

An important implication of Theorems 5 and 6 is that, within the range of equilibria considered, finding an optimal price trajectory $[p(t)]_{t\in[0,T]}$ exactly corresponds to finding its associated continuous and non-decreasing threshold function $\nu:[0,T]\to[\underline{v},\bar{v}]$ subject to either (12) or (13). Denoting by $\mathcal{C}^+$ the set of all such functions and starting with the case of a temporary option, for $\nu\in\mathcal{C}^+$ and $[p(t)]_{t\in[0,T]}$ given by (12), the seller's expected discounted revenue conditional on the first bidder arriving at $t_1=t$ when all bidders follow strategy $\mathcal{T}[\nu]$ is given by

$$\begin{aligned} u_{tmp}(\nu(t),t) \triangleq \mathbf{E}[U^S_{tmp}(\nu)|t_1=t] &= e^{-\alpha T}\mathbf{E}_t[\max(\underline{v}, v^{(2)}_{N(t,T)+1})|v_1\le\nu(t)]F(\nu(t)) \\ &\quad + e^{-\alpha t}\left(\nu(t) - e^{-(\lambda+\beta)(T-t)}\int_{\underline{v}}^{\nu(t)} e^{\lambda(T-t)F(x)}dx\right)\left(1-F(\nu(t))\right) \end{aligned} \tag{14}$$

where the definition of $\mathbf{E}_t$, $N(t,T)$ and $v^{(2)}_{N(t,T)+1}$ is the same as in (8) – note that the instantaneous buyout price $p(t)$ has been substituted with the right-hand side of (12), and that the notation $u_{tmp}(\nu(t),t)$ introduced shows explicitly that the right-hand side of (14) only depends on the value of ν at t. The seller's revenue maximization problem can thus be stated as

$$Z^*_{tmp} \triangleq \sup_{\nu\in\mathcal{C}^+}\mathbf{E}[U^S_{tmp}(\nu)] = \sup_{\nu\in\mathcal{C}^+}\int_0^T u_{tmp}(\nu(t),t)\lambda e^{-\lambda t}dt. \tag{15}$$

The next proposition establishes that a discretized version of problem (15) provides an upper bound for the maximum seller's expected discounted revenue Z^*_{tmp} just defined:

PROPOSITION 3. *Consider any partition $\tau \triangleq (\tau_j)_{j\in\{0,...,m\}}$ of $[0,T]$ into m subintervals such that $\tau_0 = 0 < \tau_1 < ... < \tau_m = T$, define $\triangle\tau_j \triangleq \tau_{j+1} - \tau_j$ for $j \in \{0,...,m-1\}$ and let $\triangle\tau \triangleq \max_j \triangle\tau_j$ be the mesh size of τ. Then*

$$Z^*_{tmp} \le \bar{Z}_{tmp}(\tau) \triangleq \begin{array}{ll} \max\limits_{(\nu_j)_{j\in\{0,...,m\}}} & \sum\limits_{j=0}^{m-1} u_{tmp}(\nu_j, \tau_j)\lambda e^{-\lambda\tau_j}\triangle\tau_j \\ \textit{subject to:} & \underline{v} \le \nu_{j-1} \le \nu_j \le \bar{v} \textit{ for all } j \in \{1,...,m\} \end{array} \tag{16}$$

From a practical standpoint, Proposition 3 provides a way to construct an upper bound for the seller's maximum expected discounted revenue by solving a nonlinear program. Note however that the function u_{tmp} appearing in the objective of (16) may not be always easy to express analytically, because of the expectation $\mathbf{E}_t$ in (14). Also, we do not provide here any description of the relationship between the mesh size of a partition τ (or size of nonlinear program (16)) and the quality of upper bound $\bar{Z}_{tmp}(\tau)$. For our numerical experiments in §5.3, we focus on the special case of uniform valuations, for which a closed-form expression for u_{tmp} is readily derived.

Turning now to the case of a permanent buyout option, the seller's revenue maximization problem can be stated similarly as $\sup_{\nu\in\mathcal{C}^+} \mathbf{E}[U^S_{prm}(\nu)]$, where $\mathbf{E}[U^S_{prm}(\nu)]$ is obtained by substituting p with (13) and $\nu_{prm}(p,t)$ with $\nu(t)$ in (9). While we were able to derive an upper bound for that optimization problem using some approximations and an approach similar to that employed when deriving $\bar{Z}_{tmp}(\tau)$, those approximations were quite coarse. Consequently, the resulting bound proved too loose to support any assertive statement, as evidenced by the fact that the piecewise constant solution obtained by solving the problem analogous to (16) for the permanent case performed significantly worse in all our simulation experiments than all other policies tested, including not using a buyout price at all. Consequently, the experimental results we report for dynamic permanent buyout prices in the next section are not quite as conclusive as for dynamic temporary buyout prices.

