Ads in Conversations: Market Thickness and Match Quality

Martino Banchio Aranyak Mehta Andres Perlroth
Google Research

Abstract

We study the optimal placement of advertisements for interactive platforms like conversational AI assistants. A key platform decision is when to show an ad during a conversation. On the one hand, by delaying ad delivery platforms can refine ad quality estimates and improve ad selection. However, delay endogenously reduces market thickness, and in turn revenue. We analyze first- and second-price auctions where the platform can commit to the format but not to its timing. First-price auctions achieve efficient allocation but they delay ad delivery. In equilibrium, advertisers forecast the delay and bid-shade aggressively, reducing the platform's revenue. Second-price auctions, on the other hand, avoid delay because their revenue relies on market thickness. The platform prefers a prior-free second-price auction to a first-price auction. Instead, with optimal reserves she prefers a first-price to a second-price auction.

1 Introduction

Conversational AI platforms, like chatbots and virtual assistants, are rapidly gaining popularity. These platforms offer users a convenient and interactive way to access information, complete tasks, and engage in entertainment. From booking flights and ordering food to playing games and getting personalized recommendations, conversational AI platforms have the potential to transform the way people interact with technology. Many of these platforms are currently offered free of charge, but as their user base grows and their capabilities expand, it is natural to envision that, among a number of other strategies,[1] advertising could play a significant role in their monetization.

[1]There are many alternative avenues for monetization of Conversational AI assistants, e.g., subscriptions, etc.

This presents a unique opportunity for targeted advertising. Unlike traditional online advertising, conversational AI platforms can glean real-time insights into user preferences through the natural flow of conversation. For example, a user interacting with a chatbot to book a hotel might reveal their budget, desired location, travel dates, and other preferences through their questions and responses. This dynamic information can be used to refine estimates of ad quality, measured in this paper by the probability of a user clicking on an ad (click-through rate).

We study how to leverage this dynamic information in advertising auctions. A platform commits to an auction format but she chooses when to run the auction and display one ad only after receiving the advertisers' bids. Her decision is then best understood as a real options problem. She can delay the ad selection to acquire more precise quality estimates, but this may reduce market thickness, potentially leading to lower revenue due to decreased competition.

We develop a theoretical model to analyze this tension between information acquisition and market thickness. A platform (the auctioneer) wants to sell a single ad to one of two advertisers (buyers) over the course of a conversation. Each advertiser's value of showing the ad is the product of two components. First, he values a click at v_i, and a user clicks on the ad with probability θ_i. We model θ_i as a binary variable, where $\theta_i = 1$ indicates a good match (high click-through rate) and $\theta_i = 0$ indicates a bad match (low click-through rate). Initially, the quality score is unknown to everyone, including the advertisers. As the user interacts with the conversational AI, the platform gradually learns whether each ad is a good or a bad match.

We assume that the platform can commit to the auction format but not to its timing. This partial commitment assumption reflects the limited monitoring available to advertisers in online advertising settings.[2] We thus begin by characterizing the optimal exercise of the platform's *option*: that is, the auction timing chosen by the auctioneer. We show that in equilibrium the platform delays a first-price auction as much as possible, and she expedites a second-price auction instead. To understand this result and its effect on the platform's objective, it is instructive to think about the drivers of revenue in the two auction formats.

First-price auctions realize revenue by selecting the efficient advertisement. The platform delays the allocation until it has precise information about the quality scores, ensuring she will display one of the highest click-through rate ads, thus maximizing her chances of generating revenue. However, this efficiency comes at the cost of aggressive

[2]While it may be possible to ex-post verify the nature of the auction format, verifying that the platform kept its timing promises seems implausible.

bid shading by the advertisers. Anticipating the delayed allocation and the resulting reduced competition, advertisers strategically lower their bids, reducing revenue for the platform. In contrast, second-price auctions rely on market thickness to generate revenue. The platform allocates the ad earlier, when information about quality scores is still imperfect. This leads to a thicker market and stronger competition, which drives up the price paid by the winning advertiser. Naturally this comes at the cost of potential inefficiency, as the platform may allocate the ad to a lower-quality advertiser due to incomplete information. Nonetheless, we show that the platform has an incentive to expedite the auction, avoiding the risk of a thin market but allocating inefficiently.

Intuitively, the divergence in equilibrium timing relies on the distinct ways each auction format leverages market thickness. In a second-price auction, the price paid by the winner is determined by the second-highest bid. Delaying the auction increases the risk of losing bidders and thinning the market, potentially leading to a drastically lower price. The platform therefore chooses to allocate early to capitalize on the thicker market, even if it means sacrificing some efficiency. Instead, in a first-price auction the winner pays their own bid regardless of the competition. The auctioneer can focus on allocating efficiently by acquiring more precise quality information. This strategy incentivizes bid shading by advertisers, as they anticipate facing less competition when the auction finally occurs.

When the distribution of advertiser's valuations are common knowledge, the auctioneer can generally reduce the impact of market thickness on revenue by carefully selecting reserve prices. We show that, because advertisers in a first-price auction no longer benefit from a thin market, the equilibrium bid-shading contracts, allowing the auctioneer to implement the optimal auction. Instead, in a second-price auction advertisers may still benefit from a thin market, even when the auctioneer selects the optimal reserve price. The auctioneer benefits from market thickness when multiple advertisers submit large bids, thus she has an incentive to expedite the auction exercise. The advertisers, on the contrary, have an incentive to delay the exercise, thus in equilibrium they will misreport to induce delay. The second-price auction will never generate as much revenue as the first-price auction.

We are motivated by the monetization problem of Generative AI providers, and the model tightly reflects our fundamental question. However, it turns out that the tradeoff between competition and quality scores appears in a number of interesting economic settings. We propose two examples below of auctions with *dynamic* scoring, where the winner selection procedure relies on factors other than the buyers' values, and such factors are dynamically updated.

Mergers & Acquisitions. Multiple firms are interested in purchasing a small business. In most competitive M&A processes, buyers submit bids to acquire the business with some contingencies that will be resolved during the due diligence period.[3] Such contingencies may result in contract termination. While due diligence is costly, a natural question is whether bidders should be encouraged to perform due diligence *before or after* the winner selection. Selecting a winner before due diligence makes the process more competitive, but exposes the business to the risk of a failed M&A deal. Viceversa, allowing due diligence to take place before selecting the winner may thin out the pool of interested parties, thereby reducing competitive pressure.

Public Procurement. Government agencies procure goods and services from suppliers in an auction. Regulators require that suppliers pass certain probity measures – for example, money-laundering and criminal record checks. Should the agencies front-load or back-load compliance verification when procuring from industry suppliers? That is, should the agencies run their auction and then audit the winner, or should they audit all participants and only after choose the winner? EU directives only require a bidder self-declaration at auction time, and then demand full compliance checks to be performed on the winner.[4] Again, by performing compliance checks after the auction, the agencies take advantage of market thickness to procure at low prices, at the risk of unfulfilled contract. Instead, running the auction after thorough checks would ensure contract fulfillment but would possibly increase prices.

The paper is structured as follows. In Section 2 we formally describe the model. Section 3 presents the main results we described in this introduction: first, we compare first- and second-price auction without reserve prices, then with reserves. We study two extensions in Appendix A, and then we discuss our assumptions in Section 4.

Our work contributes to the growing literature on online advertising auctions by explicitly incorporating the dynamic nature of information acquisition in conversational AI settings. We highlight the distinct ways in which first- and second-price auctions balance the trade-off between information and market thickness, offering valuable insights for platforms seeking to design revenue-maximizing ad auctions in this emerging space.

[3] See Marquardt and Zur (2015) and Wangerin (2019).

[4] Directive 2014/24/EU, art. 57-61.

1.1 Literature Review

We connect with a large literature on online advertising auctions starting with Edelman et al. (2007) and Varian (2009). While early papers study the static problem (Athey and Nekipelov (2010), Börgers et al. (2013)), the literature has moved on to the repeated auction problem (Balseiro et al. (2015), Chen (2017)). We abstract away from cross-auction incentives and focus instead on a single dynamic auction. The key force we study is intrinsically dynamic, and dynamics are incorporated in the auction mechanism.

A large literature studies dynamic auctions; for a non-exhaustive list of reference see Bergemann and Said (2010). The literature mainly focuses on two sources of dynamics. The first source of dynamics is an evolving population of agents, each having fixed private information. The second source of dynamics is instead a fixed population of agents whose private information evolves dynamically. In our paper we have a fixed population of agents, but private information is also fixed. The only source of dynamics is in the ad quality scores, which the auctioneer exploits to optimize her revenue. Related to our paper is Chaves and Ichihashi (2024), which also considers an auctioneer who chooses the timing to run the auction optimizing for market thickness. In our setting delay reduces market thickness, whereas there it is to increase market thickness that the auctioneer delays the auction. Importantly, instead of focusing on finding the optimal delay, we compare equilibria of different auction formats.

There is a large literature on real options in finance and economics that deals with the revenue maximizer's optimal stopping problem; there are many reviews of this literature, such as Sick (1995). In our setting, solving the optimal stopping problem is instrumental but not sufficient to characterize the auctioneer's problem, because advertisers optimally respond to the exercise timing. The equilibrium is a fixed-point problem of advertisers and auctioneer's decisions. In this sense, we connect to the literature on game-theoretic real option problems, such as those described in Grenadier (2000). Methodologically we also adopt techniques from the continuous-time literature (see Hörner and Skrzypacz (2017) for some excellent examples), and in particular we mainly adopt the "bad news" Poisson model studied in Keller and Rady (2015) for some of our extensions.

Finally, our motivation connects us to a nascent literature on Large Language Models (LLM), chatbots, and game theoretic models. As we mentioned, our problem can be thought of as the advertising decision of a LLM-based chatbot provider. Some recent papers in the literature have started thinking about the effect of LLMs on auctions and auction design, for example Duetting et al. (2023) which studies how mechanism design should aggregate several LLM-generated input in an incentive compatible way for online advertising. Feizi et al. (2023) proposes a setup for online ads on LLMs. They mention

that the system for predicting quality scores can "update the estimate in (almost) real time, which will increase the accuracy of the prediction". We consider the effect of the prediction dynamics on the auctioneer's revenue. An iterative refinement of beliefs similar to the one we model here appears in Harris et al. (2023).

