REPEATED BIDDING WITH DYNAMIC VALUE

BENJAMIN HEYMANN, ALEXANDRE GILOTTE, RÉMI CHAN-RENOUS

ABSTRACT. We consider a repeated auction where the buyer's utility for an item depends on the time that elapsed since his last purchase. We present an algorithm to build the optimal bidding policy, and then, because optimal might be impractical, we discuss the cost for the buyer of limiting himself to shading policies.

1. INTRODUCTION

1.1. **Repeated auctions with dynamic value: a planning problem.** This article analyses a *repeated second price auctions* where the buyer's utility for the item depends on the time elapsed since his last purchase of a similar item: an item value depends on the *age* of the previous purchase. Throughout the paper, We use ad auctions as a motivating application. The rise of online marketplaces and digital advertising have fuelled the study of *repeated auctions* along several research axis, in particular budget/ROI constraints, learning, and strategic interactions. However, surprisingly, we enter into much less explored territory if the buyer's utility for an item *depends on the previous auctions outcomes*, which is notwithstanding a reasonable belief for use cases such as digital marketing, where it can be beneficial to space the marketing interactions over time because of the user's *display fatigue*. Yet, finding an optimal bidding policy in this widespread setting has not been done yet, even when the auction is second price. This article aims to fill the gap. We display in Fig. 1 an example of valuation dynamic: after an auction is won, the value of the next auctions drops and then increases back to its initial value as time goes by. At this point, it should be intuitive to the reader that bidding the value (the *greedy* strategy) is suboptimal, even though the auction is second price. We illustrate this by showing in Fig. 3 the payoff of several linear scaling strategies. The greedy strategy corresponds to a scaling factor of 1 and is, in our context, **never** optimal.

This is in sharp contrast with the widespread belief that good enough feature engineering does the job. This intuition is flawed because it overlooks the underlying *planning problem*. Indeed, in the minimal model we introduce, the bidder has access to the exact immediate value of the item, but we still observe that bidding this value is suboptimal, even in the second price setting. Last, we mention that in the edge case where the buyer is interested in buying at most one item, we get a setting close in spirit to [23], who introduce a procurement version of the prophet problem, where the goal is to minimize the procurement cost.

1.2. **Display advertising auctions.** We proceed with some contextual elements on our motivating use case, display advertising auctions. Digital advertising allows the monetizing of publisher content programmatically. When a user reaches one of the publisher's pages, the right to show a display banner to the user is sold. The mechanism that elicits the deal is often an auction that takes place in real time, hence the name: Real-Time-Bidding (**RTB**). Real-Time-Bidding resents several
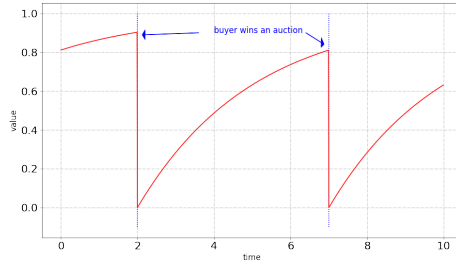

FIGURE 1. Example of dynamical value. When the bidder wins an auction, his value drops brutally, and then progressively goes back to a maximum value. An optimal bidding strategy should account for this phenomenon.

practical challenges that have motivated a growing body of scientific work at the frontier of mathematics, economics, and computer sciences [11]. The buyer typically derives a valuation formula for the display opportunity of the shape [8, 6, 11]
$Value\_Per\_Event \times \Pr(E|X)$, where $\Pr(E|X)$ is the probability that an event will take place in the future (such as click, sales. . . it depends) knowing what is available in the user's context (here denoted by X) and $Value\_Per\_Event$ is a multiplier that is independent of X. This multiplier often integrates business-related constraints such as budget [6] or ROI or cost per action. Such a valuation then becomes the input of a bidding module, tasked to find a bid that maximizes the immediate payoff (*i.e.* expected value minus cost). For a second-price auction, this bidding module simply returns the estimated value. For non-incentive compatible auctions such as first-price auctions, methods to estimate the competition have been proposed [29].
User fatigue and frequency capping. One specificity of display advertising is that bidders typically receive a sequence of auctions for opportunities to display an ad to the same user. Those auctions arise as the user browses the web and opens new pages. However, it is usually not seen as desirable to display too many similar ads to the same user in a short period of time, and a very common practice consists in defining a *frequency capping* [11], which is a maximum number of ads which can be displayed to the same user for the same campaign on a predefined period of time (such as "maximum 10 displays per day per user"). The bidder would thus stop bidding on a user when the maximum number of ads is reached, until the next period. The main rational for applying such frequency capping is essentially that ads displayed to the same user usually have diminishing return. Those return could even become negative, when too many similar ads in a short period of time could result in a degraded user experience. Several studies also point to the usefulness of spreading the ads in time [26, 9].
Fatigue as a feature of the predictors. To deal with this decreasing return of the ads, another frequent method consists in directly modelling how the *past* displayed ads affect the value of the current auctioned ad opportunity. This can be done, for example by measuring the *fatigue* as the number of ads displayed to the user in a past period of time, and using this fatigue feature as a predictor of the $\Pr(E|X)$ model. Several works point to a similar idea, such as [2, 24, 4]. We would like to note here that while these features improve the immediate value estimation, they do not account for the impact of an ad on the future opportunities, that is, they still miss the planning problem.
Impact of winning one auction on future opportunities. Indeed, one logical consequence of the diminishing return of ads is that winning an auction *now* impacts the value of the *next* opportunities for the same user. But while it is possible that the

$\Pr(E|X)$ term of the valuation recognize the fatigue effect, there is a hole in the literature on how the bidder component should address the problem. The problem is recognized in [8] but no solution is proposed. In a related vein, [14] propose a heuristic factor to decrease the bids after a click, but the method does not fully address the problem. In this paper, we propose to study how the bidder should optimally bid in a simplified model inspired by this setting. We note that the frequency capping could be interpreted as setting the value of the items to 0 during some time after the capping is reached. We simplify this by assuming that the value of the auctioned item depends solely on the time elapsed since the last won auction. We also note that today most RTB auctions are first-price. We decided to study the case of second price auction instead; both for simplicity, and to emphasize that even in this simpler second price setting the problem is non-trivial.

1.3. **Repeated auctions.** Several streams of literature related to repeated auctions have flourished around the online marketing use case. Without aiming at being complete, we mention some of them in this section and discuss how the *repetition* of the auction structure the problems under scrutiny. First, the existence of a budget or an ROI constraint over a sequence of auctions couples all the individual bidding problems together. In the literature, such constraints are typically analyzed either by using a constraint relaxation or by solving an associated knapsack problem. The typical solution is then a scaling of the value with a well-chosen factor. See for instance [6, 19, 12, 17, 10]. In this stream of research, this is mostly the coupling of the bidding problems that is under scrutiny. In this line of research, *repetition* is mostly about coupled optimization problems. Second, the level of bid can be seen as a decision that is repeated over time. The outcome (value and payment) depends on this decision. This setting can be studied with the bandit framework. See for instance [28, 1]. In this stream of research, this is mostly the explore-exploit tradeoff of the learning bidder that is under scrutiny, and *repetition* is mostly about learning. Third, there is a line of research that puts under scrutiny the interactions between the seller and the multiple buyers. In this steam of work, *repetition* is mostly about playing a repeated game [5, 16, 18, 3, 20, 25, 13].