5.3. Numerical Experiments. In this subsection we compare in different market environments the utility derived by the seller with a dynamic buyout price auction, a static buyout price auction, and an auction with no buyout price; the results presented are representative of a much larger set of experiments than those reported here.

Let $\mathbf{E}[U^S_{tmp}(p^*_{tmp})]$, $\mathbf{E}[U^S_{prm}(p^*_{prm})]$ and $\mathbf{E}[U^S_{nb}]$ be as defined in §4.4. As before, estimates of those terms are obtained by simulation, and are within 1% of the true values with 95% confidence. For both temporary and permanent options, we also consider the special case of the seller's revenue optimization problem whereby maximization is restricted to the set of fixed threshold valuation functions, i.e. $\nu(t) = v$ for all t. The optimal fixed temporary (resp. permanent) threshold valuation v^*_{tmp} (resp. v^*_{prm}) can then be computed by solving numerically the single variable concave maximization problem obtained by the above substitution. Slightly abusing notation, we denote by $\mathbf{E}[U^S_{tmp}(v^*_{tmp})]$ and $\mathbf{E}[U^S_{prm}(v^*_{prm})]$ the corresponding expected utility of the seller in a temporary and permanent buyout auction respectively. For the temporary buyout option we also compute the upper bound $\bar{Z}_{tmp}(\tau)$ defined in Proposition 3, where τ is a partition of $[0,T]$ into 500 subintervals of equal length. We report the values of all the terms above relative to the seller's expected discounted revenue from an auction with no buyout price in Table 3.

	α	0.01			0.03		
	λT	4	8	16	4	8	16
Temporary	$\frac{\mathbf{E}[U^S_{tmp}(p^*_{tmp})]-\mathbf{E}[U^S_{nb}]}{\mathbf{E}[U^S_{nb}]}$	2.87%	2.07%	1.49%	12.89%	10.34%	8.79%
	$\frac{\mathbf{E}[U^S_{tmp}(v^*_{tmp})]-\mathbf{E}[U^S_{nb}]}{\mathbf{E}[U^S_{nb}]}$	3.65%	2.02%	1.40%	13.64%	10.31%	8.60%
	$\frac{\bar{Z}_{tmp}(\tau)-\mathbf{E}[U^S_{nb}]}{\mathbf{E}[U^S_{nb}]}$	4.13%	2.87%	3.05%	13.91%	10.78%	10.38%
Permanent	$\frac{\mathbf{E}[U^S_{prm}(p^*_{prm})]-\mathbf{E}[U^S_{nb}]}{\mathbf{E}[U^S_{nb}]}$	6.55%	5.78%	6.88%	23.43%	25.22%	30.08%
	$\frac{\mathbf{E}[U^S_{prm}(v^*_{prm})]-\mathbf{E}[U^S_{nb}]}{\mathbf{E}[U^S_{nb}]}$	7.30%	6.68%	7.49%	24.57%	26.13%	30.47%

Table 3 Utility increase achieved by fixed buyout price and fixed threshold valuation auctions

A first observation from Table 3 is that the relative performances of the optimal static buyout option and the dynamic buyout option with optimal static valuation are within 1.5% of each other in both the temporary and the permanent case for all environments considered. In the temporary case, these two relative performances are furthermore always within 2% of the maximum relative performance achievable by any dynamic buyout price policy, as shown by a comparison with the relative value of the upper bound $\bar{Z}_{tmp}(\tau)$. Our results therefore support the prediction that online auction sellers stand to gain relatively little from using a dynamic temporary buyout price rather than a static one.

In the permanent case, a similar upper bound for the maximum revenue achievable with a dynamic buyout price is unfortunately not available to us, and our results are therefore not as conclusive. However, the fact that the dynamic buyout option with optimal static valuation and optimal static buyout price perform very similarly in both cases and are provably very close to optimal in the temporary case indicates that they may also be close to optimal in the permanent case. In light of the higher implementation complexity and possible negative reactions from bidders faced with the untested concept of a dynamic buyout price, our results thus suggest that one should at least have a pessimistic prior about any potential gain from a dynamic permanent buyout price relative to a static one.