2 Model

Consider a user of a content platform with a private type $\theta = (\theta_1, \theta_2)$, where each component θ_i is drawn independently from a Bernoulli distribution with parameter p. The type θ_i represents the user's interest in the ad of advertiser $i = 1, 2$.[5] Specifically, the user clicks on ad i if and only if $\theta_i = 1$; in this case we say that ad i is a *good quality ad*. Each advertiser derives a value $v_i \sim F(v)$ from a successful click, where $F(v)$ is a regular distribution over $[0, \overline{v}]$.

An auctioneer, who controls the platform, determines which advertiser's ad is displayed. She receives a payment from the winning advertiser if and only if the user clicks on the ad, and she can choose whether to allocate the ad according to a first- or second-price auction.

The user interacts with the platform in a *conversation* which unfolds over a continuous time $t \in \mathbb{R}_+$. During the conversation, the auctioneer gradually gleans insights about the quality of the ads. While good ads ($\theta_i = 1$) offer no direct signal of the user's interest, bad ads ($\theta_i = 0$) progressively reveal themselves through the user's interaction with the platform. Formally, we model the arrival of news for ad i as an exponential clock τ_i with rate $\lambda \cdot (1 - \theta_i)$. When the clock of ad i ticks (i.e. at time τ_i), the auctioneer learns that such ad is of bad quality. Absence of news is good news — as time passes, an ad becomes more likely to be of good quality if its clock has not ticked yet.[6] Formally, the auctioneer forms beliefs $\mu_t = (\mu_t^1, \mu_t^2)$ about the value of the ad qualities θ. If the exponential clock of ad i ticks, the belief jumps to $\mu_t^i = 0$. Instead, when no news arrive, the belief drifts upward according to the following law of motion:

$$\begin{cases} \dot{\mu}_t^i = \lambda \mu_t^i (1 - \mu_t^i) \\ \mu_0^i = p \end{cases} \tag{1}$$

where the prior p is the parameter of θ_i's Bernoulli distribution.

Finally, the user may interrupt his conversation on the platform at any time. This

[5] In Section 4 we consider a model with more than 2 advertisers.

[6] In the terminology of Keller and Rady (2015) this is a model of "breakdowns".

means that the advertising opportunity may disappear, making learning about qualities explicitly costly. We assume that the agent leaves at an instantaneous rate $r \geq 0$, so that conditional on the user not having left yet he will leave at an exponentially distributed time with rate r.[7]

2.1 The Auctioneer's Problem

The auctioneer strategically chooses the ad to display according to either a first- or second-price auction without reserves, and then she decides at what time $\tau \in \overline{\mathbb{R}} := \mathbb{R}_+ \cup \{+\infty\}$ to run such auction. The timing of the allocation problem is as follows.

- First, before the conversation begins, the auctioneer announces an auction format. She commits to running either a first- or second-price auction, but she cannot commit to a time τ at which such format will be run.
- Then, the auctioneer solicits bids b_1, b_2 from the bidders (advertisers). All of the advertisers' strategic bidding decisions take place before the conversation begins.
- The conversation starts at time 0, and the auctioneer begins learning about the user's type according to the beliefs in Equation (1).
- At some time $\tau \in \overline{\mathbb{R}}$, which may depends on the advertisers' bids and the realizations of the uncertainty the auctioneer decides to "stop", at which point she runs the auction format she committed to.

If the winner is advertiser i, under a first-price auction she collects revenue $\pi_{1P} = e^{-r\tau} b_i \theta_i$, while in a second-price auction she collects $\pi_{2P} = e^{-r\tau} \frac{b_{-i}\mu_\tau^{-i}}{\mu_\tau^i} \theta_i$.

We make two main assumptions with respect to timing. First, advertisers cannot adjust their bids dynamically, but instead make all their strategic considerations at time 0. Latency considerations alone justify such an assumption. Additionally, it seems unrealistic for a platform to require that advertisers monitor the news arrival process and update their bids based on the arrival or absence of news. Even more simply, allowing advertisers to dynamically adjust their bids based on calendar time would require a complex infrastructure on both the platform and the advertiser's end. Second, we assume that the auctioneer cannot commit to a time to run the auction.[8] Then, once she receives the bids of the advertisers, the auctioneer faces a real options problem. She can hold off on

[7] An alternative interpretation of the rate of departure of the agent r is that both the advertisers and the auctioneer discount the future at rate r. Most of the paper will focus on the case where $r = 0$, but we present generalizations to the $r > 0$ case in the extensions.

[8] We will relax this assumption in Appendix A.2.

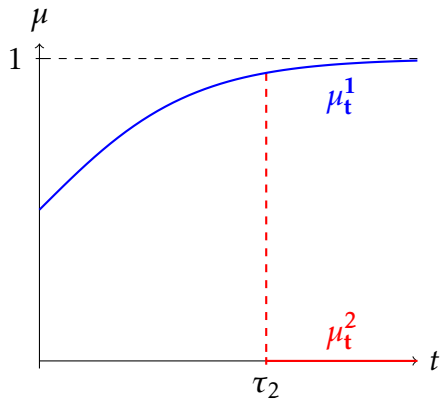


Figure 1: This figure depicts a sample path taken by the belief process μ_t. The beliefs drift upwards in the absence of news (as is the case for the process μ_t^1). Instead, news about advertiser 2 arrived at time τ_2, and thus $\mu_t^2 = 0$ from τ_2 onwards.

running the auction to learn more about the ad qualities, because she profits only when she displays a good quality ad. On the other hand, information is costly. For one, the user may leave the platform at any point in time, making the profit opportunity vanish. More subtly, delaying the auction erodes the competitive landscape. Auctions rely on competition to price the ad opportunities, but delay reduces market thickness and thus greatly diminishes competition in the auction. We show that the auction format plays an important role in the resolution of this trade-off between information and market thickness.

3 Auction formats, market thickness, and information

In order to study the auctioneer's problem, we need to characterize the optimal "exercise policy", i.e. the optimal time in the conversation to stop and run the auction from the auctioneer's perspective. To this end we solve a model in which the user never leaves the platform. That is, we assume $r = 0$. We shut down the direct costs of information acquisition, leaving only the indirect cost of a market that becomes thinner the more information is acquired.

3.1 Prior-free auctions

Suppose an auctioneer is facing bidders whose types come from unknown distributions. Instead of guessing the reserve price, she may simply run an auction without reserves. In this subsection we thus focus our attention on first- and second-price auction without reserve prices.

Consider an auctioneer who committed to running a second-price auction. She observes some reports b_1, b_2 and then must decide when to stop to maximize her expected revenue. Her problem can be written as

$$\max_{\tau} \mathbb{E}_0\Big[\min\big\{\mu_\tau^1 b_1, \mu_\tau^2 b_2\big\}\Big]$$

We are ready to prove our first result.

Lemma 1. The optimal exercise policy for the second-price auction is $\tau_{2P}^* = 0$.

Proof. Following the usual Bayesian arguments, notice that the belief process μ_t^i is a martingale: $\mathbb{E}_t[\mu_{t'}] = \mu_t$ for any $t' \geq t$. Then, by definition

$$\mathbb{E}_t\Big[\min\big\{\mu_{t'}^1 b_1, \mu_{t'}^2 b_2\big\}\Big] \leq \mathbb{E}_t\Big[\mu_{t'}^i b_i\Big] = \mu_t^i b_i$$

for all $i = 1, 2$ and for all $t' \geq t$. This implies that

$$\mathbb{E}_t\Big[\min\big\{\mu_{t'}^1 b_1, \mu_{t'}^2 b_2\big\}\Big] \leq \min\big\{\mu_t^1 b_1, \mu_t^2 b_2\big\}.$$

Thus, the revenue process is a super-martingale. Doob's Optional Stopping Theorem tells us that a super-martingale is optimally stopped at 0. □

Lemma 1 says that the auctioneer stops at time 0 *uniformly* over the reports of the advertisers. In particularly, she will not use time as a screening device, and the timing of the auction will not affect the incentives of the advertisers, who then bid truthfully. The second-price auction is thus a truthful auction in this setting, and it is executed at time 0.

The auctioneer is willing to allocate to the bidder with the highest value despite having only receiving payment from him with probability μ_0. Recall that the auctioneer will profit only if the ad she displays is of good quality. Here the auctioneer profits only with probability μ_0, while by delaying the auction she could increase her chances of profiting from the displayed ad. However, the probability that any given advertiser drops out of the race ($\mu_t^i = 0$) increases with delay. The second-price auction relies on competition to price the ad opportunity, but as soon as the belief about one advertiser drops to 0 competition vanishes, leaving a perfectly good opportunity to be sold for free to the remaining advertiser. The auctioneer runs the auction immediately to avoid such a catastrophic outcome.

Intuitively, it seems that the second-price auction suffers from lack of competition more than the first-price auction. That is because even if one advertiser drops out, in

the first-price auction the other advertiser will purchase the ad opportunity at the bid he chose at time 0, before realizing his competition disappeared. We show next that the first-price auction delays the exercise of the auction until the auctioneer is sure to allocate efficiently. Consider then an auctioneer who committed to running a first-price auction. She observes bids b_1, b_2 and then decides when to stop to maximize her expected revenue, written as

$$\max_{\tau} \mathbb{E}_0\Big[\max\left\{\mu_\tau^1 b_1, \mu_\tau^2 b_2\right\}\Big]$$

Using this expression, we can characterize the optimal exercise policy for the first-price auction.

Lemma 2. The optimal exercise policy for the first-price auction is $\tau_{1P}^* = \infty$.