1.4. **Contribution.** The applied work [8] recognizes the difficulty to handling the coupling between present and future bidding decisions. We introduce a realistic minimal model to study this coupling, and frame the design of a bidding policy as a continuous time optimal control problem over an infinite horizon. We present structuring properties of this control problem, such as the monotony of the bidding policy and a useful identity for the expected payoff. We then introduce an algorithm that iterates toward the optimal policy. Interestingly, the proof relies on a dynamical system argument, more precisely, a refined version of the Cauchy-Lipschitz Theorem. The algorithm allows us to compare through Monte-Carlo simulations the optimal bidding policy and the greedy policy, which corresponds to the economic cost of not accounting for future opportunities in the bidding policy design. We then take a look at shading policies because they constitute a natural class of simple approximations to look at. While we show that shading policies cannot be solutions to the optimal control problem, we also observe that in the numerical experiments, they
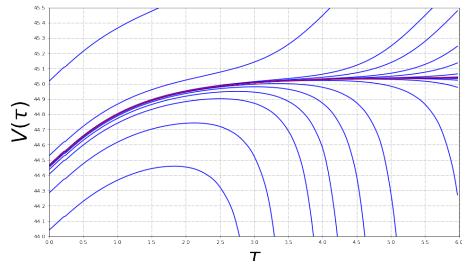

FIGURE 2. An example of Algorithm 1 run. In red is the output of the algorithm, in blue the iterates.

perform very well. We also derive a closed form
expression of the buyer payoff in a specific case.
In what follows, we use the masculine pronoun
he to refer to the buyer, and the feminine pro-
noun *she* for the seller. Proofs are postponed to the supplementary material.

2. MODEL AND BIDDING PROBLEM

This section introduces a simple environment where the value of an item depends
on the time since the last won auction, that we simply refer to as the system *age*
thereafter. The section ends with Lemma 1 that characterizes the value function
with a differential equation. An auctioneer receives items that she needs to sell
immediately to a set of buyers, and we take the perspective of one of those bidders
facing a stationary competition. The items arrivals follow a Poisson process of
intensity μ. We denote by τ the *age*, which is the time that has elapsed since the
last won auction. We assume that the value of an item for the bidder depends
solely on the age τ, and note this value $k(\tau)$. We also suppose that k is a non-
decreasing and bounded function. Fig. 1 shows an example of how the value of
buying a new item may evolve over a timeline. We suppose the auctions to be
second price, with an iid competition with a known CDF, q. This competition
may include a reserve price set by the seller. In particular, the story includes the
posted price scenario. Thus, $q(b)$ is the probability that the buyer wins with a
bid equal to b. Finally, we assume that the sequence ends at a random time T_γ,
where T_γ follows an exponential distribution of parameter γ, or equivalently we
say that the buyer has a discount rate γ. The bidder's goal is to maximize, in
expectation, the sum of the values of the auctions he wins, minus the cost he pays,
before the end of the sequence. Here, we note that this environment may be viewed
as a continuous Markov process, where the state is fully summarized by the age τ.
We denote by $p(b)$ the average payment of the user when bidding b. Because the
auction is second price, we have the relation $p(b) = q(b)b - \int_0^b q(t)\mathrm{d}t$. We denote
by $U(v,b) \overset{def}{=} q(b) \cdot v - p(b)$ the expected utility when bidding b and valuing v for
winning.

Because the state depends solely on the age τ, the bidder's policy can be fully
described by a measurable function $\tau \longrightarrow b(\tau) \in \mathbb{R}^+$ specifying the bid when the
age is τ. We note $\mathcal{B}$ the set of these bidding functions. Then, for a bidding function
$b \in \mathcal{B}$, we note $V_{b,\gamma,\mu}(\tau)$ the expectation of the bidder's future revenue when the
state is τ, or $V_b(\tau)$ for simplicity, when the context is clear. Let $T_1, T_2 \ldots T_n \ldots$,
the time of the next auctions, $C_1 \ldots C_n \ldots$ the competition at these times, and τ the
current age, then $V_b(\tau) \overset{def}{=} \mathbb{E} \sum_{i=1}^{\infty} e^{-\gamma T_i} \left(k(\tau(T_i)) - C_i \right) \mathbf{1}\{b(\tau(T_i)) > C_i\}$. We are
ready to state Lemma 1.

Lemma 1. *For $b \in \mathcal{B}$, V_b satisfies the differential equation*

$$V_b'(\tau) = -\mu \cdot U\big(k(\tau) + V_b(0) - V_b(\tau), b(\tau)\big) + \gamma V_b(\tau).$$

The proof of Lemma 1 relies on the dynamic programming principle applied to V_b.

3. OPTIMAL BID

This section contains most of the theoretical results. We define the optimal value
function (or Bellman value), $V^\star(\tau)$, as the sup, on all possible bidding policies, of
$V_b(\tau)$, namely $V^\star(\tau) \overset{def}{=} \sup_{b \in \mathcal{B}} V_b(\tau)$. It is practical to introduce the notation $\pi(v)$
for the expected profit of a bidder of value and bid v in a one shot auction, namely

$\pi(v) \overset{def}{=} U(v,v) = q(v)v - p(v)$. This corresponds to the expected profit of the incentive compatible bid in the one shot auction. Lemma 2 provides a dynamic programming characterization of the Bellman value, and provides a relation between the Bellman value and the bid. The Bellman value is then proved to be the solution of a parametrized differential equation of unknown parameter (Lemma 3). Theorem 1 ensures that when k is concave, the policy is monotone increasing. Theorem 1 can be thought of as a well posedness sufficient condition: primitives for which the optimal bid would not be monotone could arguably be considered pathological from an econometric standpoint. For Section 3.3, that presents the algorithm for computing an optimal policy, we suppose that q is continuous, and thus π is C^1. [1]

3.1. **Differential Equation.** From the proof of Lemma 1, we get the following characterizations of $V^\star$ and the optimal bid $b^\star$

Lemma 2. *We have the relation* $V_t^\star = \int_0^{+\infty} \mu e^{-(\mu+\gamma)t}\Big(\pi(k_t + V_0^\star - V_t^\star) + V_t^\star\Big)\mathrm{d}t$. *Moreover, there exists an optimal bid $b^\star$, and it satisfies the relation*
$$b^\star(\tau) = \max\left(0; k(\tau) + V^\star(0) - V^\star(\tau)\right) \ .$$

This differential equation is complicated to solve directly because its initial value $V_0^\star$ also appears as a parameter of the dynamic. To clarify this point, we define $\Phi : \mathbb{R} \times \mathbb{R}_+ \times \mathbb{R}_+ \to \mathbb{R}$ such that $\Phi(t, v, \lambda) = \gamma v - \mu\pi(k_t + \lambda - v)$. We then reformulate lemma 2 as a more usual Cauchy problem:

Lemma 3. *The value function $V^\star$ is the solution of the ordinary differential equation*

$(\mathcal{F}_{y_0})$
$$\begin{cases} \dot{Y}_t = \Phi(t, Y_t, y_0) \\ Y_0 = y_0 \end{cases}$$

for some $y_0 \in \mathbb{R}_+$.

It should be noted that in Lemma 3, the parameter y_0 is not given.

3.2. **Differential equation for b.** In this section, we provide a parametrization of the differential equation so that the $b^\star$ appears as solution. Such parametrization requires additional assumptions on k. The optimal bid also follows a differential equation:

Lemma 4. *If k is C^1, then, on every open on which $b^\star(t) > 0$, the optimal bid $b^\star$ satisfies the differential equation* $\dot{b}_t = \dot{k}_t - \gamma(V_0^\star + k_t) + \mu\pi(b_t) + \gamma b_t$.