6. Conclusion

We now summarize the answers to the three motivational questions raised in the introduction obtained from the analysis just presented:

Question 1: How should a seller using an online auction set the buyout price (if at all)? Our equilibrium analysis of an auction with a buyout option produces a prediction for the seller's expected discounted revenue resulting from a given value of the buyout price, which we can then use to formulate and analyze an optimization problem where this buyout price is the main decision variable. While this problem is difficult to solve analytically in the general case, for practical purposes its solution may still easily be computed with high precision using simulation. From a qualitative standpoint, this model and our numerical experiments confirm the intuition that the optimal buyout price for the seller increases with the expected number of bidders and the bidders' time-sensitivity, and decreases with the seller's time-sensitivity. Our results also suggest that when facing a given market environment, the value of the permanent buyout price which is optimal for the seller is higher than that for a temporary buyout price. Finally, our asymptotic analysis of the seller's optimization problem yields, in some special cases, closed-form expressions for the optimal buyout price that may be potentially useful to practitioners (the reader will find more such expressions in Gupta (2006)). But it also generates some mechanism design insights for the dynamic market environment we consider that extend those described in Harris and Raviv (1981) for a static market environment. Specifically, in our model where the relative values of the seller's and bidders' time sensitivity and the bidder arrival rate effectively capture market power and the ratio between supply and demand, a time-sensitive seller facing few patient bidders should use a fixed posted price, while a patient seller facing many bidders should bypass the buyout option and only use a regular auction mechanism; the hybrid mechanism and smooth transition enabled by a buyout option is appropriate for a range of market environments between those two extremes.

Question 2: What are the implications of using a temporary buyout option relative to a permanent one? Our equilibrium analysis suggests that with a temporary option the first bidder to submit a regular bid will also do so immediately upon arrival, but with a permanent option all regular bids should be submitted shortly before the end of the auction. Note that our model does not provide any prediction for when regular bids from the second and subsequent bidders will be submitted in an auction with a temporary buyout option. In practice, the timing of bid submissions is also affected in various ways by features not captured here; for example a high cost of monitoring the auction could hasten bid submissions, while common value could delay them. However, our model does suggest that the marginal impact of a permanent buyout option relative to a temporary one is to delay the first bid (presumably a negative for the seller if bidding activity may be attracting more bidders), and concentrate bidding activity near the end of the auction. From that perspective, we find it remarkable that Amazon's online auction site, one of the largest with a permanent buyout option, also features a rule whereby the first bidder is offered a 10% discount on the final selling price should he win the auction. However, this obvious incentive for early bidder involvement is not used on any site with a temporary buyout option we are aware of, most prominently eBay. Another relevant remark is that automatic activity-based bidding period extension rules, which Roth and Ockenfels (2002) show to reduce bidding concentration near the end of the auction, are predominantly featured by auction sites using a permanent buyout option (e.g. Amazon and Yahoo!), and conspicuously absent from the site eBay. Taken together, these observations lend support in our view to the validity of our analysis and robustness of our model predictions. This paper thus sheds some light on, but does not resolve, the issue of which type of buyout option is preferable from a seller's standpoint. A first insight we obtained is that the relative attractiveness for the seller of a temporary buyout option decreases with the expected number of bidders, whereas it increases in the case of a permanent buyout option. Furthermore, the seller's expected discounted revenue derived from an optimal permanent buyout option was larger than that obtained with an optimal temporary option in all the numerical experiments we performed with our optimization model. In practice however, the higher incentives for late bidding associated with the permanent option may negatively impact the seller's revenue for reasons that our model does not capture (e.g. signaling effect of bidding activity). The theoretical results just mentioned thus do not justify in our view an unambiguous recommendation to always use a permanent option over a temporary one, except perhaps for very time-sensitive sellers in environments with a high expected number of bidders, the conditions under which the predicted difference in expected discounted revenue was largest in our experiments. This nuanced interpretation also seems justified by the continued use by eBay (the largest and arguably most successful auction site currently operating) of a temporary buyout option.

Question 3: What is the potential benefit associated with using a dynamic buyout price that may vary as the auction progresses? While our results are not quite as conclusive in the permanent case as in the temporary one, they still suggest that the potential revenue increase enabled by such dynamic buyout price is small, seemingly not justifying the associated implementation complexity and possible negative reactions from bidders; the fact that to the best of our knowledge no dynamic buyout price has ever been used in any actual auction site may also be corroborating these findings.