Proof. By definition

$$\mathbb{E}_t\Big[\max\left\{\mu_{t'}^1 b_1, \mu_{t'}^2 b_2\right\}\Big] \geq \mathbb{E}_t\Big[\mu_{t'}^i b_i\Big] = \mu_t^i b_i$$

for all $i = 1, 2$ and for all $t' \geq t$. This implies that

$$\mathbb{E}_t\Big[\max\left\{\mu_{t'}^1 b_1, \mu_{t'}^2 b_2\right\}\Big] \geq \max\left\{\mu_t^1 b_1, \mu_t^2 b_2\right\}.$$

Thus the revenue process of the first-price auction is a sub-martingale until the clock of an advertiser ticks, if ever. A sub-martingale's expected value increases over time, so the auctioneer will never stop before one of the advertisers' clock ticks. If it does, the revenue process turns into the (bounded) martingale $\mu_t^i \cdot b_i$ for some i, and the auctioneer is indifferent between stopping and continuing at any time t. We say the auctioneer "stops at ∞", where the value of $\mu_\tau^i b_i$ at $\tau = \infty$ is given by the pointwise limit of the process $\mu_t^i b_i$ which exists almost surely because $\mu_t^i b_i$ is bounded. □

The first-price auction is efficient: the advertiser that displays his ad is always the one with the largest $b_i \cdot \theta_i$. This auction is not truthful however: the advertisers at time 0 foresee that, conditional of having a good quality ad, they will have no competition with probability $1 - \mu_0$. Only with probability μ_0 they will need to outbid their opponent, and their shading will reflect the probability of this event.[9]

How does the bid shading affect the revenue to the auctioneer? The following proposition responds to this question.

Theorem 1. *The revenue π_{2P}^* generated by the truthful optimally-stopped second-price auction dominates the revenue π_{1P}^* generated by an optimally-stopped first-price auction for any value*

[9] We will show that there is a unique symmetric equilibrium in Proposition 4.

distribution F. In particular,

$$\frac{\pi^*_{2P}}{\pi^*_{1P}} = \frac{1}{\mu_0}$$

Proof. To prove the first part of the theorem, notice that the second-price auction is truthful regardless of its stopping time, as long as the stopping time does not depend on the values reported by the advertisers. In particular, consider a second-price auction stopped at τ^*_{1P}. This stopping time depends solely on the path of the beliefs μ_t, but cannot be influenced by the reports v_1, v_2, as shown in Lemma 2. Then, the static Revenue Equivalence Theorem (RET) applies to auctions that run at time τ^*_{1P}: the first- and second-price auction that stop at τ^*_{1P} allocate identically, and therefore generate the same revenue. But the auctioneer receiving truthful reports prefers to stop at time $\tau^*_{2P} < \tau^*_{1P}$, which implies that her revenue must be (weakly) larger by stopping earlier than by delaying the auction until τ^*_{1P}. Therefore, the second-price auction generates weakly higher revenue than the first-price auction.

To prove the second part of the theorem, we leverage the virtual value characterization of the auctioneer's revenue. The expected revenue of an optimally-exercised first-price auction is given by $\mu_0^2 \mathbb{E}_v\Big[\max\{\phi(v_1), \phi(v_2)\}\Big] + 2\mu_0(1-\mu_0)\mathbb{E}_v\Big[\phi(v)\Big]$ where $\phi(v)$ is the virtual value function. Intuitively, the first-price auction allocates efficiently: with ex-ante probability $2\mu_0(1-\mu_0)$ there is only one ad which is a good match, and with probability μ_0^2 the auctioneer can choose the good match with the largest virtual value. The usual integration by part argument shows that $\mathbb{E}_v[\phi(v)] = 0$, thus implying that $\pi^*_{1P} = \mu_0^2 \mathbb{E}_v\Big[\max\{\phi(v_1), \phi(v_2)\}\Big]$. Similarly, the expected revenue of the optimally-exercised second-price auction is $\pi^*_{2P} = \mu_0 \mathbb{E}_v\Big[\max\{\phi(v_1), \phi(v_2)\}\Big]$. Intuitively, the second-price auction allocates immediately, with probability 1 to the bidder with the largest virtual value. She will collect her revenue only when that bidder is a good match, which happens with probability μ_0. Therefore,

$$\frac{\pi^*_{2P}}{\pi^*_{1P}} = \frac{\mu_0 \mathbb{E}_v\Big[\max\{\phi(v_1), \phi(v_2)\}\Big]}{\mu_0^2 \mathbb{E}_v\Big[\max\{\phi(v_1), \phi(v_2)\}\Big]} = \frac{1}{\mu_0}$$

□

The bid shading induced by the first-price auction is detrimental to revenue, so much so that the auctioneer would rather implement a second-price auction and allocate inefficiently. In fact, regardless of the value distribution, the revenue obtained by a first-price auction are exactly a fraction μ_0 of the revenue obtained by a second-price auction . Thus the benefit from adopting an optimally-exercised second-price auction increases as the

probability of encountering a good-quality ad decreases.

The case in which the user does not leave a conversation leads to a crisp, stark result. When information would otherwise be costless, the only disincentive to information acquisition is market thickness. Different auction formats leverage market thickness differently: a second-price auction requires competition to price the good for sale. A first-price auction instead requires competition to limit the extent of the bidders' equilibrium shading. Because shading takes place ex-ante, this limits the ability of prices to react to a changing competitive landscape. The second-price auction retains this flexibility, so that the auctioneer can commit in a sequentially rational way not to learn anything about the competitive landscape. These intuitions carry over to the world in which information is costly, i.e. when $r > 0$, as we show in Appendix A.1.

3.2 Auctions with prior information

Our previous findings illustrate that the auctioneer's timing decision heavily depends on market thickness. Both first- and second-price auctions leverage market thickness for revenue generation, albeit through distinct methods. If the auctioneer knew the distribution F of the bidder's values, she may carefully choose a reserve price to try and limit the impact of market thickness on her revenue. In this section, we'll examine the design of first- and second-price auctions with reserve prices, again under the assumption of costless information (r=0). With a reserve price, the auctioneer mitigates the concern of a thinning market, relying on this price floor to maintain revenue.

First, the reserve price does not affect the information acquisition strategy of the auctioneer in a first-price auction.

Lemma 3. The optimal exercise policy for the first-price auction with reserve price R is $\tau^*_{1P,R} = \infty$.

Proof. The auctioneer's expected revenue when $b_1, b_2 > R$ can again be written as

$$\max_{\tau} \mathbb{E}_0\Big[\max\left\{\mu^1_\tau b_1, \mu^2_\tau b_2\right\}\Big],$$

and we know from Lemma 2 that the stopping time takes the form $\tau^*_{1P} = \infty$ Instead, if $b_i > R > b_{-i}$, the auctioneer's revenue at time t is the martingale $\mathbb{E}_t[\theta_i b_i]$, so she is indifferent between stopping and continuing at all times — in particular, we can assume she will stop at infinity with a similar argument to Lemma 2. Finally, if $b_1, b_2 < R$ the auctioneer's revenue is 0 regardless of her stopping decision, so she is again indifferent between all stopping times, and we assume she stops at infinity. □

Bidders can no longer shade as aggressively because the reserve price acts as a backstop. The first-price auction is no longer efficient, because of the introduction of a reserve price, but it maximizes the revenue for the auctioneer. To see this, notice that the direct revelation mechanism that maximizes revenue in this context is one such that

$$x(v_1, v_2|\theta_1, \theta_2) \in \operatorname*{argmax}_x \Big\{x\theta_1\phi(v_1), x\theta_2\phi(v_2), 0\Big\}$$

That is, the mechanism maximizes revenue if it allocates to the bidder with the highest *quality-weighted* virtual value, provided it is positive. Quality-weighted virtual values can be negative only when the virtual values are negative, thus the auctioneer can implement such an allocation with an auction with a reserve price r such that $\phi(r) = 0$. In particular, the allocation function of the first-price auction when there exists a symmetric equilibrium in increasing bidding functions corresponds to $x(v_1, v_2)$, and thus the first-price auction maximizes the auctioneer's revenue.

Perhaps surprisingly, in general the second-price auction instead cannot implement the optimal auction. This is because an auctioneer running a second-price auction cannot commit to stopping at infinity. Instead, the auctioneer sometime will have an incentive to stopping early, to capitalize on competition.

Lemma 4. Suppose that the auctioneer is running a second-price auction with reserve price R. If $b_1, b_2 > 2R$, then $\tau^*_{2P,r}(b_1, b_2) = 0$.

Proof. Let without loss bidders submit bids $b_1 > b_2 > 2R$. Then the auctioneer's expected revenue from stopping at time t (provided that no clock has ticked so far) is $\mu_t b_2$. Her value function at belief μ_t satisfies

$$V(\mu) = \max\Big\{\mu_t b_2, V(\mu_{t+\Delta})\Big\}$$

for some small $\Delta > 0$. Suppose by contradiction that the auctioneer decided to continue at belief μ_t. Then her value function would satisfy[10]

$$V'(\mu)\mu = 2\Big(V(\mu) - \mu R\Big).$$

If the auctioneer decided to stop at a belief $\overline{\mu} > \mu_t$, by continuous pasting at such a belief it must be that

$$\overline{\mu} b_2 = 2\overline{\mu} b_2 - 2\overline{\mu} R \iff b_2 = 2R.$$

[10] See Appendix B for a full derivation.

When $b_2 > 2R$, there is no such belief $\overline{\mu}$. Thus, either the value function $V(\mu)$ is larger than μb_2 for all μ, or $V(\mu) = \mu b_2$ and the auctioneer stops immediately. We know that $V(1) = b_2$, and therefore by continuity $\lim_{\mu \to 1} V(\mu) = b_2$. This holds when $V(\mu) = \mu b_2$ but cannot possibly hold when $V(\mu)$ solves the ODE $V'(\mu)\mu = 2(V(\mu) - \mu R)$: if that was the case, then

$$V'(\mu) = 2(b_2 - R) \geq 2b_2 - b_2 = b_2$$

which by continuity implies that $V(\mu) < \mu b_2$ for μ sufficiently close to 1. This is impossible, and concludes the proof. □

The auctioneer's incentive to stop early improved her revenue with respect to a first-price auction without reserve prices; instead, with reserves the first-price auction dominates the second-price auction.