Since k is increasing with τ, it may seem natural to expect that the optimal bid would also increase with τ. The following result confirms this intuition under the additional assumption that k is concave.

Theorem 1. *If k is concave, then $b^\star$ is increasing with τ, and strictly increasing on any interval where k strictly increases.*

---

[1]Indeed, by definition, $\pi(v) = q(v)v - p(v) = \int_0^v q(t)\mathrm{d}t$

We insist that this is not true without the concavity hypothesis: Fig. 5 proposes a counter example where k is increasing but not concave.

3.3. **Algorithm.** Lemma 3 does not allow us to derive $V^\star$ from a Cauchy problem because the initial condition is also a parameter of the dynamics. Still, by the Cauchy-Lipschitz Theorem, the solution of the ordinary differential equation

$$(\mathcal{F}_{y_0,\lambda}) \qquad\qquad \begin{cases} \dot{Y}_t = \Phi(t, Y_t, \lambda) \\ Y_0 = v_0 \end{cases}$$

admits a unique maximal solution $Z^{y_0,\lambda} : t \to Z^{y_0,\lambda}(t)$ for any $y_0 > 0$ and $\lambda > 0$. We also observe that the problem $\mathcal{F}_{y_0,y_0}$ is the same at the problem of equation $\mathcal{F}_{y_0}$, and we define $Z^v(t) \overset{def}{=} Z^{v,v}(t)$ for all $t \geq 0$. In the following, we assume that q is continuous, and thus π is C^1. This allows to apply the differentiable dependency Theorem, which is the key to our approach. Lemma 5 is the most technical result of the paper, and is leveraged in the proof of Theorem 2.

Lemma 5. *Suppose q is continuous. The value $V_0^\star$ is the unique v for which $\lim_{t\to+\infty} Z^v(t))$ is finite.*

We then use Lemma 5 to prove that a simple dichotomy on the parameter allows us to solve the bidding problem for repeated auctions with dynamic value.

Theorem 2 (Algorithm). *Suppose q is continuous. Let $y^{(n)}$ the iterations of Algorithm 1, then $||V_0^\star - y^{(n)}(0)|| = O(1/2^n)$.*

Algorithm 1 Computing V_0

Require: (a, b, N)
 1: **for** $n \leq N$ **do**
 2:     $c \leftarrow (a + b)/2$
 3:     Compute $Z^{c,c}$, solution of $\mathcal{F}_c$
 4:     $y^{(n)} \leftarrow Z^{c,c}$
 5:     **if** $t \to y^{(n)}(t)$ **then**
 6:         $a \leftarrow c$
 7:     **else**
 8:         $b \leftarrow c$
 9:     **end if**
10: **end for**

3.4. **Numerical estimation of the cost of impatience.** Algorithm 1 is implemented in Python 3. We use SOLVE_IVP, an ODE solver from SCIPY [27] to numerically solve the ODE. We rely on the default Runge-Kutta method [15, 22]. To test the different bidding policies, we use Lemma 8 in the Appendix to say that, for $\gamma \to 0$, we can compute the time average performance over a long enough period. For the experiments, reported in Table 1, we used $\gamma = 0.01$ to construct the optimal policy. We ran simulation on this optimal policy and on the greedy policy. We tested several values for μ, with $k(\tau) = 1 - e^{-\tau}$ and $k(\tau) = 1 - 1/(1 + \tau)$.

μ	0.1	1	5	10	100	μ	0.1	1	5	10	50
optimal	0.045	0.22	0.43	0.52	0.72	optimal	0.035	0.16	0.32	0.40	0.69
greedy	0.043	0.20	0.34	0.38	0.46	greedy	0.035	0.14	0.26	0.31	0.44

TABLE 1. Estimated value per time unit with the greedy bidder and with the optimal bidder, for different values of μ, and γ close to 0 with $k(t) = 1 - e^{-t}$(left) and $k(t) = 1 - 1/(1 + t)$ (right)

4. Shading strategy

A shading strategy consists in bidding a constant factor α (smaller than 1) of the value, *i.e.* $b(\tau) = \alpha k(\tau)$. This class of strategy is of great practical and theoretical interest because it is simple to implement and to analyze. It also satisfies some theoretical properties. For instance, as mentioned in the introduction, shading is used as a way to implement budget or ROI constraints. We use in this section the shorthand notation $V_\alpha = V_{\alpha k, \gamma, \mu}$. We found it notable that on the settings we tested, the class of shading strategies performs very well. It can be proven however that it is not optimal in general. Indeed, by Lemma 4, the optimal bidding policy satisfies $\dot{b}_t = \dot{k}_t - \gamma(V_0^\star + k_t) + \mu\pi(b_t) + \gamma b_t$, which implies that a shading policy $b_t = \alpha k_t$ that would outperform all other strategies should satisfy $\alpha\dot{k}_t = \dot{k}_t - \gamma(V_0^\star + k_t) + \mu\pi(\alpha k_t) + \gamma\alpha k_t$, which in general does not hold.

4.1. **Special case.** We mention that there is an instance of the problem where the time average payoff of a shading strategy can be computed explicitly. We use Theorem 3 to build a visualization of the average expected payoff, that we display in Fig. 3. We check that shading strategies outperform truthfulness in the presence of a value dynamic. We also check that the greater the arrival rate μ, the stronger the shading should be (the α multiplier close to zero).

Theorem 3. *If the competition is a uniform distribution on* $[0,1]$, *if* $k_t = 1 - 1/(1 + t)$ *then for all* α *satisfying* $\mu\alpha > 1$

$$\lim_{\gamma \to 0} \gamma V_0^\alpha = (1 - 1/2\alpha)(\mu\alpha)^2 \frac{\Gamma(\mu\alpha - 1, \mu\alpha)}{\Gamma(\mu\alpha + 1, \mu\alpha)},$$

where $\Gamma(\_, \_)$ *is the upper incomplete gamma function.*[2]

4.2. **Numerical comparison.** Figure 4 displays the proportion $V_\alpha/V^\star$ of the optimal value retrieved by bidding αk, for several values of α and μ, and two distinct concave functions k (and with γ was close to 0). We note that on all these settings, the best value of α gives results very close to the optimal value. (it was higher than 99%) Another observation is that for smaller values of μ, the greedy bidder (*i.e.* $\alpha = 1$) performs near optimally. This observation is easily explained: a small μ means that there are few auctions per time unit; it is thus likely that the first auction only arrives a long time after 0 , when the function $k(t)$ is already close to its sup. This is thus not so different

FIGURE 3. Illustration of Theorem 3

from the case when k is constant, for which the greedy bidder is optimal. On the

---

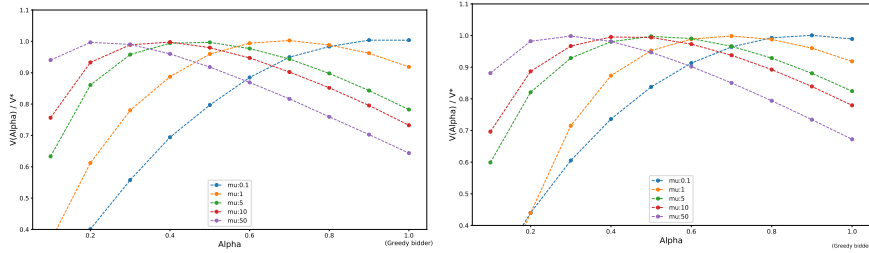[2]defined as $\Gamma(s, x) = \int_x^\infty t^{s-1}e^{-t}\mathrm{d}t$

