Repeated Auctions with Budgets in Ad Exchanges: Approximations and Design

Santiago R. Balseiro*, Omar Besbes†, Gabriel Y. Weintraub†

*Fuqua School of Business, Duke University
†Graduate School of Business, Columbia University
srb43@duke.edu, ob2105@columbia.edu, gyw2105@columbia.edu

First version: September, 2012
This version: March, 2014

Abstract

Ad Exchanges are emerging Internet markets where advertisers may purchase display ad placements, in real-time and based on specific viewer information, directly from publishers via a simple auction mechanism. Advertisers join these markets with a pre-specified budget and participate in multiple second-price auctions over the length of a campaign. This paper studies the competitive landscape that arises in Ad Exchanges and the implications for publishers' decisions. The presence of budgets introduces dynamic interactions among advertisers that need to be taken into account when attempting to characterize the bidding landscape or the impact of changes in the auction design. To this end, we introduce the notion of a Fluid Mean Field Equilibrium (FMFE) that is behaviorally appealing, computationally tractable, and in some important cases yields a closed-form characterization. We establish that a FMFE approximates well the rational behavior of advertisers in these markets. We then show how this framework may be used to provide sharp prescriptions for key auction design decisions that publishers face in these markets. In particular, we show that ignoring budgets, a common practice in this literature, can result in significant profit losses for the publisher when setting the reserve price.

Keywords. auction design, revenue management, ad exchange, display advertising, internet, budget constraints, dynamic games, mean field, fluid approximation.

1 Introduction

The market for display ads on the internet, consisting of graphical content such as banners and videos on web pages, has grown significantly in the last decade, generating about 11 billion dollars in the United States in 2011 (Internet Advertising Bureau, 2012). This growth has been accompanied by the emergence of alternative channels for the purchase of display ads. While traditionally, advertisers would purchase display ad placements by negotiating long term contracts directly with publishers (webpage owners), spot markets for ad slots, called Ad Exchanges, have emerged and the ad spending through these continues to grow (Wall Street Journal, 2013). Google's DoubleClick, OpenX, and Yahoo!'s Right Media are examples of such exchanges.

An Ad Exchange is a platform that operates as an intermediary between online publishers and advertisers. When a user visits a web page (e.g., the New York Times), the publisher may post an ad slot in the exchange together with potentially some user information known to her; e.g., the user's geographical location and her *cookies*. Based on the latter, and in conjunction with their targeting criteria, interested advertisers (or bidders) post bids. Then, an auction is run to determine the winner and the ad to be shown to the user. The latter process happens in milliseconds, between the time a user requests a page and the time the page is displayed to her. As viewers visit her web-site, the publisher repeatedly offers slots to display advertisements; typically, a given publisher runs millions of these auctions per day. On their part, advertisers participate in the exchange with the objective of fulfilling marketing campaigns. In practice, such campaigns are commonly based on a pre-determined budget and extend for a fixed amount of time over which advertisers participate in a large volume of auctions. Given the large number opportunities and the time scale on which decisions are made, bidding is fully automated. See Muthukrishnan (2009) for a more detailed description of Ad Exchanges.

The prevalence of advertisers' budget constraints in these markets links the different auctions over time, and therefore, traditional equilibrium and revenue optimization analysis for static auctions do not apply in this setting. Thus motivated, this paper introduces a new approach to study the key auction design decisions that publishers face, while considering the strategic response of *budget-constrained* advertisers. In particular, the framework captures some key characteristics of an exchange, and allows to start quantifying some central trade-offs faced by publishers and advertisers in this new channel.

1.1 Main Contributions

Advertisers participate in repeated auctions subject to budget constraints, and therefore they typically require dynamic bidding strategies to optimize the allocation of budget to incoming impressions in order to maximize cumulated profits over the length of the campaign. In many cases, advertisers may have similar targeting criteria and bid for the same inventory of ads. Thus, the dynamic bidding strategy an advertiser adopts impacts the competitive landscape for other advertisers in the market. Moreover, the publisher's auction design decisions, such as the reserve price, also impact these interactions. Thus motivated, we formulate our Ad Exchange model as a game among advertisers and the publisher.[1] First, the publisher defines the parameters of a second-price auction that become common knowledge.

[1] In practice, Ad Exchanges may be operated by third-parties; for simplification, in this paper we assume that the publisher and the party running the exchange constitute a single entity.

Then, given the auction format, advertisers compete in a *dynamic* game. In order to quantify the impact of auction design parameters, the first question pertains to the competitive landscape that emerges for fixed auction decisions.

An important challenge in our analysis is solving for the equilibrium of the dynamic game among advertisers induced by the auction rules. At one extreme of agent sophistication, one may consider traditional game theoretic notions of equilibrium such as Perfect Bayesian Equilibrium (PBE), in which advertisers maintain priors on the states of all other bidders, and update them accordingly using Bayes' rule. Such an approach presents two main drawbacks. First, the analysis of the resulting game is, in most cases, intractable from both analytical and computational stand-points. Second, such sophistication and informational requirements on the part of agents is highly unrealistic.