Theorem 2. *The optimal mechanism can be implemented as a first-price auction with reserves. Instead, there exist distributions F such that no second-price auction with reserve implements the optimal mechanism.*

Proof. We have already shown that the first-price auction with reserve R such that $\phi(R) = 0$ implements the optimal mechanism. We are left to prove that the second-price auction cannot possibly implement the optimal mechanism for some distribution F. To do this, take any distribution F over the support $[0, \overline{v}]$ such that $\overline{v} > 2R$ where $\phi(R) = 0$, and fix $\varepsilon < \frac{\overline{v} - 2R}{2}$. We can restrict attention to the second-price auction with reserve R, since any other second-price auction cannot implement the optimal mechanism.[11]

First, note that the second-price auction with a reserve R is not truthful. To see this, consider the incentives of the first advertiser with type $2R + \varepsilon$ if the other advertiser bids truthfully. The auctioneer will allocate at infinity if the other advertiser's bid is lower than $2R$, in which case bidding truthfully is optimal. However, if the other advertiser's bid is larger than $2R$, the auctioneer stops immediately, at time 0. Then, the advertiser has a profitable deviation from truth-telling: she can willfully shade her report as to induce the auctioneer to wait and allocate the item at infinity, once she fully learns the quality scores. This is a profitable deviation because with probability $1 - p$ the competitor's quality score will be $\theta_2 = 0$, and the advertiser wins the item, pays R, and collects an expected payoff $(1 - p)(2R + \varepsilon - R) = (1 - p)(R + \varepsilon)$ which is larger than the payoff she could have obtained by instead reporting truthfully and getting allocated the object with probability $F(2R + \varepsilon) - F(2R)$ and paying $\mathbb{E}[v | v \in (2R, 2R + \varepsilon)] > 2R$.

[11] To see this, note that the reserve cannot be lower because types R may then get an allocation, and cannot be higher because the auctioneer must allocate to types immediately above R.

Naturally there is no mixed strategy equilibrium that implements the optimal mechanism, so we focus our attention on symmetric equilibria in increasing bidding functions $\beta(-)$. If such an equilibrium implemented the optimal mechanism, it may not stop before fully learning at least one of the quality scores. This can only happen if all bids are uniformly bounded above by $2R$. But then, an advertiser with type $\overline{v} - \varepsilon$ has an incentive to bid $x > 2R$ instead of $\beta(\overline{v} - \varepsilon) \leq 2R$: the auctioneer will run the auction at infinity, but the advertiser will now win the item even when his opponent of type $\overline{v}$ has quality score $\theta = 1$. He will pay $\beta(\overline{v}) < 2R$ and make a profit of $\overline{v} - \varepsilon - \beta(\overline{v}) > 0$. □

4 Discussion

In this section we extend the insight we derived in the previous section to more general models.

4.1 The Information Acquisition Process

All the analysis so far focused on the exponential "bad news" model, in part because it delivers tractability to this complex setting, and in part because it represents quite a good model of dynamic quality scores for online advertising. A platform is unlikely to discover whether an ad will receive a click *before* showing the ad, but it may instead become obvious that the ad will not receive a click, maybe because the user's queries steer the conversation away from the advertiser's products. Fortunately, some of our results hold for a general information acquisition process.

Let the news learning process be a (possibly multidimensional) stochastic process $(X_t)_{t=0}^{\infty}$ with law ν^{θ} which depends on the user's type θ.[12]

The auctioneer's information is given by the natural filtration of X, denoted by $\mathbb{F} = (\mathcal{F}_t)_{t \in \mathbb{R}_+}$. She makes Bayesian inference about the value θ. That is, the auctioneer's belief at time t about the quality of advertiser i's ad is $\mu_t^i = \mathbb{E}_t\left[\theta_i\right] := \mathbb{E}\left[\theta_i | \mathcal{F}_t\right]$. We drop the script i when we refer to vectors. For example, the vector of beliefs is denoted by $\mu_t = (\mu_t^1, \mu_t^2)$. The Bayesian prior $\mu_0^i = \mathbb{E}_0[\theta_i]$. We require that a long enough conversation will identify the qualities θ_i.

Definition 1. A belief μ_t is *consistent* if $P\Big(\lim_{t\to\infty} \mu_t = \theta \mid \theta\Big) = 1$ almost surely in θ.

[12] Formally, let Ω be a Polish metric space and $\mathcal{F}$ a σ-algebra over Ω. For each $\theta \in \Theta$ let P^{θ} be a probability measure over $(\Omega, \mathcal{F})$. The collection $X = \left(X_t \colon \Omega \to \mathbb{R}\right)_{t \in \mathbb{R}_+}$ is such that every X_t is measurable with respect to the Borel σ-algebra on $\mathbb{R}$. We denote by $\nu_t^{\theta} = P^{\theta} \circ X_t^{-1}$ the law of the random variable X_t for a given parameter θ.

Definition 1 says that the beliefs should converge in probability to the true quality θ. Assumption 1 ensures that consistency is guaranteed.

Assumption 1. Assume

1. For any $t > 0$ such that $\mathcal{F}_t \subset \mathcal{F}$, there exists a $t' > t$ such that $\mathcal{F}_t \subset \mathcal{F}_{t'}$.
2. $\theta \to P^\theta(A)$ is measurable with respect to the Borel σ-algebra on Θ for every $A \in \mathcal{F}$.
3. There exists a measurable function $f : \Omega \to \Theta$ such that $f(\omega) = \theta$ almost surely with respect to P^θ.

The first condition ensures that the stochastic process is truly informative by requiring that the filtration expand over time. The second condition is technical, and it guarantees that we can compute probabilities with respect to the user's type. The third condition ensures identifiability: for any point in the sample space there can only be one parameter θ that generates such sample.

Under these assumptions, we can prove a version of Doob's consistency theorem (see Schwartz (1965)):

Proposition 1. If Assumption 1 is satisfied, then the beliefs $\mu_t = \mathbb{E}[\theta|\mathcal{F}_t]$ are consistent.

Assumption 1 is satisfied in various common settings. For example, the exponential clock process we studied in earlier sections fulfills all requirements and is indeed consistent. The assumptions are similarly verified by a 2-dimensional Brownian motion with drift parameter θ, in which case the beliefs take the form of Kalman filters. All results in Section 3.1 generalize immediately to this alternative model.

Theorem 3. *Let the news process $(X_t)_t$ satisfy Assumption 1, and let $r = 0$. Then,*

(Lemma 1) The optimal exercise policy for the second-price auction without reserves is $\tau^*_{2P} = 0$

(Lemma 2) The optimal exercise policy for the first-price auction without reserves is $\tau^*_{1P} = \infty$

(Theorem 1) The revenue π^*_{2P} generated by an optimally-stopped second-price auction dominates the revenue π^*_{1P} generated by an optimally-stopped first-price auction for any value distribution F. In particular

$$\frac{\pi^*_{2P}}{\pi^*_{1P}} = \frac{1}{\mu_0}.$$

The only subtle part of the proof here is that the Revenue Equivalence Theorem no longer holds for all $t > 0$. This is because allocations now depend on the quality scores, which can be asymmetric between the two bidders. Nonetheless, the Revenue Equivalence Theorem holds both at time 0 and in the limit, which allows us to conclude. Some of the results in Section 3.2 are more complex to derive in this setting — for instance, the counterexample in Lemma 4 relies on a form of the value function specific to the Poisson bad news model. Nonetheless, it remains true that the first-price auction weakly dominates the second-price auction when the auctioneer has prior information, because the first-price auction implements the optimal auction. The counterexample provided in Lemma 4 shows that the inequality may in general be strict.

4.2 The Number of Advertisers

Another direction in which we can generalize the results is the number of advertisers. Consider again the exponential news model with discount rate $r = 0$, and let there be $n > 2$ advertisers competing for the ad slot. We can show the following:

Lemma 5. For any realization of the exponential clocks $\tau_1, \ldots, \tau_n$ the second-price auction optimally stops earlier than the first-price auction, i.e.

$$\tau^*_{2P} \leq \tau^*_{1P} \quad \text{pointwise.}$$

Proof. To see this, first note that the revenue from the first-price auction is a maximum of n martingales, and thus is a sub-martingale. Let K_t be the number of advertisers i such that $\mu^i_t \neq 0$. Then,

$$\tau^*_{1P} = \inf_t \{K_t | K_t = 1\}.$$

That is, the auctioneer stops in the limit when there are at least two advertisers who are a good match. She stops earlier when there are less than two good matches and only one competitor remains in the race.

Instead, the revenue from a second-price auction is no longer a super-martingale. The auctioneer's optimal stopping can be characterized via the following HJB:

$$nV(\mu) = V'(\mu)\mu + \mu\Big((n-1)b^{(1)} + b^{(2)}\Big).$$

Continuous-pasting techniques imply that the auctioneer never stops, so long as $n \geq 3$. When $n = 2$ instead the problem becomes a super-martingale, and the auctioneer ends

the auction. Therefore,

$$\tau^*_{2P} = \inf_t \{K_t | K_t = 2\}.$$

The lemma's statement follows immediately from this description of the stopping times. □

From Lemma 5 and the Revenue Equivalence Theorem (which again holds because of the exponential news assumption) immediately follows that the revenue generated by the second-price auction is larger than the revenue generated by a first-price auction.

Our results hold also when the auctioneer has a prior about the advertisers' values. The first-price auction with optimal reserve dominates the second-price auction with optimal reserve. To see this, we can show that the first-price auction implements the optimal auction in this domain, and instead the second-price auction does not. The proof is a replica of the the one discussed in the previous section and is thus omitted.