FIGURE 4. Ratio $V_\alpha/V^\star$ as a function of α with $k_\tau = 1 - e^{-t}$ and $\mu = 5$ (left) and with $k(t) = 1 - {}^1\!/_{1+t}$
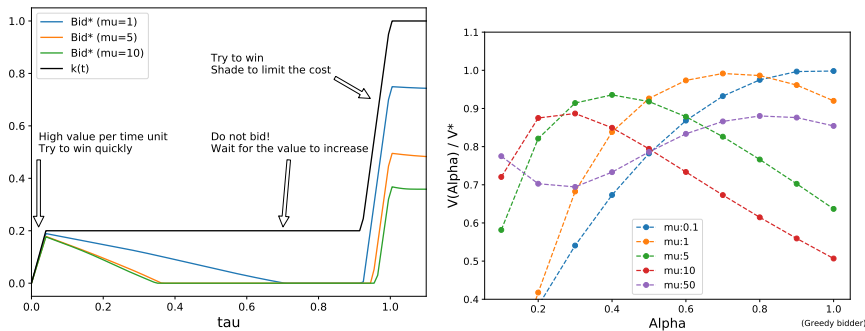

FIGURE 5. (left) k and optimal bid as a function of τ, for different values of μ (right). Optimal bids and V_α with the two steps function $k(t) = \min(5t, 0.2) + \max(0, \min(0.8, 10t - 9.2)$

other hand, for larger values of μ, the greedy bidder performs poorly; and the best shading factor decreases with the density of auctions μ.

Example with $k(\tau)$ non concave. We have seen in Lemma 2 that $b^\star$ is an increasing function when k is concave. We propose here an example where k is not concave to observe what can happen in this case. The function k we used is depicted in Fig. 5. It is a simple piece-wise linear function, making two steps. The optimal bids as a function of τ are plotted for different values of μ: we can observe that they are not monotonous. The intuition for this behavior is simple: since the function k increases quickly near 0, it is worth winning early because this generates a high value per time unit. Thereafter, the function stops increasing until the second step, where it increases very sharply again. Just before the increase, it is better not to bid at all : the bidder should indeed avoid resetting the state to 0 just before this strong increase. Fig. 5 shows the value obtained when submitting a shaded bid $k(t) \times \alpha$, with the same two steps k function. We note that the best V_α here is quite sensibly smaller than $V^\star$. We also note that for large values of μ, V_α is not a concave function of α and may have several local maxima. This example also contradicts another assumption which could seem intuitive: $b^\star$ is *not* an increasing function of k. Indeed, if we removed the second step on the k function above, then we would get a concave function, and the optimal bid would be positive for all $t > 0$.

4.3. **Asymptotic convergence.** Theorem 4 provides a complementary perspective on why shading policies might work well. Basically, as μ goes to $+\infty$, the best shading policy performs closer and closer to the optimal policy.

Theorem 4. *Set* $\mathfrak{p}(b) = \mathbb{E}(C|C \leq b)$, *where* C *is the price to beat. Suppose* k *concave and such that* $k \in C^1$ *and* $k_0 = 0$. *Suppose* $q \in C^1$, *and* $\dot{q}(0) > 0$, *then*

$$\sup_{b \in \mathcal{B}} \lim_{\mu \to +\infty} \lim_{\gamma \to 0} \frac{V_{b,\gamma,\mu}(0) - V_{k,\gamma,\mu}(0)}{V_{b,\gamma,\mu}(0)} = \dot{\mathfrak{p}}(0) = \lim_{t \to 0} \frac{\dot{q}(t) \int_0^t q(s)\mathrm{d}s}{q(t)^2} \ ,$$

Moreover,

$$\sup_{b \in \mathcal{B}} \lim_{\mu \to +\infty} \lim_{\gamma \to 0} V_{b,\gamma,\mu}(0) = \sup_{\alpha \in [0,1]} \lim_{\mu \to +\infty} \lim_{\gamma \to 0} .V_{\alpha \cdot k,\gamma,\mu}(0)$$

5. Conclusion

Auction systems with dynamic values are everywhere. This work is a first step toward better understanding them. Further research could include studying what happens when the bidder does not know the dynamics in advance, and needs to learn, as well as extending the results to more general auctions and more general value dynamics.

References

[1] Juliette Achddou, Olivier Cappé, and Aurélien Garivier. Efficient algorithms for stochastic repeated second-price auctions. In *Algorithmic Learning Theory*, pages 99–150. PMLR, 2021.

[2] Deepak Agarwal, Bee-Chung Chen, and Pradheep Elango. Spatio-temporal models for estimating click-through rate. In *Proceedings of the 18th international conference on World wide web*, pages 21–30, 2009.

[3] Shipra Agrawal, Eric Balkanski, Vahab Mirrokni, and Balasubramanian Sivan. Robust repeated first price auctions. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, EC '21, page 4, New York, NY, USA, 2021. Association for Computing Machinery.

[4] Michal Aharon, Yohay Kaplan, Rina Levy, Oren Somekh, Ayelet Blanc, Neetai Eshel, Avi Shahar, Assaf Singer, and Alex Zlotnik. Soft frequency capping for improved ad click prediction in yahoo gemini native. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, pages 2793–2801, 2019.

[5] Kareem Amin, Afshin Rostamizadeh, and Umar Syed. Repeated contextual auctions with strategic buyers. *Advances in Neural Information Processing Systems*, 27, 2014.

[6] Santiago R. Balseiro and Yonatan Gur. Learning in repeated auctions with budgets: Regret minimization and equilibrium. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, EC '17, page 609, New York, NY, USA, 2017. Association for Computing Machinery.

[7] D. Bertsekas. *Reinforcement Learning and Optimal Control*. Athena Scientific optimization and computation series. Athena Scientific, Nashua, 2019.

[8] Martin Bompaire, Alexandre Gilotte, and Benjamin Heymann. Causal models for real time bidding with repeated user interactions. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pages 75–85, 2021.

[9] Michael Braun and Wendy W Moe. Online display advertising: Modeling the effects of multiple creatives and individual impression histories. *Marketing science*, 32(5):753–767, 2013.

[10] Matteo Castiglioni, Andrea Celli, Alberto Marchesi, Giulia Romano, and Nicola Gatti. A unifying framework for online optimization with long-term constraints. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022.

[11] Hana Choi, Carl F Mela, Santiago R Balseiro, and Adam Leary. Online display advertising markets: A literature review and future directions. *Information Systems Research*, 31(2):556–575, 2020.

[12] Vincent Conitzer, Christian Kroer, Eric Sodomka, and Nicolas E. Stier-Moses. Multiplicative pacing equilibria in auction markets. *Operations Research*, 70(2):963–989, 2022.

[13] Yuan Deng, Vahab Mirrokni, and Hanrui Zhang. Posted pricing and dynamic prior-independent mechanisms with value maximizers. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022.

[14] Eustache Diemert, Julien Meynet, Pierre Galland, and Damien Lefortier. Attribution modeling increases efficiency of bidding in display advertising. In *Proceedings of the ADKDD'17*, ADKDD'17, New York, NY, USA, 2017. Association for Computing Machinery.

[15] John R Dormand and Peter J Prince. A family of embedded runge-kutta formulae. *Journal of computational and applied mathematics*, 6(1):19–26, 1980.

[16] Alexey Drutsa. Horizon-independent optimal pricing in repeated auctions with truthful and strategic buyers. In *Proceedings of the 26th International Conference on World Wide Web*, WWW '17, page 33–42, Republic and Canton of Geneva, CHE, 2017. International World Wide Web Conferences Steering Committee.