Fluid Mean Field Equilibrium. The main contribution of this paper is the introduction of an equilibrium notion that is tractable and provides a good approximation to the strategic interactions among budget-constrained bidders in an Ad Exchange. Our notion of equilibrium combines in a novel way two different approximations to address the limitations in PBE. First, we consider a *Mean Field* approximation to relax the informational requirements of agents. The motivation behind the latter is that, when the number of players is large, there is little value in tracking the specific actions of all agents and one may rely on some aggregate and stationary representation of the competitors' bids. The mean-field approximation assumes that, even when the overall number of advertisers in the market is large, only a small fraction of them participates in every auction, which closely reflects the existing competitive landscape in today's Ad Exchanges. This type of approximations have appeared in other auction and industrial organization applications (see, e.g., Iyer et al. (2011); Weintraub et al. (2008); Adlakha et al. (2011)). Second, borrowing techniques from the revenue management literature (see, e.g., Gallego and van Ryzin (1994)), we consider a stochastic *fluid* approximation to handle the complex dynamics of the advertisers' control problem. Such approximations are suitable when the number of opportunities is large and the payment per opportunity is small compared to the budget; hence, these are well motivated in the context of Ad Exchanges (see, e.g, Netmining (2011)).

Using the two approximations above, we define the notion of a *Fluid Mean Field Equilibrium* (FMFE).[2] We show that FMFE provides a good approximation to the rational behavior of agents as markets become large, yielding theoretical support for the use of FMFE as an equilibrium concept in this setting. Moreover, we show through a combination of theoretical and numerical results that the FMFE strategy is typically close to being a best response among a large class of strategies that keep track of all available information in the market, even in small markets with few advertisers (e.g., 5 to 10), providing further practical support to the concept. Specifically, in small markets a bidder may have incentives to overbid and deplete competitors' budgets to decrease competition in the future. We

[2] In Section 1.2 we compare and contrast FMFE with related notions of equilibria introduced in previous work.

show, however, that the incentives to exercise such strategic behavior are low relative to playing FMFE even in small markets.

From a structural perspective, when a second-price auction is conducted, quite remarkably, the resulting FMFE strategy has a simple, yet appealing, form: an advertiser needs to shade her values by a *constant factor.* Furthermore, in equilibrium, advertisers will deplete their budget at an essentially constant rate, a typical practical requirement known as "smooth budget depletion". Intuitively, when budgets are tight, advertisers shade their bids because there is an option value for future good opportunities. In addition, we show that an FMFE always exists and provide a set of sufficient conditions that guarantee its uniqueness. We also provide a characterization for FMFE that suggests a simple and efficient algorithm for its computation. Lastly, we derive a closed-form characterization of FMFE in the case of homogeneous bidders. These succinct characterizations of equilibria are remarkably rare and one may significantly leverage them when studying the publisher's problem.

Auction Design. We show how a publisher that maximizes expected profits can use FMFE as a tool for backtesting different auction designs, while accounting for the strategic response of budget-constrained advertisers. In particular, we focus on optimally setting the reserve price. When solving her optimization problem the publisher trades-off the revenues extracted from the auction with the opportunity cost of selling the impressions through an alternative channel. In addition, she needs to consider that changing the auction parameters may change the FMFE strategies played by advertisers. In particular, we show through a combination of theoretical and numerical results that ignoring budgets typically results in reserve prices that are lower than optimal, and may result in significant profit losses for the publisher. We believe these results are particularly relevant, because budgets are typically ignored in the literature when setting optimal reserve prices in the Ad Exchange, despite their prevalence in practice (see the related literature below). We further highlight that other levers may be optimized through the proposed framework, such as the allocation of impressions to the exchange or the extent of user information to disclose to the advertisers.

Overall, the paper is the first in the literature (with the exception of Gummadi et al. (2012) that we discuss below) that provides a framework for profit optimization in repeated auctions, considering the strategic response of budget-constrained bidders. As such, we believe this work can have a practical impact on the design of Ad Exchange auctions. More broadly, we expect that FMFE may have additional applications beyond the one presented in this paper.