5 Conclusion

In this paper, we investigate how dynamic scoring influences optimal auction design. Two key countervailing incentives for the auctioneer determine when the auction should be held:

- Information Acquisition: Over time, signals about ad quality improve, allowing the auctioneer to make better allocation decisions.
- Market Thickness: As quality signals improve, advertisers separate between more and less competitive. As the market thins, revenue deteriorates.

We prove that the second-price auction is most sensitive to market thickness: the auction happens early when quality scores are very uncertain. In the first-price auction instead, the auctioneer is most sensitive to information. She wants to learn as much as possible before running the auction, but the bidders anticipate the auctioneer's behavior and shade their bids aggressively. The auctioneer generates more revenue by running a second-price auction. Interestingly, the introduction of reserve prices wipes the second-price auction's advantage. The first-price auction implements the optimal auction, while the second-price auction is not truthful and sometimes ends too early.

A natural question is how introducing an explicit cost of information acquisition would affect the results. We find that most of our results are robust to small information acquisition costs, which can also be interpreted as the instantaneous rate of departure of

the conversational AI user. The analysis is instructive and we include it in Appendix A.1. Additionally, we discuss in Appendix A.2 the problem of a platform that has instead full commitment.

Our theory has interesting implications for ad auction design. In light of the rapid adoption of interactive systems where platforms can dynamically learn the preferences of users, platforms that are planning to monetize using ad-auctions should think carefully about the dynamic forces we highlight when selecting an auction format.

References

Susan Athey and Denis Nekipelov. 2010. *A structural model of sponsored search advertising auctions*. Mimeo.

Santiago R Balseiro, Omar Besbes, and Gabriel Y Weintraub. 2015. Repeated auctions with budgets in ad exchanges: Approximations and design. *Management Science* 61, 4 (2015), 864–884.

Dirk Bergemann and Maher Said. 2010. Dynamic Auctions. *Wiley Encyclopedia of Operations Research and Management Science* (2010).

Tilman Börgers, Ingemar Cox, Martin Pesendorfer, and Vaclav Petricek. 2013. Equilibrium bids in sponsored search auctions: Theory and evidence. *American economic Journal: microeconomics* 5, 4 (2013), 163–187.

Isaias N Chaves and Shota Ichihashi. 2024. Auction timing and market thickness. *Games and Economic Behavior* 143 (2024), 161–178.

Shuchi Chawla and Jason D. Hartline. 2013. Auctions with unique equilibria. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce (EC '13)*. Association for Computing Machinery, New York, NY, USA, 181–196.

Ying-Ju Chen. 2017. Optimal dynamic auctions for display advertising. *Operations Research* 65, 4 (2017), 897–913.

Paul Duetting, Vahab Mirrokni, Renato Paes Leme, Haifeng Xu, and Song Zuo. 2023. Mechanism Design for Large Language Models. *arXiv preprint arXiv:2310.10826* (2023).

Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. 2007. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American economic review* 97, 1 (2007), 242–259.

Soheil Feizi, MohammadTaghi Hajiaghayi, Keivan Rezaei, and Suho Shin. 2023. Online Advertisements with LLMs: Opportunities and Challenges. arXiv:cs.CY/2311.07601

Steven R Grenadier. 2000. Option exercise games: the intersection of real options and game theory. *Journal of Applied Corporate Finance* 13, 2 (2000), 99–107.

Keegan Harris, Nicole Immorlica, Brendan Lucier, and Aleksandrs Slivkins. 2023. Algorithmic Persuasion Through Simulation: Information Design in the Age of Generative AI. *arXiv preprint arXiv:2311.18138* (2023).

Johannes Hörner and Andrzej Skrzypacz. 2017. Learning, experimentation and information design. *Advances in Economics and Econometrics* 1 (2017), 63–98.

Godfrey Keller and Sven Rady. 2015. Breakdowns. *Theoretical Economics* 10, 1 (2015), 175–202.

Carol Marquardt and Emanuel Zur. 2015. The role of accounting quality in the M&A market. *Management Science* 61, 3 (2015), 604–623.

Kalyanapuram Rangachari Parthasarathy. 2005. *Probability measures on metric spaces*. Vol. 352. American Mathematical Soc.

Philip J. Reny. 1999. On the Existence of Pure and Mixed Strategy Nash Equilibria in Discontinuous Games. *Econometrica* 67, 5 (1999), 1029–1056.

Lorraine Schwartz. 1965. On bayes procedures. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete* 4, 1 (1965), 10–26.

Gordon Sick. 1995. Real options. *Handbooks in operations research and management science* 9 (1995), 631–691.

Hal R Varian. 2009. Online ad auctions. *American Economic Review* 99, 2 (2009), 430–434.

Daniel Wangerin. 2019. M&A due diligence, post-acquisition performance, and financial reporting for business combinations. *Contemporary Accounting Research* 36, 4 (2019), 2344–2378.

A Extensions

In this appendix we explore two natural extensions to the baseline model. First, we show that our main insights carry over to a model where information acquisition has a direct cost, in addition to the indirect, market-thickness, cost. Second, we study how could the auctioneer gain from being able to contract on time. Naturally the auctioneer would be better off, but we show that first- and second-price auction remain different even with full commitment.

A.1 Costly Information Acquisition ($r > 0$)

The problem of the auctioneer becomes more complex when she faces two separate costs of information. As before, she pays for information indirectly, by reducing market thickness through delay. Additionally, now information is costly in and of itself, because the user may leave the platform at a rate $r > 0$.

This additional force seems to push the auctioneer toward less delay. However, the same discounting applies to the advertisers, who may be tempted to shade their bids less. Receiving the allocation in the future is less valuable in this world, so they may want to try and influence the timing of allocation by reporting different bids. We begin by showing that such incentives are not sufficient to change the auctioneer's decision in the optimally-stopped second-price auction.

Proposition 2. The optimal exercise policy for the second-price auction is $\tau^*_{2P} = 0$.

Proof. For any pair of reports v_1, v_2 the auctioneer faces the following optimization problem:

$$\max_{\tau} \mathbb{E}_0\left[e^{-r\tau} \min\{\mu^1_\tau b_1, \mu^2_\tau b_2\}\right]$$

We know from Lemma 1 that $\min\{\mu^1_\tau b_1, \mu^2_\tau b_2\}$ is a super-martingale, and the auctioneer here is trying to maximize a discounted super-martingale, which is itself a super-martingale. The result then follows exactly as in Lemma 1. □

Because this result again holds uniformly over reports, the second-price auction simply allocates at time 0 to the highest bidder. This result is independent of the discount factor, which is particularly striking: the auctioneer allocates immediately, and is unwilling to take any risk with respect to the realization of the quality process. In a world in which competition may suddenly collapse (as it happens when one of the clocks ticks) the second-price auction hedges against such risk by sacrificing efficiency. With probability $1 - \mu_0$ the auctioneer will not generate any revenue, but with probability μ_0 she will do

so at a time in which competition is strong. Finally, note that because the second-price auction allocates at a time independent of the reports of the advertisers, advertisers have no incentive to misreport, and the auction is truthful.

Next, we describe the optimal exercise policy of the first-price auction as a function of the advertisers' bids. First note that the auctioneer will never stop after $\min\{\tau_1, \tau_2\}$. In particular, suppose the auctioneer has not stopped yet at $\tau_1 < +\infty$, when the clock of advertiser 1 ticks. Then the auctioneer will stop exactly at τ_1, because the auctioneer's expected revenue at time $t > \tau_1$ is given by $\mathbb{E}_{\tau_1}\left[e^{-rt}\mu_t^2 b_2\right]$ (which is a super-martingale) since advertiser 1 is out of the race. Thus, we only need to consider the case in which no news has arrived yet. In this case, the state space is completely characterized by a single number μ, because $\mu_t^1 = \mu_t^2$ before the arrival of news. The auctioneer's value at beliefs $\mu_t^1 = \mu_t^2 = \mu$ solves the Hamilton-Jacobi-Bellman equation below. Its derivation is in Appendix B.

$$\underbrace{V'(\mu)\lambda\mu(1-\mu)}_{\text{gain from delay}} = \overbrace{rV(\mu)}^{\text{direct cost of delay}} + \underbrace{(1-\mu)\lambda}_{\text{clock ticks}}\Big(\overbrace{V(\mu)-\mu b_1}^{\text{cost if ad 2 is bad}} + \overbrace{V(\mu)-\mu b_2}^{\text{cost if ad 1 is bad}} \Big)$$

The terms in the HJB highlight the balance between benefits and costs of information. The left-hand side represents the increase in continuation value induced by a marginal delay —since beliefs are drifting upwards, if the clock deosn't tick, the expected revenue increases. The right-hand side represents the decrease in continuation value over the marginal delay — first because of information cost (in terms of discounting) and second because a tick of the clock shatters competition. To find out when the auctioneer will run the auction, without loss we set $b_1 \geq b_2$. By requiring continuous pasting on the stopping boundary we find that, if no clock has ticked yet, the auctioneer will stop and run the auction when the beliefs $\mu_t^1 = \mu_t^2$ reach the threshold $\overline{\mu} = \max\left\{1 - \frac{r}{\lambda}\frac{b_1}{b_2}, \mu_0\right\}$. That is, unless a clock ticked the auctioneer runs the auction at

$$\tau^{no-news}(b_1, b_2) = \begin{cases} \frac{1}{\lambda}\log\left(\frac{\lambda}{r}\frac{\min_i b_i}{\max_i b_i} - 1\right) & \text{if } 1 - \frac{r}{\lambda}\frac{\max_i b_i}{\min_i b_i} > \mu_0 \\ 0 & \text{otherwise.} \end{cases}$$

and allocates to the highest bidder. From this equation we glean some insight into the trade-offs faced by the auctioneer. First, the auctioneer runs the auction with the largest delay when both bidders' bids coincide. Intuitively, when the bids are close to each other,

the auctioneer has little to lose from one of the advertisers dropping out of the race. Instead, she wants to anticipate the auction when one bid is substantially larger than the other, because her revenue is now more at risk. This can be seen by noting that the threshold $\overline{\mu}$ is decreasing in the ratio $\frac{b_1}{b_2}$.