[17] Yuan Gao, Kaiyu Yang, Yuanlong Chen, Min Liu, and Noureddine El Karoui. Bidding agent design in the linkedin ad marketplace. *arXiv preprint arXiv:2202.12472*, 2022.

[18] Negin Golrezaei, Adel Javanmard, and Vahab Mirrokni. Dynamic incentive-aware learning: Robust pricing in contextual auctions. *Advances in Neural Information Processing Systems*, 32, 2019.

[19] Benjamin Heymann. Cost per action constrained auctions. In *Proceedings of the 14th Workshop on the Economics of Networks, Systems and Computation*, pages 1–8, 2019.

[20] Yash Kanoria and Hamid Nazerzadeh. Dynamic reserve prices for repeated auctions: Learning from bids. In Tie-Yan Liu, Qi Qi, and Yinyu Ye, editors, *Web and Internet Economics*, pages 232–232, Cham, 2014. Springer International Publishing.

[21] Raphaël Krikorian. M1 systèmes dynamiques, 2021.

[22] F Shampine Lawrence. Some practical runge-kutta formulas. *Mathematics of Computation*, 46:135–150, 1986.

[23] Vasilis Livanos and Ruta Mehta. Prophet inequalities for cost minimization. *arXiv preprint arXiv:2209.07988*, 2022.

[24] Hao Ma, Xueqing Liu, and Zhihong Shen. User fatigue in online news recommendation. In *Proceedings of the 25th International Conference on World Wide Web*, pages 1363–1372, 2016.

[25] Thomas Nedelec, Clément Calauzènes, Noureddine El Karoui, Vianney Perchet, et al. Learning in repeated auctions. *Foundations and Trends® in Machine Learning*, 15(3):176–334, 2022.

[26] Navdeep S Sahni. Effect of temporal spacing between advertising exposures: Evidence from online field experiments. *Quantitative Marketing and Economics*, 13(3):203–247, 2015.

[27] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020.

[28] Jonathan Weed, Vianney Perchet, and Philippe Rigollet. Online learning in repeated auctions. In *Conference on Learning Theory*, pages 1562–1583. PMLR, 2016.

[29] Tian Zhou, Hao He, Shengjun Pan, Niklas Karlsson, Bharatbhushan Shetty, Brendan Kitts, Djordje Gligorijevic, San Gultekin, Tingyu Mao, Junwei Pan, Jianlong Zhang, and Aaron Flores. An efficient deep distribution network for bid shading in first-price auctions. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, New York, aug 2021. ACM.

Appendix

—

Proof of Section 2

Expected value when an auction starts in state τ. Let $W_b(\tau)$ be the expected revenue of the bidder when an auction is received in state τ. Then $W_b(\tau)$ can be written by taking the expectation on the outcome of the immediate auction. On average, the bidder pays the second price $p(b(\tau))$. We then consider two possible outcomes. Either the auction is lost: then the bidder pays 0, stays in state τ, and its expected future reward is $V(\tau)$. Or he wins the auction: the bidder then receives a reward $k(\tau)$ and the state is reset to 0, which means that the bidder's expected reward after the auction is $V_b(0)$. Wrapping all this together, we get

$$W_b(\tau) = q(b(\tau)) \cdot \big(k(\tau) + V_b(0))\big) - p(b(\tau)) + (1 - q(b(\tau))) \cdot V_b(\tau)$$
$$= V_b(\tau) + q(b(\tau)) \cdot \big(k(\tau) + V_b(0) - V_b(\tau)\big) - p(b(\tau))$$
$$(1) \qquad = V_b(\tau) + U\big(k(\tau) + V_b(0) - V_b(\tau), b(\tau)\big).$$

Bellman equation. We note T the waiting time until the next auction. For our stationary Poisson process, T always follows an exponential distribution, of parameter μ. We can write the evolution of the system as an expectation on T:

$$(2) \qquad V_b(\tau) = \mathbb{E}_T\big(e^{-\gamma T} \cdot W_b(\tau + T)\big) = \int_0^{+\infty} e^{-(\gamma+\mu)t} \cdot \mu \cdot W_b(\tau + t)dt.$$

Lemma 1.

Proof. This is obtained directly by writing the evolution of V from τ to $\tau + \delta$. Similarly to (2), let T the time of the next auction, and δ a positive time. We split the expectation of V in two parts, when T is either less than δ or larger than δ. We have $V_b(\tau) = \int_0^\delta e^{-(\gamma+\mu)t} \cdot \mu \cdot W_b(\tau + t)dt + e^{-(\gamma+\mu)\delta} \cdot V_b(\tau + \delta)$. From this we deduce that for all $\tau, \delta > 0$ $\frac{V_b(\tau+\delta) - V_b(\tau)}{\delta} = -\frac{1}{\delta}\int_0^\delta e^{-(\gamma+\mu).t} \cdot \mu \cdot W_b(\tau + t)dt + \frac{1 - e^{-(\gamma+\mu)\delta}}{\delta} V_b(\tau + \delta)$. The right-hand side is bounded, from which we first conclude that V_b is continuous (and thus W is continuous too); we may now take the limit when $\delta \to 0$ of the right-hand size, which is $-\mu \cdot W_b(\tau) + (\gamma + \mu) \cdot V_b(\tau)$. This proves both the derivability and the formula. □

Proof of Section 3

Lemma 2.

Proof. If a bidding function achieves the sup in the definition of $V^\star$, then, since the auction is second price, we can increase the expected return at any point τ by bidding the true value $k(\tau) + V^\star(0) - V^\star(\tau)$. This implies that if there is an optimal bidding function, it can only be $b^\star$ as defined above. The existence of an optimal bidding function is a consequence of the classical result that the Bellman operator is a contraction: For any value function $\tau \to V(\tau)$, we define the bidding policy $B(V) := \tau \to \max(0; k(\tau) + V(0) - V(\tau))$. Because the auction is second price, bidding with $B(V)$ is the optimal action in any state τ if the value we get after this auction if defined by V. From this, we deduce that for any b, $V_{B(V_b)} \geq V_b$, by recurrence on the number of steps where we use $B(V_b)$ instead of b. We also note that the operator $V \to V_{B(V)}$ is a contraction for the norm ∞ (eg: [7]), from which we deduce that any sequence defined by $V_{i+1} = B(V_i)$ will converge

to the same value. Since this sequence is also increasing, it must converge to $V^\star$, and the associated bids converges to $b^\star := B(V^\star)$. We plug $b^\star$ in Lemma (1) to get $V_t^\star = \int_0^{+\infty} \mu e^{-(\mu+\gamma)t} \Big(U(b_t^\star, k_t + V_0^\star - V_t^\star) + V_t^\star \Big) \mathrm{d}t$. This is maximized for $b_t^\star = k_t + V_0^\star - V_t^\star$, in which case, $U(b_t^\star, k_t + V_0^\star - V_t^\star) = \pi(k_t + V_0^\star - V_t^\star)$ $\square$

Lemma 4.