1.2 Related Work

This work contributes to various streams of literature. By accounting for advertisers' budget constraints and the resulting inter-temporal dependencies and dynamic bidding strategies they induce, we contribute to the internet advertising literature in particular, and more generally, to the literature on auction design in dynamic settings. To gain tractability, some papers have also used mean-field approximations in these settings. In this vein, Iyer et al. (2011) study repeated auctions in which bidders learn about their private value over time. Our mean-field approximation builds on theirs. However, in our setting dynamics are driven by budget constraints as opposed to learning, resulting in a different model. Moreover, in addition to the mean-field approximation, we impose a fluid approximation to the bidders' control problem. Relative to Iyer et al. (2011), this yields a more succinct characterization of equilibria based on shading factors that (1) brings computational advantages; (2) provides closed-form solutions in some settings even for the optimal auction decisions; and (3) allows using univalence theorems to provide broad sets of conditions under which FMFE is unique. Furthermore, for special cases, we provide approximation results under a sharper scaling, more in line with the typical scales observed in practice. In summary, the *combination* of the two approximations yields results that are extremely rare in the analysis of dynamic games, even after imposing (only) a mean-field approximation.

Closest to our paper is the study of Gummadi et al. (2012) that, in simultaneous and independent work, also study budget-constrained bidders in repeated auctions and introduce a similar equilibrium concept to FMFE. However, the studies differ along many important dimensions. Gummadi et al. (2012) study a more general class of online budgeting problems in an asymptotic regime in which the spending per interaction is small relative to the budget; repeated second-price auctions with budget constraints is a particular case of this general formulation. The present paper, in contrast, focuses on second-price auctions and provides the following sharper results for them that are not present in Gummadi et al. (2012). First, we rigorously justify FMFE as a solution concept through an asymptotic result for large markets and numerical results for finite markets, and provide sufficient conditions for uniqueness of FMFE. Furthermore, we also study various auction design decisions of the publisher, providing important insights on, e.g., reserve price optimization.

More broadly, our work contributes to the growing literature on display advertising, and in particular on that with Ad Exchanges. From the publisher's perspective, various studies analyze display ad allocation with both guaranteed contracts and spot markets. See, e.g., McAfee et al. (2009), Balseiro et al. (2011), Yang et al. (2010) and Alaei et al. (2009). These papers, however, take the actions of the advertisers as exogenous in the auction design. Chen (2011) employs a mechanism design approach to characterize the optimal dynamic auction for the publisher in the presence of guaranteed contract constraints. In this work, however, the publisher faces short-lived advertisers and budget constraints are also ignored. Vulcano et al. (2002) considers a related problem in the context of a single-leg revenue management problem. Celis et al. (2011) introduce a new randomized auction mechanism that experimentally performs better than an optimized second-price auction in markets that become thin due to targeting. They consider, however, a one shot auction and do not take into account the dynamics introduced by budget constraints. Arnon and Mansour (2011) consider an abstraction of a repeated budgeted second price auction in which the repeated interactions are collapsed into a single-shot auction with multiple identical copies of the same item, and study the pure Nash equilibrium of this game. They do not provide, however, a rigorous justification of the solution concept. From the advertiser's perspective, Ghosh et al. (2009) study the design of a bidding agent for a campaign in the presence of an exogenous market.

There is some body of literature on display advertising from a revenue management angle that focuses exclusively on guaranteed contracts (see, e.g., Araman and Fridgeirsdottir (2011), Fridgeirsdottir and Najafi (2010), Roels and Fridgeirsdottir (2009), and Turner (2012)). In the related area of TV broadcasting, Araman and Popescu (2010) study the allocation of advertising space between forward contracts and the spot market when the planner faces supply uncertainty. Our work also relates from a methodological standpoint to some stream of work in revenue management. The single agent fluid approximation we use and some of the intuition underlying it is related to that of, e.g., Gallego and van Ryzin (1994). Building on the latter, Gallego and Hu (2011) focusing on price competition, use a related notion of fluid, or open-loop, equilibrium. Other papers studying dynamic games in revenue management (all focusing on price competition) include Farias et al. (2011), de Albéniz and Talluri (2011), and Dudey (1992).

Our work is related to various streams of literature in auctions. First, previous work has studied auctions with financially constrained bidders in static one-shot settings (see, e.g, Laffont and Robert (1996), Che and Gale (1998), Che and Gale (2000), Maskin (2000), and Pai and Vohra (2011)). In §5 we show that in our dynamic model we obtain drastically different results to some of the main results in that stream. In addition, while our focus is on the impact of budget constraints on second price auctions, our work is somewhat related to the recent literature in optimal dynamic mechanism design (see Bergemann and Said (2010) for a survey). Finally, our work relates to previous papers in repeated auctions, such as Jofre-Bonet and Pesendorfer (2003), in which similarly to our model, bidders shade their bids to incorporate the option value of future auctions. However, in contrast to our work, the latter paper assumes Markov perfect equilibrium behavior in an empirical setting.