This concludes the characterization of the optimal stopping time for the first-price auction , which we report below.

Proposition 3. The optimal exercise policy for the first-price auction is

$$\tau^*_{1P}(b_1, b_2) = \min\left\{\tau^{no-news}(b_1, b_2), \tau_1, \tau_2\right\}$$

Unfortunately, computing the equilibrium of the first-price auction requires aggregating exercise policies over the set of possible reported pairs of bids. This is a largely intractable problem, even for simple value distributions. We prove however that when the probability of the user leaving is sufficiently small, the revenue from a second-price auction dominates the revenue from a first-price auction . In other words, when the advertisers and the auctioneer become more patient (or the beliefs converge sufficiently fast to the truth) the first-price auction allocates almost efficiently, and thus its revenue cannot possibly be close to the revenue of the inefficient second-price auction .

Proposition 4. For every discount rate $r \geq 0$, there exists a pure strategy equilibrium for the first-price auction game. Moreover, for $r = 0$ the unique equilibrium is symmetric.

The key idea of the existence proof relies on the continuity of $\tau^*_{1P}(b_1, b_2)$ on the bids for a fixed discount rate r. In some sense the dynamic first-price auction then inherits the same continuity properties of the static first-price auction , which is sufficiently continuous to apply the machinery of Reny (1999).[13]

To obtain equilibrium uniqueness for the case when $r = 0$, we leverage Theorems 3.1 and 4.6 in Chawla and Hartline (2013). However, a key assumption for their results is that the allocation depends only on the ranking of the bids. When $r > 0$ instead the auction timing is bid-dependent, and thus from an ex-ante perspective bids affect the auction rule. This is not the case when $r = 0$, allowing us to conclude uniqueness and symmetry of equilibrium in that case.

With the existence of equilibrium at hand and the fact that the optimal stopping time τ^*_{1P} is also continuous on the discount rate, we show in Appendix C that the equilibrium

[13] See for instance the paper's Example 5.2, that shows that for static first-price auctions there exists a pure strategy equilibrium with strictly increasing bidding strategies.

bidding strategies are continuous on the discount rate. Therefore we can show a generalized version of Theorem 1 for the case where the information cost – measured by the discount rate – is small.

Theorem 4. *Fix a value distribution F, there exists a discount rate $\underline{r} > 0$ such that for any $r < \underline{r}$ the revenue generated by the second-price auction dominates the revenue generated by any equilibrium of the first-price auction.*

A.2 Contractibility of time

Suppose now that the auctioneer can commit to an auction timing. This commitment power allows the auctioneer to optimize information acquisition without concerns about dynamic inconsistency.

The bidding strategies of advertisers at time 0 hinge on the auctioneer's commitment to a mechanism and its timing. For example, bids will differ if the auctioneer commits to an immediate allocation versus allocating only after identifying a low-quality ad. In this section, the auctioneer can commit to a random variable τ which denotes the time at which she will run the auction. We focus on how first- and second-price auctions drive a balance between acquiring information (by delaying allocation) and ensuring market thickness and robust competition.

Theorem 5. *Let the auction designer commit to the optimal stopping for each format. Then, a first-price auction generates more revenue than a second-price auction.*

We build up to this theorem by proving several results of independent interest. We do this in steps:

- First, we characterize the timing of the optimal mechanism.
- Second, we show that the optimal mechanism can be implemented as a first-price auction.
- Third, we prove that there is no implementation of the optimal mechanism as a second-price auction .

A.3 Ideas of the proof

Step 1: Optimal Mechanism Given a general mechanism $x(v,\mu), p(v,\mu)$ the incentive compatibility constraint can be written as

$$v \in \operatorname*{argmax}_z \mathbb{E}_0\Big[e^{-r\tau}\mu_t\Big(v_i x(z, v_{-i}, \mu_\tau) - p(z, v_{-i}, \mu_\tau)\Big)\Big] \qquad \forall i, v_i, v_{-i}$$

If the allocation rule is monotone, we can rewrite the utility of the bidder as

$$U(v) = \int_0^v \mathbb{E}_{0,v_{-i}}\left[e^{-r\tau(z,v_{-i},\mu_\tau)}x(z,v_{-i},\mu_\tau)\mu_\tau\right]dz$$

and therefore the constraint becomes

$$\mathbb{E}_{0,v}\left[e^{-r\tau}\mu^i_\tau p(v_i,v_{-i},\mu_\tau)\right] = \mathbb{E}_{0,v}\left[e^{-r\tau}\mu^i_\tau v_i x(v_i,v_{-i},\mu_\tau) - \int_0^{v_i} e^{-r\tau(z,v_{-i},\mu_\tau)}\mu^i_\tau x(z,v_{-i},\mu_\tau)dz\right]$$

With the usual Myersonian steps, one can characterize the optimal mechanism as the solution to the following problem:

$$\max_{x,p}\left[\max_\tau \mathbb{E}_0\left[\sum_{i=1}^n e^{-r\tau}\mu_\tau x(v_i,\mu_\tau)\phi(v_i)\right]\right]$$

subject to

$$\sum_{i=1}^n x(v_i,\mu) \le 1 \quad \forall v,\mu$$

where $\phi(v)$ is the virtual value function associated with the distribution F. If both bidders submit values v_1, v_2 such that their virtual values $\phi(v_1), \phi(v_2)$ are negative, the optimal mechanism does not allocate the item. If one of the bidders has a positive virtual value, the auctioneer runs the auction at time 0. This is a direct consequence of the fact that the auctioneer's problem is a discounted martingale (and thus a supermartingale) in this case. Finally, if both bidders report positive virtual values, the auctioneer's value at given beliefs $\mu^i = \mu^j = \mu$ solves the same HJB equation as the one reported in Appendix A.1, with bids replaced by the virtual values. Then, unless a clock ticked, the auctioneer runs the auction at

$$\tau^*(v_1,v_2) = \begin{cases} \frac{1}{\lambda}\log\left(\frac{\lambda}{r}\frac{\min_i \phi(v_i)}{\max_i \phi(v_i)} - 1\right) & \text{if } 1 - \frac{r}{\lambda}\frac{\max_i \phi(v_i)}{\min_i \phi(v_i)} > \mu_0 \\ 0 & \text{otherwise.} \end{cases}$$

and allocates to the highest-virtual-value bidder. This completes the description of the optimal mechanism.

Step 2: The first-price auction implements the optimal mechanism We show that the auctioneer can commit to a stopping time such that the first-price auction achieves the maximum revenue possible. This is the revenue generated by the optimal mechanism,

and we leverage the optimal mechanism's characterization to construct such a stopping time.

Proposition 5. There exists a stopping time τ^*_{1P} such that the first-price auction implements the optimal mechanism.

Intuitively, τ^*_{1P} replicates the stopping of the optimal mechanism. However, the first-price auction without reserves must allocate the item even when the virtual values are negative. To replicate the allocation rule of the optimal mechanism, when virtual values are negative the auctioneer will choose $\tau^*_{1P} = +\infty$, because even if both advertisers have good quality ads, running the auction "at infinity" discounts the contribution of the allocation to the optimization problem. Of course, we still need to show that a first-price auction can indeed implement the allocation for positive virtual values. A complete proof of this result is in Lemma 6 in Appendix B.

Step 3: The second-price auction cannot implement the optimal mechanism Consider the first-price implementation of the optimal mechanism. Discounting and commitment power allow the auctioneer to replicate a reserve price with the timing rule. Moreover, because the optimal τ^* has to solve the HJB problem, if the second-price auction could implement the optimal mechanism then it would have to be the case that $\tau^*_{2P} = \tau^*_{1P}$. Suppose by contradiction that there exists a bidding equilibrium of the second-price auction that implements the optimal mechanism according to τ^*_{1P}. Since the optimal allocation is different for each of the valuations $v \geq \hat{v}$, where $\phi(\hat{v}) = 0$, in equilibrium there is no pooling across types $v \geq \hat{v}$. Second, because $\tau^*_{1P}(\hat{v}, \cdot) = 0$ (i.e., SPA runs at time 0 if one bidder reports $\hat{v}$) and in the optimal mechanism $v_i = \hat{v}$ wins whenever $v_{-i} < \hat{v}$, we have that $\mathbb{E}[b(\hat{v})] > \mathbb{E}[b(v_{-i})|v_{-i}] = \hat{v}$. We show instead that $\mathbb{E}[b(\hat{v})] = \hat{v}$, which completes the contradiction.

Indeed, notice that for any ϵ there exists a sufficiently small ε such that $\tau^*_{1P}(\hat{v}+\varepsilon, \cdot) \leq \epsilon$. That is, the stopping time is sufficiently continuous, so the auction runs close to time 0 (a.s.) when the type lies between $\hat{v}$ and $\hat{v}+\varepsilon$. Suppose by contradiction that $\mathbb{E}[b(\hat{v})] > \hat{v}+\delta$, for some $\delta > 0$, then by the monotonicity of the equilibrium bids we must have that $\mathbb{E}[b(v)] > \hat{v} + \delta$. But then, types $v \in (\hat{v}, \hat{v} + \epsilon)$ have an incentive to mimic type $\hat{v}$, because they would rather pool with type $\hat{v}$ than book a loss whenever their opponent's type lies in $(\hat{v}, \hat{v} + \varepsilon)$. This proves that $\mathbb{E}[b(\hat{v})] = \hat{v}$, which contradicts the assumption and shows that the second-price auction cannot implement the optimal mechanism.