Proof. From Lemma 2, $b_t^\star = k_t + V_0^\star - V_t^\star$, hence since by Lemma 1, $V^\star$ is differentiable, $\dot{b}_t^\star = \dot{k}_t - \dot{V}_t^\star$. Replacing in the latter expression $\dot{V}_t^\star$ by $\gamma V_t^\star - \mu\pi(k_t + V_0^\star - V_t^\star)$ from Lemma 3, we get $\dot{b}_t^\star = \dot{k}_t - \gamma V_t^\star + \mu\pi(k_t + V_0^\star - V_t^\star) = \dot{k}_t - \gamma(k_t + V_0^\star - b_t^\star) + \mu\pi(b_t^\star)$. Hence $\dot{b}_t = \dot{k}_t - \gamma(V_0^\star + k_t) + \mu\pi(b_t) + \gamma b_t$. $\square$

Theorem 1.

Proof. When k is smooth, the proof follows from the fact that $\dot{b}^\star(t)$ is the sum of $\dot{k}_t - \gamma(V_0^\star + k_t)$, which is (strictly) decreasing with t under the hypotheses, and $\mu\pi(b^\star) + \gamma b^\star$ which varies in the same direction as $b^\star$. We treat the case when k is strictly increasing, (the case k increasing is similar), and we proceed ad absurdum. Suppose that $b^\star$ is non-increasing on an interval $]t_1, t_2[$. Then b' must be strictly decreasing on this interval, and thus $b'(t_2) < 0$. This ensures that the interval can be extended, and thus $b^\star$ would be decreasing up to $+\infty$, at a rate at least $b'(t_2) < 0$. This implies that the solution of the differential equation 2 would go to $-\infty$. We remind here that this differential equation 2 is no longer valid when $b^\star$ takes negative values (the bid cannot become negative). But this still means that the bid would become 0 at some point t_0, and the same argument proves it cannot become positive again after that: the bid would thus be 0 on $]t_0; +\infty[$. But bidding 0 on $]t_0, +\infty[$ would generate a return of 0; which clearly is not optimal. We conclude that $b^\star$ must be strictly increasing. For non-smooth k, the concavity still implies that k must be C^1 with k' decreasing on a dense subset of R^+. The result still holds by slightly adapting the previous proof, noting that b' is defined on the same set as k', and is equal to k' plus a continuous function. This implies that if b' decreasing and negative just before t_2, it must still be negative just after t_2. $\square$

Lemma 3 does not allow us to derive $V^\star$ from a Cauchy problem because the initial condition is also a parameter of the dynamics. Still, by the Cauchy-Lipschitz Theorem, the solution of the ordinary differential equation

$$(\mathcal{F}_{y_0, \lambda}) \qquad \begin{cases} \dot{Y}_t = \Phi(t, Y_t, \lambda) \\ Y_0 = v_0 \end{cases}$$

admits a unique maximal solution $Z^{y_0, \lambda} : t \to Z^{y_0, \lambda}(t)$ for any $y_0 > 0$ and $\lambda > 0$. We also observe that the problem $\mathcal{F}_{y_0, y_0}$ is the same at the problem of equation $\mathcal{F}_{y_0}$, and we define $Z^v(t) := Z^{v,v}(t)$ for all $t \geq 0$.

Lemma 6 (Taylor extension). *Set $u_t = u_t(v) = k_t + v - Z^v(t)$, then*

$$\frac{\partial Z^v(t)}{\partial v} = e^{\gamma t + \mu \int_0^t q(u_s(v))\mathrm{d}s} \Big(1 - \mu \int_0^t e^{-\gamma s - \mu \int_0^s q(u_l(v))\mathrm{d}l} q(u_s(v))\mathrm{d}s \Big)$$

Proof. By the differentiable dependency Theorem [21], for any $t > 0$, $(y_0, \lambda) \to Z^{y_0, \lambda}(t)$ is differentiable, and its differential maps a perturbation $(\Delta v, \Delta \lambda)$ to $\Delta Z_{y_0, \lambda}^{\Delta v, \Delta \lambda}$, which is the solution of

$$(3) \qquad \begin{cases} \dot{\Delta Z}_{y_0, \lambda}^{\Delta v, \Delta \lambda}(t) = D_y\Phi(t, Z^{y_0, \lambda}(t), \lambda).\Delta Z_{y_0, \lambda}^{\Delta v, \Delta \lambda}(t) + D_\lambda\Phi(t, Z^{y_0, \lambda}(t), \lambda).\Delta \lambda \\ \Delta Z_{y_0, \lambda}^{\Delta v, \Delta \lambda}(0) = \Delta v. \end{cases}$$

By definition $Z^v(t) = Z^{\cdot\cdot}(t) \circ m(v)$ where $m(v) = (v, v)$, so that by derivation of a composition,

$$D_v(Z^v(t))[\Delta v] = D_{m(v)}Z^{\cdot\cdot}(t).D_v m[\Delta v] = D_{m(v)}Z^{\cdot\cdot}(t).(\Delta v, \Delta v) = \Delta Z_{v,v}^{\Delta v, \Delta v}(t)$$

therefore, using (3),

(4)
$$\dot{D_v Z}^v(t)[\Delta v] = \dot{\Delta Z}_{v,v}^{\Delta v, \Delta v}(t) = D_y \Phi(t, Z^{v,v}(t), v).\Delta Z_{v,v}^{\Delta v, \Delta v}(t) + D_\lambda \Phi(t, Z^{v,v}(t), v).\Delta v.$$

We now compute the two partial derivatives. First since $\Phi(t, y, \lambda) = \gamma y - \mu\pi(k_t + \lambda - y)$, $D_y \Phi(t, y, \lambda) = \gamma + \mu\dot{\pi}(k_t + \lambda - y)$. Now since $\pi(v) = q(v)v - p(v) = \int_0^v q(t)dt$ we have $\dot{\pi}(v) = q(v)$, therefore $D_y \Phi(t, y, \lambda) = \gamma + \mu q(k_t + \lambda - y)$ Similarly, $D_\lambda \Phi(t, y, \lambda) = -\mu\dot{\pi}(k_t + \lambda - y) = -\mu q(k_t + \lambda - y)$. Therefore, combining with (4)

$$\dot{D_v Z}^v(t)[\Delta v] = D_y \Phi(t, Z^{v,v}(t), v).\Delta Z_{v,v}^{\Delta v, \Delta v}(t) + D_\lambda \Phi(t, Z^{v,v}(t), v).\Delta v$$
$$= (\gamma + \mu q(k_t + v - Z^{v,v}(t))).\Delta Z_{v,v}^{\Delta v, \Delta v}(t) - \mu q(k_t + v - Z^{v,v}(t))\Delta v$$
$$= (\gamma + \mu q(k_t + v - Z^v(t))).D_v(Z^v(t))[\Delta v] - \mu q(k_t + v - Z^v(t))\Delta v$$

We observe that the equation on $D_v Z^v(t)[\Delta v]$ is of the shape $\dot{X}_t = a_t X_t + b_t$ with $X_0 = \Delta v$. The solution of such ODE is $X_t = e^{\int_0^t a_s ds}(\Delta v + \int_0^t e^{-\int_0^s a_u du}b_s ds)$. Therefore,

$$D_v Z^v(t)[\Delta v] =$$

$$e^{\int_0^t (\gamma + \mu q(k_s + v - w_s(v)))ds}\left(\Delta v - \int_0^t e^{-\left(\int_0^s (\gamma + \mu q(k_l + v - w_l(v)))dl\right)}\mu q(k_s + v - w_s(v))\Delta v ds\right)$$

$$= \left(e^{\gamma t + \mu \int_0^t q(u_s(v))ds}\left(1 - \mu \int_0^t e^{-\gamma s - \mu \int_0^s q(u_l(v))dl}q(u_s(v))ds\right)\right)\Delta v$$