2 Model Description

We study a continuous-time infinite horizon setting in which users arrive to an online publisher's web-page according to a Poisson process $\{N(t)\}_{t\geq 0}$ with intensity η. We index the sequence of arriving users by $n \geq 1$, and we denote the sequence of arrival times by $\{t_n\}_{n\geq 1}$. When a user requests the web-page, the publisher may display one advertisement; an event referred to as an *impression*. The publisher may decide to send the impression to an *Ad Exchange*, where an auction among potentially interested advertisers is run to decide which ad to show to the user. The exchange determines the winning bid via a second-price auction with a reserve price, and returns a payment to the publisher. The rules of the auction and the characteristics of the users' arrival process are common knowledge.

Advertisers. Advertisers arrive to the exchange according to a Poisson process $\{K(t)\}_{t\geq 0}$ with intensity λ. We index the sequence of arriving advertisers by $k \geq 1$, and denote by $\{\tau_k\}_{k\geq 1}$ the arrival times.[3]

Advertiser k is characterized by a type vector $\theta_k = (b_k, s_k, \alpha_k, \gamma_k) \in \mathbb{R}^4$. The first component of the type, b_k, denotes the budget and the second component, s_k, denotes the campaign length. That is, the k^{th} advertiser's campaign takes place over the time horizon $[\tau_k, \tau_k + s_k)$ and her total expenditure cannot exceed b_k. Once the advertiser leaves the exchange she never comes back.

When the publisher contacts the exchange she submits some partial information about the user visiting the website, that for example, could include *cookies*. This information, in turn, may heterogeneously affect the targeting criteria and the value an advertiser perceives for the impression, which are captured by α_k and γ_k as we now explain. When the n^{th} user arrives, the advertisers in the exchange observe the user information disclosed by the publisher, and determine whether they will participate or not in the auction based on their targeting criteria. We assume that the k^{th} advertiser matches a user with probability α_k independently and at random (both across impressions and advertisers). Conditional on a match, advertisers have independent private values for an impression. In particular, all values for advertiser k are independent and identically distributed random variables with a continuous cumulative distribution $F_v(\cdot; \gamma_k)$, parameterized by $\gamma_k \in \mathbb{R}$. The distributions have compact support $[\underline{V}, \bar{V}] \subset \mathbb{R}_+$ and continuously differentiable density.[4]

At the moment of arrival, an advertiser's type is drawn independently from a common knowledge distribution with support $\mathbf{\Theta}$, a finite subset of the strictly positive orthant $\mathbb{R}^4_{++}$. This distribution characterizes the heterogeneity among advertisers in the market.

Advertisers have a quasilinear utility function given by the difference between the sum of the valuations generated by the impressions won minus the expenditures corresponding to the second price rule over all auctions they participate during the length of their campaign. The objective of each advertiser is to maximize her expected utility subject to her budget constraint.

Publisher. On the supply side, the publisher has an opportunity cost for selling her inventory of impressions in the exchange; that is, the publisher obtains some fixed amount $c > 0$ for each impression not won by some advertiser in the exchange. The publisher's payoff is given by the long-run average

[3] We note that our approach does not rely in the assumption of Poisson arrivals. In fact, our framework is general and, as shown in Section 6, it also applies for example to the case of synchronous campaigns, when all campaigns start and end at the same time (e.g., weekly or monthly campaigns). In reality, arrivals may lie in a middle ground with a combination of some campaigns repeating over time through a regular schedule, a random inflow of new advertisers (launching, for example, a one-off campaign) and exits of existing advertisers. Our framework could be applied to this setting as well.

[4] By assuming private values, we will ignore the effects of adverse selection and cherry-picking in common value auctions when some advertisers have superior information. See Levin and Milgrom (2010) and Abraham et al. (2012) for work that discusses and analyzes this setting.

profit rate generated by the auctions, where the profit is measured as the difference between the payment from the auction and the lost opportunity cost c when the impression is won by an advertiser in the exchange. The publisher's objective is to maximize its payoff by adjusting the reserve price r to set for the auctions.

Notation. Given a random variable X, we denote a realization x with lower case, its sample space $\mathbf{X}$ with bold capitals, the cumulative distribution function by $F_x(\cdot)$, and the law by $\mathbb{P}_x\{\cdot\}$.

Note. Due to space considerations only selected proofs are presented in the main appendix. All other proofs are presented in a supplementary appendix.

3 Equilibrium Concept

Given the auction design decisions of the publisher, the advertisers participate in a game of incomplete information. Moreover, because the budget constraints couple advertisers' decisions across periods, the game is dynamic and does not reduce to a sequence of static auctions.

A standard solution concept used for dynamic games of incomplete information is that of weak perfect Bayesian equilibrium (WPBE) (Mas-Colell et al., 1995). Roughly speaking, in such a game, a pure *strategy* for advertiser k is a mapping from histories to bids, where the histories represent past observations. A strategy specifies, given a history and assuming the advertiser participates in an auction at time t, an amount to bid. A strategy profile in conjunction with a belief system constitute a WPBE if the following holds. First, given a belief system and the competitors' strategies, an advertiser's bidding strategy maximizes expected future payoffs. Second, beliefs must be consistent with the equilibrium strategies and Bayes' rule whenever possible.