We mention in closing several possible extensions of this work. Although as stated in §2 we focused in this paper on time-sensitivity as a primary driver for the use of buyout options, it turns out that the structure of the equilibrium strategy derived in §4.1 remains the same when bidders are also assumed risk averse with CARA utility function $U(v,t,\tau) \triangleq 1 - e^{-re^{-\beta(\tau-t)}(v-x)}$ ($r > 0$ is the coefficient of risk aversion, see §3.1 for other notation). That is, Theorem 1 can be extended to show that for a temporary buyout price auction with such bidders, there exists a threshold function $\nu_{tmp}^{(r)}$ such that $\mathcal{T}[\nu_{tmp}^{(r)}]$ defines a Bayesian Nash equilibrium. Furthermore, it follows from Jensen's inequality that $\nu_{tmp}^{(r)} \leq \nu_{tmp}$. Intuitively, the riskless buyout option is more attractive to risk averse bidders relative to regular bidding (which involves both winning and selling price uncertainty), resulting in a lower buyout valuation threshold. The results for the permanent buyout price case stated in §4.2 can be similarly generalized for such bidders. A more complete study of the impact of seller and bidders' risk aversion on the buyout price in this dynamic environment remains to be conducted however. While we conjecture that the qualitative impact of seller and bidder risk-aversion on the optimal buyout price is similar to that of our time-sensitivity discount factors α and β because in this game the riskier outcome (auction, as opposed to buyout) is also more distant in time, we leave this issue aside for future research. Also, while focusing on the seller's perspective seemed justified in this first study because sellers typically choose auction sites and parameters, it would be valuable to explore the impact of buyout options on bidders' utilities. Finally, we would like to extend our analysis to the case of multi-item auctions, and also consider dynamic buyout prices that would not be pre-determined but rather modified according to actual bidding activity during the auction.

Acknowledgments
We are grateful to Paul Zipkin, the Associate Editor and two anonymous referees for their many comments and insightful suggestions. We would like to thank Alexandre Belloni, Pranava Goundan, Raghavendran Sivaraman, Dan Stratila, Georgia Perakis, John Sterman, Yossi Sheffi, James Dana, Michael Rothkopf, Wedad Elmaghraby and the participants of the Operations Management Seminar of the MIT Sloan School of Management for helpful discussions, reference suggestions and feedback. The second author is also grateful to his colleagues James B. Orlin and Stephen C. Graves for their help. This work was partially supported by the Singapore-MIT Alliance, Robert M. Freund, and the J. Spencer Standish Career Development Chair of the MIT Sloan School of Management.

References

Budish, E.B., L.N. Takeyama. 2001. Buy prices in online auctions: irrationality on the internet? *Economic Letters* **72** 325–333.

Caldentey, R., G. Vulcano. 2004. Online Auction and List Price Revenue Management. *Stern School of Business, New York University* .

eBay. 2006. Learn about Buy It Now. http://pages.ebay.com/services/buyandsell/buyitnow.html.

Fudenberg, D., J. Tirole. 1991. *Game Theory*. The MIT Press.

Gallien, J. 2006. Dynamic Mechanism Design for Online Commerce. *Operations Research* **54**(2) 291–310.

Gupta, S. 2006. MIT Operations Research Center. Ph.D. thesis, in preparation.

Harris, M., A. Raviv. 1981. A Theory of Monopoly Pricing Schemes with Demand Uncertainty. *American Economic Review* **71**(3) 347–65.

Hidvégi, Z., W. Wang, A.B. Whinston. 2003. Buy-Price English Auction. *Center of Research on Electronic Commerce, University of Texas at Austin* .

Kirkegaard, R., P.B. Overgaard. 2003. Buy-Out Prices in Online Auctions: Multi-Unit Demand. *Department of Economics, University of Aarhus* .

Klemperer, P. 1999. Auction Theory: A Guide to the Literature. *J. of Economic Surveys* **13**(3) 227–286.

Lucking-Reiley, D. 2000. Auctions on the Internet: What's Being Auctioned, and How? *Journal of Industrial Economics* **48**(3) 227–252.

Mathews, T. 2003a. Bidder Welfare in an Auction with a Buyout Option. *Department of Economics, California State University-Northridge* .

Mathews, T. 2003b. A Risk Averse Seller in a Continuous Time Auction with a Buyout Option. *Brazilian Journal of Economics* **5**(2).

Mathews, T. 2004. The Impact of Discounting on an Auction with a Buyout Option: a Theoretical Analysis Motivated by eBay's Buy-It-Now Feature. *Journal of Economics* **81**(1) 25–52.

McAfee, R.Preston, John McMillan. 1987. Auctions and Bidding. *Journal of Economic Literature* **25**(2) 699–738.

Reynolds, S.S., J. Wooders. 2003. Auctions with a Buy Price. *Department of Economics, Eller College of Business & Public Administration* .

Roth, A., A. Ockenfels. 2002. Last-Minute Bidding and the Rules for Ending Second-Price Auctions: Evidence from eBay and Amazon on the Internet. *American Economic Review* **92**(4) 1093–1103.

Vakrat, Y., A. Seidmann. 2001. Implications of the bidders' arrival process on the design of online auctions. *Proceeding of the 33rd Hawaii International Conference on System Sciences* .

van Damme, E. 1987. *Stability and Perfection of Nash Equilibria*. Springer-Verlag.

Wan, W., H. Teo, L. Zhu. 2003. Determinants of Buy-It-Now Auction Choice on eBay. *7th Pacific-Asia Conference on Information Systems* .