Hence $D_v Z^v(t) = e^{\gamma t + \mu \int_0^t q(u_s(v))ds}\left(1 - \mu \int_0^t e^{-\gamma s - \mu \int_0^s q(u_l(v))dl}q(u_s(v))ds\right)$. $\qquad\square$

Lemma 7 (Bound on the derivative). *For all $t \geq 0$, $\frac{\partial Z^v(t)}{\partial v} \geq e^{\gamma t}$.*

Proof. By Lemma 6,

$$\frac{\partial Z^v(t)}{\partial v} = e^{\gamma t + \mu \int_0^t q(u_s(v))ds}\left(1 - \mu \int_0^t e^{-\gamma s - \mu \int_0^s q(u_l(v))dl}q(u_s(v))ds\right)$$

$$\geq e^{\gamma t + \mu \int_0^t q(u_s(v))ds}\left(1 - \int_0^t e^{-\mu \int_0^s q(u_l(v))dl}\mu q(u_s(v))ds\right)$$

$$= e^{\gamma t + \mu \int_0^t q(u_s(v))ds}\left(1 + [e^{-\mu \int_0^s q(u_l(v))dl}]_0^t\right)$$

$$= e^{\gamma t + \mu \int_0^t q(u_s(v))ds}\left(e^{-\mu \int_0^t q(u_l(v))dl}\right) = e^{\gamma t}.$$

$\qquad\square$

Corollary 1 (Uniqueness of the stable solution). *There can be at most one $y_0 \in \mathbb{R}$ such that $\lim_{t \to +\infty} Z^{y_0}(t)$ is finite.*

Proof. We deduce from Lemma 7, that for any t, and any $v_1 < v_2$, $Z^{v_2}(t) > Z^{v_1}(t) + (v_2 - v_1)e^{\gamma t}$. Now, if $Z^{v_1}(t)$ converges to a finite value, then $Z^{v_2}(t)$ goes to $+\infty$. This guarantees that $\lim_{t \to +\infty} Z^v(t)$ is finite for at most one v. $\qquad\square$

Lemma 5.

Proof. The value function satisfies $\mathcal{F}_{V_0}$, and is bounded and positive. By Corollary 1, it is the only bounded solution. $\qquad\square$

Theorem 2.

Proof. By corollary 1, the condition "$t \to y^{(n)}(t)$ is diverging to $+\infty$" is equivalent to $y^{(n)}(0) > V_0^\star$. Hence, if $V_0^\star$ is in $[a, b]$ at step (n), it will also be in $[a, b]$ at step $(n+1)$. □

PROOF OF SECTION 4

Let T_w the random variable "age at the time of the next won auction". We set $A(t) = \int_0^t \gamma + \mu q(b(s))\mathrm{d}s$. We also define the random variable N_b as the number of won auctions until the end of the sequence T_γ. We may now state how to compute V_b and its initial value:

Lemma 8. *The solution V_b is given by:*

$$(5) \qquad V_b(\tau) = e^{A(\tau)} \left(V_b(0) - \int_0^\tau e^{-A(t)} \cdot \mu \cdot U(V_b(0) + k(t), b(t))dt \right)$$

$$(6) \qquad V_b(0) = \mathbb{E}(N_b) \times \mathbb{E}\left(U(k(T_w), b(T_w))|T_w < T_\gamma \right).$$

Proof. Equation (5) is the integral solution of the linear differential equation of Lemma 1. Equation (6) states that $V(0)$ is precisely the expected number of won auctions multiplied by the average return on those won auctions, which seem intuitively obvious. We may prove it more formally as follows, by first rewriting (5) as :

$$e^{-A(\tau)} \cdot V_b(\tau) = V_b(0) \cdot \left(1 - \int_0^\tau e^{-A(t)} \cdot \mu \cdot q(b(t))dt \right) - \int_0^\tau e^{-A(t)} \cdot \mu \cdot U(k(t), b(t))\, dt$$

Since we know that $V_b(\tau)$ is bounded, and $A(\tau) > \gamma.\tau \xrightarrow{\tau \to \infty} +\infty$, the left-hand side goes to 0 when $\tau \to +\infty$ Thus

$$V_b(0) = \frac{\int_0^\infty e^{-A(t)} \cdot \mu \cdot U(k(t), b(t))\, dt}{1 - \int_0^\infty e^{-A(t)} \cdot \mu \cdot q(b(t))dt}$$

$$= \frac{\int_0^\infty e^{-A(t)} \cdot \mu \cdot U(k(t), b(t))\, dt}{\Pr(T_w < T_\gamma)} \times \frac{\Pr(T_w < T_\gamma)}{1 - \Pr(T_w < T_\gamma)}.$$

Here we used the formula $\int_0^\infty e^{-A(t)} \cdot \mu \cdot q(b(t))dt = \Pr(T_w < T_\gamma)$. We conclude by identifying the two factors on the right side of equation (7) as some expectations:

$$\mathbb{E}\left(U(k(T_w), b(T_w))|T_w < T_\gamma \right) = \frac{\int_0^\infty e^{-A(t)} \cdot \mu \cdot U(k(t), b(t))\, dt}{\Pr(T_w < T_\gamma)}$$

is the average return per won auction, and $\mathbb{E}(N_b) = \frac{\Pr(T_w < T_\gamma)}{1 - \Pr(T_w < T_\gamma)}$ is the expected number of won auctions. (We note that it follows a geometric law of parameter $\Pr(T_w < T_\gamma)$, hence the formula.) □

Limit when γ goes to 0. Since the process ends after a time T_γ following an exponential law, the expected lifetime is $^1/\gamma$. Intuitively, increasing the length of the sequence should increase linearly the expected return, as with the same bidding strategy we will increase the number of won items without changing their average price or value. The following lemma confirms this intuition for long sequences, by proving that the average return per time unit, $\gamma \cdot V_{b,\gamma}(0)$ is converging:

Lemma 9.

$$(7) \qquad \lim_{\gamma \to 0} \gamma V_{b,\gamma}(0) = \mu \frac{\int_0^{+\infty} e^{-\mu \int_0^t q(b_s)\mathrm{d}s} U(k_t, b_t)\mathrm{d}t}{\int_0^{+\infty} e^{-\mu \int_0^t q(b_s)\mathrm{d}s}\mathrm{d}t}$$

Proof. We have $\Pr(T_w < T_\gamma) = \int_0^{+\infty} e^{-\gamma t} e^{-\int_0^t q(b_s)\mathrm{d}s} \mu q(b_t))\mathrm{d}t = 1 - \gamma \int_0^{+\infty} e^{-\mu \int_0^t q(b_s)\mathrm{d}s - \gamma t}\mathrm{d}t$ by integration by parts. We then take the limit in Equation (7). □

Theorem 3.