WPBE and commonly used refinements, such as perfect Bayesian equilibrium and sequential equilibrium, require advertisers to hold beliefs about the entire future dynamics of the market, including the future market states. With more than few competitors in the market this imposes a very strong rationality assumption over advertisers as these belief distributions are high-dimensional. Moreover, to find a best response, advertisers need to solve a dynamic programming problem that optimizes over history-dependent strategies. This optimization problem can be high-dimensional and intractable both analytically and computationally. Hence, solving for WPBE for most markets of interest is not possible. More importantly, WPBE imposes informational requirements and a level of sophistication on the part of agents that seems highly unrealistic. This motivates the introduction of alternative equilibrium concepts. After some background in §3.1, we introduce such an alternative in §3.2.

3.1 Mean Field and Fluid Approximation

When selecting an amount to bid, an advertiser needs to form some expectation of the distribution of bids she will compete against. There are various possible bases for such an expectation as a function of the sophistication of the advertiser and the type of information she would have access to. In practice, the number of advertisers in an exchange is often large, in the order of hundreds or even thousands. The first approximation we make is based on the premise that, given a large number of advertisers in the market, the distribution of competitors' bids is stationary and that these random quantities are uncorrelated among periods. Moreover, the bids of any particular advertiser do not affect this distribution. It is common that in these markets, auctioneers provide a "bid landscape" based on aggregated historical data that inherently assumes stationarity, at least for some significant time horizon. This information is commonly used by advertisers to set their bids, and therefore, our assumption about the distribution of competitors' bids may naturally arise in practice (Ghosh et al., 2009; Iyer et al., 2011). In the present paper, while our approximation is predicated on the overall number of advertisers in the market being large, the average number of bidders *per auction* need not be large. For this reason running auctions remains useful in this regime; a small number of bidders with heterogeneous valuations participate in each one of them.

To win an auction, an advertiser competes against all other bidders and against the reserve price r. We denote by D the steady-state maximum of the "competitors' bids", where we assume the publisher is one more competitor that submits a bid equal to r. Assume for a moment that D is i.i.d. across different auctions and distributed according to a c.d.f $F_d(\cdot)$. (Note that $F_d(\cdot)$ will be endogenously determined in equilibrium in §3.2).

In this setting, the advertiser's dynamic bidding problem in the repeated auctions can be casted as a revenue management-type stochastic dynamic programming problem, in which bidding decisions across periods are coupled through the budget constraint. However, the resulting Hamilton-Jacobi-Bellman is a partial differential equation that, in general, does not have a closed-form solution. To get a better handle on the bidder's dynamic optimization problem we introduce a second level of approximation motivated by the fact that a given advertiser has a very large number of bidding opportunities (campaigns span for weeks or months, and thousands of impressions arrive per day). In such an environment, the advertiser's stochastic dynamic programming problem can be well approximated through a stochastic fluid model. In particular, the approximation we focus on is predicated on assuming that bidders solve a control problem in which the budget constraint need only be satisfied in expectation. Under the latter assumption, it is possible to show that one can restrict attention to stationary bidding strategies that ignore the individual state and are only dependent on the actual realization of the bidder's value without loss of optimality. We emphasize here that the budget constraint is imposed almost surely when we conduct performance analysis in §6. The main point is that the stationary bidding strategies derived above can be shown to provide advertisers with provably good policies in the real system (with constraints imposed a.s.) when both the number of impressions and budgets are large, so the number of bidding opportunities over the campaign length also grows large.

Now, the control problem, for a bidder with type $\theta = (b, s, \alpha, \gamma)$ is one of finding a *fluid-based bidding strategy* $\beta_\theta^{\mathrm{F}}(v; F_d)$ that bids depending solely on her value v for the impression. A bidder with total campaign length s observes, in expectation, a total number of $\alpha\eta s$ impressions during her stay in the exchange. By conditioning on the impressions' arrival process, and using our assumption of the stationarity of the maximum bids and the valuations, the bidder's optimization problem is given by

$$J_\theta^{\mathrm{F}}(F_d) = \max_{w(\cdot)} \ \alpha\eta s \mathbb{E}\Big[\mathbf{1}\{D \le w(V)\}(V - D)\Big] \tag{1a}$$

$$\text{s.t.} \ \alpha\eta s \mathbb{E}\Big[\mathbf{1}\{D \le w(V)\} D\Big] \le b, \tag{1b}$$

where the expectation is taken over both the maximum bids D and the valuations V, which are independently distributed according to $F_d(\cdot)$ and $F_v(\cdot; \gamma)$ respectively. Note that the payments in the bidders' problem are consistent with a second-price rule. The bidder optimizes over a bidding strategy that maps its own valuation to a bid; hence, the resulting problem is an infinite-dimensional optimization problem. The next result provides, however, a succinct characterization of an optimal fluid-based bidding strategy.