B Calculations for Poisson Model

The evolution of beliefs in a Poisson model follow the derivation below.

$$\begin{aligned}\mu^i_{t+\Delta} &= \mathbb{P}(\theta_i = 1 | \text{ no news before } t+\Delta) \\ &= \frac{\mathbb{P}(\text{no news in } [t, t+\Delta))\mathbb{P}(\theta_i = 1 | \text{ no news before } t)}{\mathbb{P}(\text{no news in } [t, t+\Delta | \text{ no news before } t)} \\ &= \frac{\mu^i_t}{\mu^i_t + (1-\mu^i_t)e^{-\lambda\Delta}}\end{aligned}$$

Then

$$\begin{aligned}\dot{\mu}^i_t = \lim_{\Delta\to 0}\frac{\mu^i_{t+\Delta} - \mu^i_t}{\Delta} &= \lim_{\Delta\to 0}\frac{\mu^i_t}{\Delta}\left(\frac{1}{\mu^i_t + (1-\mu^i_t)e^{-\lambda\Delta}} - 1\right) \\ &= \lim_{\Delta\to 0}\frac{\mu^i_t}{\Delta}\left(\frac{1-\mu^i_t - (1-\mu^i_t)e^{-\lambda\Delta}}{\mu^i_t + (1-\mu^i_t)e^{-\lambda\Delta}}\right) \\ &= \lambda\mu^i_t(1-\mu^i_t)\end{aligned}$$

Fix $b_1 \geq b_2$. The value function for non-zero beliefs is given by

$$\begin{aligned}V(\mu_t, \mu_t) = \max\Bigg\{ & b_1\mu_t, e^{-r\Delta}V(\mu_{t+\Delta}, \mu_{t+\Delta})(\mu_t^2 + 2\mu_t(1-\mu_t)e^{-\lambda\Delta} + e^{-2\lambda\Delta}(1-\mu_t)^2) \\ & + e^{-r\Delta}(\mu_t(1-\mu_t)(1-e^{-\lambda\Delta})(b_1\mu_t + b_2\mu_t) + (1-\mu_t)^2(1-e^{-2\lambda\Delta})\Big(\frac{b_1\mu_t + b_2\mu_t}{2}\Big)\Bigg\}\end{aligned}$$

At a belief where the auctioneer wants to continue,

$$\begin{aligned}V(\mu_t) = & e^{-r\Delta}V(\mu_{t+\Delta})\Big(\mu_t + (1-\mu_t)e^{-\lambda\Delta}\Big)^2 \\ & + e^{-r\Delta}(b_1\mu_t + b_2\mu_t)\left(\mu_t(1-\mu_t)(1-e^{-\lambda\Delta}) + \frac{(1-\mu_t)^2}{2}(1-e^{-2\lambda\Delta})\right)\end{aligned}$$

$$\begin{aligned}0 = & \frac{e^{-r\Delta}V(\mu_{t+\Delta})\Big(\mu_t + (1-\mu_t)e^{-\lambda\Delta}\Big)^2 - V(\mu_t)}{\Delta} \\ & + \frac{e^{-r\Delta}}{\Delta}(b_1\mu_t + b_2\mu_t)\left(\mu_t(1-\mu_t)(1-e^{-\lambda\Delta}) + \frac{(1-\mu_t)^2}{2}(1-e^{-2\lambda\Delta})\right)\end{aligned}$$

Taking the limit for $\Delta \to 0$, and writing $\rho = \frac{r}{\lambda}$ we get

$$V(\mu_t)\Big(r + 2\lambda(1-\mu_t)\Big) = V'(\mu_t)\dot{\mu}_t + (b_1+b_2)\mu_t\Big(\mu_t(1-\mu_t)\lambda + (1-\mu_t)^2\lambda\Big) \iff$$
$$\iff V'(\mu)\mu(1-\mu) = \rho V(\mu) + (1-\mu)\Big(2V(\mu) - \mu(b_1+b_2)\Big)$$

Note that if $r = 0$ the ODE becomes

$$V'(\mu)\mu = 2V(\mu) - \mu(b_1+b_2),$$

as we used in Lemma 4. The solution to the general ODE is

$$V(\mu) = \frac{\mu(1-\mu)}{\rho+1}(b_1+b_2) + K\mu^2\left(\frac{\mu}{1-\mu}\right)^{\rho}$$

where

$$K = \frac{b_1}{(1+\rho)}\left(\frac{b_2}{\rho b_1} - 1\right)^{-\rho}$$

which results in

$$V(\mu) = \frac{\mu(1-\mu)}{\rho+1}(b_1+b_2) + \frac{b_1}{(1+\rho)}\left(\frac{b_2-\rho b_1}{\rho b_1}\right)^{-\rho}\mu^2\left(\frac{\mu}{1-\mu}\right)^{\rho}$$

We find this value of K by requiring continuous pasting at the threshold $\bar{\mu}$ such that $\bar{\mu}b_1 = V(\bar{\mu})$. This threshold turns out to be such that $1-\bar{\mu} = \frac{\rho b_1}{b_2}$. Then we can rewrite

$$K = \frac{b_1}{(1+\rho)}\left(\frac{1}{1-\bar{\mu}} - 1\right)^{-\rho} = \frac{b_1}{(1+\rho)}\left(\frac{1-\bar{\mu}}{\bar{\mu}}\right)^{\rho}$$

and

$$V(\mu) = \frac{\mu(1-\mu)}{\rho+1}(b_1+b_2) + \frac{b_1}{(1+\rho)}\mu^2\left(\frac{\mu(1-\bar{\mu})}{\bar{\mu}(1-\mu)}\right)^{\rho}$$

B.1 Optimal Mechanism

The optimal mechanism problem is the solution to the following problem:

$$\max_{x,p}\left[\max_{\tau}\mathbb{E}_0\Big[\sum_{i=1}^{n}e^{-r\tau}\mu_\tau p(v_i,\mu_\tau)\Big]\right]$$

subject to

$$\sum_{i=1}^{n}x(v_i,\mu)\le 1 \qquad \forall v,\mu$$

$$\text{(DC)}\quad \tau\in\operatorname*{argmax}_t\mathbb{E}_s\Big[\sum_{i=1}^{n}e^{-rt}\mu_t p(v_i,\mu_t)\Big] \qquad \forall v,s,\mu_s$$

$$\text{(IC)}\quad v\in\operatorname*{argmax}_z\mathbb{E}_0\Big[e^{-r\tau}\mu_t\Big(vx(z,v_{-i},\mu_\tau)-p(z,v_{-i},\mu_\tau)\Big)\Big] \qquad \forall i,v_i,\hat{v},v_{-i}$$

$$\text{(IR)}\quad \mathbb{E}_0\Big[e^{-r\tau}\mu_\tau\Big(v_i x(v_i,v_{-i},\mu_\tau)-p(v_i,v_{-i},\mu_\tau)\Big)\Big]\ge 0 \qquad \forall i,v_i,v_{-i}$$

Using the envelope formula (and assuming that the allocation is monotone) we can rewrite the utility of a bidder as

$$U(v)=\int_0^v\mathbb{E}_{0,v_{-i}}\Big[e^{-r\tau(z,v_{-i},\mu_\tau)}x(z,v_{-i},\mu_\tau)\mu_\tau\Big]dz$$

and by substituting and integrating by parts, we get

$$\mathbb{E}_{0,v}\Big[e^{-r\tau}\mu_\tau p(v,\mu_\tau)\Big]=\mathbb{E}_{0,v}\Big[e^{-r\tau}\mu_\tau x(v,\mu_\tau)\phi(v)\Big]$$

Thus, the optimal mechanism problem reduces to

$$\max_{x,p}\left[\max_{\tau}\mathbb{E}_0\Big[\sum_{i=1}^{n}e^{-r\tau}\mu_\tau x(v_i,\mu_\tau)\phi(v_i)\Big]\right]$$

subject to

$$\sum_{i=1}^{n}x(v_i,\mu)\le 1 \qquad \forall v,\mu$$

$$\text{(DC)}\quad \tau\in\operatorname*{argmax}_t\mathbb{E}_s\Big[\sum_{i=1}^{n}e^{-rt}\mu_t x(v_i,\mu_t)\phi(v_i)\Big] \qquad \forall v,s$$

Therefore, we can use the mathematics we developed earlier in this section to describe the optimal mechanism. If both bidders submit values such that $\phi(v_1),\phi(v_2)<0$ then the auctioneer commits to a stopping time $\tau=\infty$. This is equivalent to not allocating the object at all when the discount factor is positive. If instead one of the bidders has a positive virtual value, the auctioneer stops immediately and allocates to that bidder. This is a direct consequence of the fact that the auctioneer's problem in this case is a

discounted martingale. Finally, if both agents have positive virtual values the auctioneer may wait: her decision depends on the ratio $\frac{\rho\phi(v_1)}{\phi(v_2)}$ where $\phi(v_1) \geq \phi(v_2)$.

Lemma 6. The first-price auction implements the optimal mechanism for positive virtual values.

Proof. We use the envelope formula to produce the first-price auction bids that implement the optimal mechanism. Recall that the expected payment at time 0 for the advertisers is given by

$$\mathbb{E}_{0,v}\Big[e^{-r\tau}\mu_\tau p(v,\mu_\tau)\Big] = \mathbb{E}_{0,v}\Big[e^{-r\tau}\mu_\tau x(v,\mu_\tau)\phi(v)\Big]$$

We define

$$p(v_i) = \frac{\mathbb{E}_{0,v}\Big[e^{-r\tau}\mu_\tau p(v,\mu_\tau)\Big|v_i\Big]}{\mathbb{E}_{0,v}\Big[e^{-r\tau}\mu_\tau\Big|v_i\Big]}$$

These payments are increasing, which can be seen by taking a simple derivative. Thus, they are an equilibrium implementation of the optimal mechanism through a first-price auction. □

C Proofs of Section A.1

Proof of Proposition 4. Observe that the our dynamic first-price auction , can be summarized as pay your bid auction where the payoff for bidder i with valuation v_i and given bids b_1, b_2 is $u_i(v_i, b_i, b_{-i} = \mu_0 \cdot (v_i - b_i)x_i(b_i, b_{-i}|r)$ where

$$\begin{aligned} x_i(b_i, b_{-i}|r) = \mathbb{E}_0\Big[&\mathbf{1}_{\{\tau_{-i}\geq\tau^{no-news}(b_i,b_{-i}|r)\}}e^{-r\tau^{no-news}(b_i,b_{-i}|r)}\Big(\mathbf{1}_{\{b_i>b_{-i}\}} + \frac{1}{2}\cdot+\mathbf{1}_{\{b_i=b_{-i}\}}\Big) \\ &+ \mathbf{1}_{\{\tau_{-i}\geq\tau^{no-news}(b_i,b_{-i}|r)\}}e^{-r\tau_{-i}}\Big]. \end{aligned}$$

Observe that for bids $b_i \neq b_{-i}$ the payoff is continuous on b_i as the clock τ_1, τ_2 are independent of the bids and the optimal policy $\tau^{no-news}(b_i, b_{-i}|r)$ is continuous on b_i. Therefore, the same proof of the multi-unit pay your bid auction in Example 5.2 in Reny (1999) applies to our setting which guarantees the existence of a pure strategy equilibrium.