Proof. From Lemma 9, we have:

$$\lim_{\gamma \to 0} \gamma V_0^\alpha = (1 - 1/2\alpha) \frac{\int_0^\infty \alpha \mu k_t^2 e^{-(\mu\alpha \int_0^t k_s \mathrm{d}s)} \mathrm{d}t}{\int_0^\infty e^{-(\mu\alpha \int_0^t k_s \mathrm{d}s)} \mathrm{d}t}.$$

Since $k_t = 1 - \frac{1}{1+t}$, we have $\int_0^t k_s \mathrm{d}s = t - \ln(1+t)$ and $e^{-(\mu\alpha \int_0^t k_s \mathrm{d}s)} = e^{-\alpha\mu t}(1+t)^{\alpha\mu}$. For $m \geq 0$, we have

$$\int_0^{+\infty} (1+t)^{-m} e^{-\alpha\mu t} (1+t)^{\alpha\mu}$$

$$= \int_0^{+\infty} e^{-\alpha\mu t} (1+t)^{\alpha\mu - m} \mathrm{d}t$$

$$= \int_1^{+\infty} e^{-\alpha\mu(x-1)} x^{\alpha\mu - m} \mathrm{d}x$$

$$= e^{\alpha\mu} \int_1^{+\infty} e^{-\alpha\mu x} x^{\alpha\mu - m} \mathrm{d}x$$

we set $\theta = (\alpha\mu)^{-1}$, $k = (\alpha\mu - m + 1)$, $\tilde{\alpha} = \alpha\mu$

$$= e^{\tilde{\alpha}} \int_1^{+\infty} x^{k-1} e^{-\frac{x}{\theta}} \mathrm{d}x$$

$$= e^{\tilde{\alpha}} \theta^k \int_{1/\theta}^{+\infty} t^{k-1} e^{-t} \mathrm{d}t$$

$$= e^{\tilde{\alpha}} \theta^k \Gamma(k, 1/\theta) = e^{\tilde{\alpha}} \tilde{\alpha}^{m-1-\tilde{\alpha}} \Gamma(\tilde{\alpha} - m + 1, 1/\theta).$$

Therefore

$$(1 - 1/2\alpha) \frac{\int_0^\infty \alpha \mu k_t^2 e^{-(\mu\alpha \int_0^t k_s \mathrm{d}s)} \mathrm{d}t}{\int_0^\infty e^{-(\mu\alpha \int_0^t k_s \mathrm{d}s)} \mathrm{d}t}$$

$$= \alpha\mu(1 - 1/2\alpha) \frac{\tilde{\alpha}^{-1-\tilde{\alpha}} \Gamma(\tilde{\alpha} + 1, 1/\theta) + \tilde{\alpha}^{1-\tilde{\alpha}} \Gamma(\tilde{\alpha} - 1, 1/\theta) - 2\tilde{\alpha}^{-\tilde{\alpha}} \Gamma(\tilde{\alpha}, 1/\theta)}{\tilde{\alpha}^{-1-\tilde{\alpha}} \Gamma(\tilde{\alpha} + 1, 1/\theta)}$$

$$= (1 - 1/2\alpha) \frac{\tilde{\alpha} \Gamma(\tilde{\alpha} + 1, 1/\theta) + \tilde{\alpha}^3 \Gamma(\tilde{\alpha} - 1, 1/\theta) - 2\tilde{\alpha}^2 \Gamma(\tilde{\alpha}, 1/\theta)}{\Gamma(\tilde{\alpha} + 1, 1/\theta)}$$

$$= (1 - 1/2\alpha) \frac{\tilde{\alpha}(\Gamma(\tilde{\alpha} + 1, 1/\theta) - \tilde{\alpha}\Gamma(\tilde{\alpha}, 1/\theta)) - \tilde{\alpha}^2(\Gamma(\tilde{\alpha}, 1/\theta) - \tilde{\alpha}\Gamma(\tilde{\alpha} - 1, 1/\theta))}{\Gamma(\tilde{\alpha} + 1, 1/\theta)}$$

Now, we use the property $\Gamma(s + 1, 1/\theta) = s\Gamma(s, 1/\theta) + \theta^{-s} e^{-1/\theta}$ to get

$$(1 - 1/2\alpha) \frac{\tilde{\alpha}(\Gamma(\tilde{\alpha} + 1, 1/\theta) - \tilde{\alpha}\Gamma(\tilde{\alpha}, 1/\theta)) - \tilde{\alpha}^2(\Gamma(\tilde{\alpha}, 1/\theta) - (\tilde{\alpha} - 1)\Gamma(\tilde{\alpha} - 1, 1/\theta)) + \tilde{\alpha}^2 \Gamma(\tilde{\alpha} - 1, 1/\theta))}{\Gamma(\tilde{\alpha} + 1, 1/\theta)}$$

$$= (1 - 1/2\alpha) \frac{\tilde{\alpha}(\theta^{-\tilde{\alpha}} e^{-1/\theta}) - \tilde{\alpha}^2(\theta^{-(\tilde{\alpha}-1)} e^{-1/\theta}) + \tilde{\alpha}^2 \Gamma(\tilde{\alpha} - 1, 1/\theta))}{\Gamma(\tilde{\alpha} + 1, \tilde{\alpha})}$$

$$= (1 - 1/2\alpha) \frac{\tilde{\alpha}(\tilde{\alpha}^{\tilde{\alpha}} e^{-\tilde{\alpha}}) - \tilde{\alpha}^2(\tilde{\alpha}^{\tilde{\alpha}-1} e^{-\tilde{\alpha}}) + \tilde{\alpha}^2 \Gamma(\tilde{\alpha} - 1, \tilde{\alpha})}{\Gamma(\tilde{\alpha} + 1, \tilde{\alpha})}$$

$$= (1 - 1/2\alpha) \tilde{\alpha}^2 \frac{\Gamma(\tilde{\alpha} - 1, \tilde{\alpha})}{\Gamma(\tilde{\alpha} + 1, \tilde{\alpha})}.$$

$\square$

Theorem 4.

Proof. Let a smooth bid policy b satisfying $b_t \leq k_t$. We have $U(k_t, b_t) = q(b_t)(k_t - \mathfrak{p}(b_t))$. Set $Q_t = \int_0^t q(b_s)\mathrm{d}s$. Then by integration by part, and because $0 = k_0 = b_0$ by assumption,

$$\int_0^{+\infty} \mu e^{-\mu \int_0^t q(b_s)\mathrm{d}s} U(k_t, b_t)\mathrm{d}t$$

$$= \int_0^{+\infty} \mu e^{-\mu Q_t} q(b_t)(k_t - \mathfrak{p}(b_t))\mathrm{d}t$$

$$= \int_0^{+\infty} e^{-\mu Q_t}(\dot{k}_t - \dot{b}_t \dot{\mathfrak{p}}(b_t))\mathrm{d}t.$$

Now observing that

$$\lim_{\mu \to +\infty} \frac{\int_0^{+\infty} e^{-\mu Q_t}(\dot{k}_t - \dot{b}_t \dot{\mathfrak{p}}(b_t))\mathrm{d}t}{\int_0^{+\infty} e^{-\mu Q_t}\mathrm{d}t} = \dot{k}_0 - \dot{b}_0 \dot{\mathfrak{p}}(b_0) = \dot{k}_0 - \dot{b}_0 \dot{\mathfrak{p}}(0),$$

and combining with Lemma 9, we get

$$\lim_{\mu \to +\infty} \lim_{\gamma \to 0} \gamma V_{b,\gamma,\mu}(0) = \dot{k}_0 - \dot{b}_0 \dot{\mathfrak{p}}(0)$$

We see that only the derivative of b at 0 matters here, so we can without loss of generality suppose that $b_t = \alpha k_t$ for some $\alpha \leq 1$. We end up with $\dot{k}_0(1 - \alpha \dot{\mathfrak{p}}(0))$. Since $\dot{\mathfrak{p}}(0) > 0$, this quantity is minimized for $\alpha = 1$ and maximized for $\alpha = 0$, otherwise said, the left-hand side of the Theorem is equal to $\dot{\mathfrak{p}}(0)$. Then we observe that $\mathfrak{p} = \frac{\int_0^t s\mathrm{d}(q(s))}{q(s)}$. The remaining follows after integration by part and taking the limit at 0. $\square$