Proposition 3.1. *Suppose that* $\mathbb{E}[D] < \infty$. *Let* μ_θ^* *be an optimal solution of the dual problem* $\inf_{\mu \ge 0} \Psi_\theta(\mu; F_d)$ *with* $\Psi_\theta(\mu; F_d) = \alpha\eta s \mathbb{E}\Big[V - (1+\mu) D\Big]^+ + \mu b$. *Then, an optimal bidding strategy that solves* (1) *for type* θ *is given by*

$$\beta_\theta^{\mathrm{F}}(v; F_d) = \frac{v}{1 + \mu_\theta^*}.$$

The optimal bidding strategy has a simple form: an advertiser of type θ needs to *shade* her values by the *constant factor* $1 + \mu_\theta^*$, and this factor guarantees that the advertiser's expenditure does not exceed the budget. In the previous expression, μ_θ^* is the optimal dual multiplier of the budget constraint and gives the marginal utility in the advertiser's campaign of one extra unit of budget. Intuitively, when budgets are tight, advertisers shade their bids, because there is an option value for future good opportunities. When budgets are not tight, the optimal dual multiplier is equal to zero and advertisers bid truthfully as in a static second price auction. The proof of the result relies on an analysis of the dual of problem (1). While the latter is not a convex program, the proof establishes from first principles that no duality gap exists in the present case.

3.2 Fluid Mean Field Equilibrium

We now define the dynamics of the market as a prelude to introducing the equilibrium concept we focus on. At any point in time there can be an arbitrary number of advertisers in the exchange, and these dynamics are governed by the patterns of arrivals and departures. In particular, the number of advertisers in the exchange behaves as an $M/G/\infty$ queue. We denote by $\mathcal{Q}(t)$ the set of indices of the advertisers in the exchange at time t, and by $Q(t) = |\mathcal{Q}(t)|$ the total number of advertisers in the system. The *market state* at time t is given by the set of bidders in the exchange, together with their individual states and types, $\Omega(t) = \{\mathcal{Q}(t), \{b_k(t), s_k(t), \theta_k\}_{k \in \mathcal{Q}(t)}\}$, where we denote by $b_k(t)$ and $s_k(t)$ the k^{th} advertiser remaining budget and residual time in the market by time t. When advertisers implement fluid-based strategies the market state encodes all the information relevant to determine the evolution of the market, and the process $\Omega = \{\Omega(t)\}_{t \geq 0}$ is Markov.

In our equilibrium concept we will require the consistency of the distribution of the maximum bid that bidders conjecture they compete against with the bidding strategies they use. A difficulty with this consistency check is that the number of *active* bidders, those that match the target criteria and have remaining budgets, depends on the market dynamics. In particular, the budget dynamics depend on who wins and how much the winner pays in each auction. Hence, in principle, characterizing the resulting steady-state distribution of the maximum bid of active competitors' (that have remaining budgets) is complex. However, it is reasonable to conjecture that, when the number of opportunities during the campaign length is large, rational advertisers would deplete their budgets close to the end of their campaign with high probability. For analytical tractability we impose that, in our equilibrium concept, any bidder currently in the exchange that matches the targeting criteria, without regard of her budget, gets to bid. Under this assumption, the number of bidders in an auction equals the number of advertisers matching the targeting criteria, denoted by $M(t)$, which is just an independent sampling from the process $Q(t)$.[5] In the proof of Theorem 6.1 and in the technical report Balseiro et al. (2012) we show this layer of approximation is in fact asymptotically correct. Indeed, the performance analysis in Section 6 takes into account that when advertisers implement the FMFE strategies stochastic fluctuations in their expenditure may actually induce them to run out of budget before the end of the campaign, at which point they cannot continue to participate in any auction.

Since arrival and departures of advertisers are governed by an $M/G/\infty$ queue and campaign lengths are bounded, it is not hard to show that under fluid-based strategies the market process Ω is Harris recurrent, so it is ergodic and admits a unique invariant steady-state distribution (see, e.g., Asmussen (2003) p. 203). Let $\left(\hat{M}, \{\hat{\Theta}_k\}_{k=1}^{\hat{M}}\right)$ be a random vector that describes the number of matching bidders,

[5] We note that an important difference between our FMFE and the related equilibrium concept proposed in parallel by Gummadi et al. (2012) is that they do not impose this additional layer of approximation. This plays a key role to obtain tractability in our analysis.