For $r = 0$, notice that $\tau^{no-news}(b_1, b_2|r) = \infty$, and therefore the allocation rule of bidder i only depends on the ranking of the bids but not on the value of the bids. Thus, using the language of Chawla and Hartline (2013), the auction is a rank-based allocation rule. Moreover, the first price nature of the auction implies that it satisfies the bid-based payment and win-vs-tie-strict properties required for Theorems 3.1 and 4.6 in Chawla and

Hartline (2013) which imply that there is a unique equilibrium for the game and bidders use the same pure strategy bidding function. □

To show that that the equilibrium is continuous as function of the discount rate, we consider the mixed extension of the auction game.[14] We denote the Borel measures on $[0,1]^2$ by $\mathcal{B}([0,1]^2)$ so that a strategy $B \in \mathcal{B}([0,1]^2)$ corresponds to a probability measure $(v,b) \sim B$ with $\mathbb{P}_B[v' \leq v] = F(v)$ for all $v \in [0,\overline{v}]$ (almost surely). For example, a pure strategy $b(v)$ has mixed strategy representation B_b where $\mathbb{P}_{B_b}[(v,b) : b = b(v)] = 1$.

We endow $\mathcal{B}([0,1]^2)$ with the weak topology, so that $B_n \Rightarrow B$ if for all continuous function $f : [0,1]^2 \to R$ we have that $\int f dB_n \to \int f dB$. Because $\mathcal{B}([0,1]^2)$ is compact and the subsspace of mixed strategies is closed, we conclude that the subspace of mixed bidding strategies is compact.

Claim 1. *Given a discount rate $r \geq 0$, consider an equilibrium $(B_1^*(r), B_2^*(r))$. Then the probability of having a tie on the equilibrium is zero.*

Proof. Suppose that by contradiction that there are ties on the equilibrium. So consider a type v_i on where there is a tie. By equilibrium optimality we must have that $\mathbb{P}_{B_i^*(r)}[(v_i,b) : b \geq v_i] = 0$: else, the bidder obtains a payoff of zero while by bidding small ϵ obtains a positive payoff since it wins the object with positive probability whenever the other bidder has a valuation $v_{-i} \leq \epsilon$. Since type v_i is bidding less that its value, if there is a tie it can raise its bids by a small ϵ. By doing that it removes the ties and gets a mass increase on their payoff, and because τ_{1P}^* is continuous the time allocation changes smoothly so that would be a profitable deviation. Therefore, there are no ties on the equilibrium. □

Claim 2. *Consider $r_n \to 0$ and equilibrium bids $(B_1^*(r_n), B_2^*(r_n))$ then $(B_1^*(r_n), B_2^*(r_n)) \Rightarrow (B_{b^*}, B_{b^*})$, where $b^*(\cdot)$ is the symmetric bidding strategy in the first-price auction game when $r = 0$. Furthermore, if $B_i^*(r_n)$ is a a pure strategy for all n, then $b^*(v|r_n) \to b^*(v|0)$ for all $v \in [0,\overline{v}]$ (a.s.).*

Proof. Because the space of actions is compact, consider a subsequence of $(B_1^*(r_{n_k}), B_2^*(r_{n_k})) \Rightarrow (B_1^*(0), B_2^*(0))$. We assert that $(B_1(0), B_2(0))$ is a nash equilibrium in the mixed extension of the game, which by Proposition 4 implies that $B_i(0) = B_{b^*}$. This further implies that is the unique accumulation point of the sequence $(B_1^*(r_n), B_2^*(r_n))$ which, given that the space is compact, implies that $B_i^*(r_n) \Rightarrow B_{b^*}$.

Next, observe that $e^{-r\tau^{no-news}(b_i,b_{-i}|r_n)}$ is increasing as $n \to \infty$ and converges pointwise to 1, thus by Dini's Theorem we have that it converges uniformly to 1. The same convergence result can be used for the indicator functions $\mathbf{1}_{\{\tau_{-i} \geq \tau^{no-news}(b_i,b_{-i}|r)\}}$, $\mathbf{1}_{\{\tau_{-i} \geq \tau^{no-news}(b_i,b_{-i}|r)\}}$

[14] We refer to Chapter 2 of (Parthasarathy, 2005) for textbook treatment on the space of Borel measures and the weak topology on it.

which implies that $x_i(b_i, b_{-i}|r_n)$ converges uniformly to $x_i(b_i, b_{-i}|r_n)$ and therefore that $u_i(v_i, b_i, b_{-i}|r_{n_k})$ converges uniformly to $u_i(v_i, b_i, b_{-i}|r_{n_k})$.

Since there are no ties on the equilibrium (Claim 1) so that the payoff functions are continuous when there are no ties (see proof of Proposition 4), and the payoff function has uniform convergence, we have that $U_i(B_1^*(r_{n_k}), B_2^*(r_{n_k})|r_{n_k}) \to U_i(B_1(0), B_2(0)|r = 0)$ where

$$U_i(B_1(r), B_2(r), r) = \int \int u_i(v_i, b_i, b_{-i}|r) dB_{-i}(v_{-i}, b_{-i}) dB_i(v_i, b_i).$$

To show that $(B_1(0), B_2(0))$ is a Nash Equilibrium, consider a pure strategy b. First, suppose that b is increasing on v. Observe that $U_i(B_i^*(r_{n_k}), B_{-i}^*(r_{n_k})|r_{n_k}) \geq U_i(B_b, B_{-i}^*(r_{n_k})|r_{n_k})$ since $(B_i^*(r_{n_k}), B_{-i}^*(r_{n_k}))$ is a Nash Equilibrium when the discount rate is r_{n_k}. Because b is increasing we know that there are no ties so the payoff function is continuous and has uniform converge as $n_k \to \infty$. Therefore, by taking the limit we conclude that $U_i(B_1(0), B_2(0)|r = 0) \geq U_i(B_b, B_2(0)|r = 0)$. If the function b is non-decreasing, then for every $\epsilon > 0$ an increasing function $\tilde{b} \geq b$ exists such $U_i(B_b, B_2(0)|r = 0) \leq U_i(B_{\tilde{b}}, B_2(0)|r = 0) + \epsilon$. We conclude by taking $\epsilon \to 0$, that for all feasible bidding strategies $U_i(B_1(0), B_2(0)|r = 0) \geq U_i(B_b, B_2(0)|r = 0)$. Therefore, the uniqueness of equilibrium when $r = 0$ implies that $B_i(0) = B_{b^*}$.

Finally, if $B_i^*(r_n)$ is a sequence of pure strategy equilibria. Consider f_ϵ a continuous approximation of the indicator function $\mathbf{1}_{[v-\epsilon, v+\epsilon]}$. Because $B_i^*(r_n) \Rightarrow B_{b^*}$ we have that $\int f_\epsilon dB_i^*(r_n) \to \int f_\epsilon dB_{b^*}$. Thus, by taking ϵ we get the desired pointwise convergence result. □

Proof of Theorem 4. From Proposition 2, we have that the optimal stopping is independent of the discount rate $r = 0$, which in turn implies that the auctioneer's revenue when optimally running a second-price auction is independent of r. Thus, $\pi_{2P}^*(r) = \pi_{2P}^*(0)$.

For the first-price auction case, we want to claim that for $r_n \to 0$, and any mixed equilibrium of the game when $r_n > 0$ we have that $\pi_{1P}^*(r_n) \to \pi_{1P}^*(0)$. Indeed, consider a sequence of mixed bidding equilibrium $(B_1^*(r_n), B_2^*(r_n))$. From Claim 2, we have that $B_1^*(r_n) \Rightarrow B_{b^*}$, where b^* correspond to the symmetric bidding strategy in the unique equilibrium of the first-price auction game when $r = 0$ (Proposition 4). Fixing τ_1, τ_2, we have that $\tau_{1P}^*(b1, b2|r_n) \to \min\{\tau_1, \tau_2\} := \tau_1 \wedge \tau_2$ converges pointwise and $\tau_{1P}^*(b1, b2|r_n)$ is increasing on r_n. Then, Dini's Theorem implies that the convergence is uniformly on b_1, b_2.

Therefore, the following limit holds

$$\begin{aligned}
&\lim_{n\to\infty} \int\int e^{-r\tau^*_{1P}(b_1,b_2|r_n)} \min\{\mu^1_\tau b_1, \mu^2_\tau b_2\} dB^*_1(r_n) dB^*_2(r_n) \\
&= \int\int \min\{\mu^1_{\tau_1\wedge\tau_2} b_1, \mu^2_{\tau_1\wedge\tau_2} b_2)\} dB^*_{b^*} dB^*_{b^*} \\
&= \min\{\mu^1_{\tau_1\wedge\tau_2} b^*(v_1), \mu^2_{\tau_1\wedge\tau_2} b^*(v_2)\}
\end{aligned}$$

By taking expectations on τ_1, τ_2, we conclude that $\pi^*_{1P}(r_n) \to \pi^*_{1P}(0)$.

Finally, we invoke Theorem 1 that shows that $\pi^*_{1P}(0) < \pi^*_{2P}(0)$ and conclude that a $\underline{r} > 0$ exists such that for $r \le \underline{r}$, $\pi^*_{1P}(r) < \pi^*_{2P}(r)$. □