together with their respective types when sampling a market state according to the invariant distribution. Notice that advertisers with longer campaign lengths and higher matching probability are more likely to participate in an auction, and thus the law of a type sampled from the invariant distribution does not coincide with the law of the types in the population. Indeed, by exploiting the fact that arrival-time and service-time pairs constitute a Poisson random measure on the plane (see, e.g., Eick et al. (1993)), one can show that $\hat{M}$ is Poisson with parameter $\mathbb{E}[\alpha_\Theta \lambda s_\Theta]$, and that each component of the vector of types is independently and identically distributed as $\mathbb{P}_{\hat{\Theta}}\{\theta\} = \frac{\alpha_\theta s_\theta}{\mathbb{E}[\alpha_\Theta s_\Theta]}\mathbb{P}_\Theta\{\theta\}$ for each type $\theta \in \mathbf{\Theta}$, and independent of $\hat{M}$.[6]

For a fluid-based strategy profile $\boldsymbol{\beta} = \{\beta_\theta(\cdot) : \theta \in \mathbf{\Theta}\}$ with $\beta_\theta : [\underline{V}, \bar{V}] \to \mathbb{R}$, we denote by $F_d(\boldsymbol{\beta})$ the distribution of the following random variable

$$\max\left(\left\{\beta_{\hat{\Theta}_k}(V_{\hat{\Theta}_k})\right\}_{k=1}^{\hat{M}}, r\right), \tag{2}$$

which represents the steady-state maximum bid. Note that here V_θ are independent valuations sampled according to $F_v(\cdot; \gamma_\theta)$. We are now in a position to formally define the notion of a *Fluid Mean Field Equilibrium* (FMFE).

Definition 3.1 (Fluid Mean Field Equilibrium)**.** *A fluid-based strategy profile $\boldsymbol{\beta}$ constitutes a FMFE if for every advertiser's type $\theta \in \mathbf{\Theta}$, the bidding function β_θ is optimal for problem* (1) *given that the distribution of the maximum bid of other advertisers is given by $F_d(\boldsymbol{\beta})$ (equation* (2)*).*

Essentially a FMFE is a set of bidding strategies such that (i) these strategies induce a given competitive landscape as represented by the steady-state distribution of the maximum bid, and (ii) given this landscape, advertisers' optimal fluid-based bidding strategies coincide with the initial ones. We focus on symmetric equilibria in the sense that all bidders of a given type adopt the same strategy. Note that in the FMFE, upon arrival to the system an advertiser is assumed to compete against the market steady-state maximum bid D.[7]

Remarks. We introduced FMFE by heuristically arguing that it should be a sensible equilibrium concept for large markets when the number of bidding opportunities per advertiser are also large. In Theorem 6.1, we show that when all advertisers implement the FMFE strategy, the relative profit increase of any unilateral deviation to a strategy that keeps track of all information available to the advertiser becomes negligible as the scale of the market increases. This provides a theoretical justification for using FMFE as an approximation of advertisers' behavior.

[6]For a type $\theta \in \mathbf{\Theta}$ we denote, with some abuse of notation, the corresponding budget by b_θ, the campaign length by s_θ, the matching probability by α_θ, and the valuation parameter by γ_θ. Additionally, we denote by Θ a random variable distributed according to the law of types in the population.

[7]Note that by the PASTA property of a Poisson arrival process this assumption is in fact correct.

In the asymptotic regime described above the matching probabilities are decreased so that the number of bidders per auction remains constant, and therefore, the probability that two advertisers participate repeatedly in the same auctions becomes negligible. In real-world markets, it might be the case that similar advertisers compete repeatedly in the same auctions to advertise to the same users. Nonetheless, in Section 6.2 we show through a combination of theoretical and numerical results that even with a moderate number of advertisers (e.g., 5 to 10) FMFE strategies are typically close to being a best response. Naturally, in these markets two advertisers may interact repeatedly over time and our results show that FMFE provides a good approximation to the rational behavior of agents even in these cases.

4 Fluid Mean Field Equilibrium Characterization

In this section we prove the existence, provide conditions for uniqueness, and characterize the FMFE. Proposition 3.1 will significantly simplify our analysis, because it allows one to formulate the equilibrium conditions in terms of a vector of multipliers instead of bidding functions. By doing so, the problem of finding the equilibrium strategy function for a given type will be reduced to finding a single multiplier.

4.1 Equilibrium Existence and Sufficient Conditions for Uniqueness

We first prove the existence of a FMFE for a fixed reserve price. Recall from Proposition 3.1 that, in an optimal fluid bidding strategy, advertisers of type θ shade their bids using a fixed multiplier μ_θ. In the following we denote by $\boldsymbol{\mu} = \{\mu_\theta\}_{\theta \in \boldsymbol{\Theta}}$ a vector of multipliers in $\mathbb{R}_+^{|\boldsymbol{\Theta}|}$ for the different advertisers' types. Given a postulated profile of multipliers $\boldsymbol{\mu}$, let $F_d(\boldsymbol{\mu})$ denote the steady-state distribution of the maximum bid and let $\Psi_\theta(\mu; \boldsymbol{\mu}) \triangleq \Psi_\theta(\mu; F_d(\boldsymbol{\mu}))$ be the dual objective for one θ-type advertiser (as defined in Proposition 3.1) when all other bidders adopt a strategy given by the vector $\boldsymbol{\mu}$ (including those of her own type). In the dual formulation, a vector of multipliers $\boldsymbol{\mu}^*$ constitutes a FMFE if and only if it satisfies the best-response condition

$$\mu_\theta^* \in \arg\min_{\mu \geq 0} \Psi_\theta(\mu; \boldsymbol{\mu}^*), \qquad \text{for all types } \theta \in \Theta. \tag{3}$$

One may establish that the system of equations (3) always admits a solution to obtain the following.

Theorem 4.1. *There always exists an FMFE.*

The proof shows that the dual strategy space can be reduced to a compact set, and that the dual objective function is jointly continuous in its arguments, and convex in the first argument. Then, a standard result that relies on Kakutani's Fixed-Point Theorem implies existence of an FMFE.

We now turn to sufficient conditions for uniqueness. Let $\mathbf{G} : \mathbb{R}_+^{|\mathbf{\Theta}|} \times \mathbb{R}_+ \to \mathbb{R}_+^{|\mathbf{\Theta}|}$ be a vector-valued function that maps a profile of multipliers and a reserve price to the steady-state expected expenditures per auction of each type. The expected expenditure of a θ-type bidder in a second-price auction when advertisers implement a profile of multipliers $\boldsymbol{\mu}$ is given by $G_\theta(\boldsymbol{\mu}, r) \triangleq \mathbb{E}\left[\mathbf{1}\{(1+\mu_\theta)D \leq V\}D\right]$, where the maximum competing bid is given by $D = \max\left(\left(V_{\hat{\Theta}} / \left(1+\mu_{\hat{\Theta}}\right)\right)_{1:\hat{M}}, r\right)$.[8]

Assumption 4.1 (P-matrix)**.** *The Jacobian of $-\mathbf{G}$ with respect to $\boldsymbol{\mu}$ is a P-matrix for all $\boldsymbol{\mu}$ in $\mathbb{R}_+^{|\mathbf{\Theta}|}$*

A matrix is P-matrix if all its principal minors are positive (Horn and Johnson, 1991, p.120). Assumption 4.1 can be shown to hold for various cases of interest. For example, it is easy to see that it always holds for the case of homogeneous advertisers, i.e., when the space of types $\mathbf{\Theta}$ is a singleton. In Appendix C, we provide an important class of settings with two types of bidders in which it also holds. The P-matrix condition can be interpreted as a monotonicity condition on the expected expenditures. Namely, if a group of types increases its multipliers simultaneously, then the expenditures cannot increase for every type in the group. The next theorem shows that the equilibrium is unique under the P-matrix assumption.

Theorem 4.2. *Suppose Assumption 4.1 holds. Then, there is a unique FMFE of the form $\beta_\theta(v) = v/(1+\mu_\theta)$, θ in Θ.*

We prove the result by formulating the FMFE conditions as a Non-linear Complementarity Problem (NCP) as presented in Corollary 4.1 below, and employing a Univalence Theorem to show that the expected expenditure mapping is injective (Facchinei and Pang, 2003*a*). We note that results regarding uniqueness of equilibria in dynamic games are extremely rare (Doraszelski and Pakes, 2007).

Providing conditions for which Assumption 4.1 holds is challenging for more than two types of bidders. In our numerical experiments we use a myopic best response algorithm, presented in detail in Section 5.2.1, that could naturally describe how agents learn to play the game and reach an FMFE. It is reassuring that in our computational experience, for a given model instance with two or more types, this algorithm always found the same FMFE even when starting from different initial points.

We finish this subsection by noting that under further mild regularity conditions one can establish that any set of continuous increasing bidding functions that constitute an FMFE necessarily yield the same outcome (in terms of auctions' allocations and payments) as that of the FMFE in Theorem 4.2. In the rest of the paper, we focus on the simple and intuitive FMFE strategies that can be described by a vector of dual multipliers.

[8] Note that consistent with the FMFE assumption and the PASTA property, the bidder competes against the *market* steady-state maximum bid.

4.2 Equilibrium Characterization

A direct corollary of the earlier results and their proofs yields the following succinct characterization.

Corollary 4.1. *Any FMFE characterized by a vector of multipliers $\boldsymbol{\mu}^*$, such that $\beta_\theta(v) = v/(1+\mu_\theta^*)$ for all $v \in [\underline{V}, \bar{V}]$ and $\theta \in \boldsymbol{\Theta}$, solves*

$$\mu_\theta^* \geq 0 \quad \perp \quad \alpha_\theta \eta s_\theta G_\theta(\boldsymbol{\mu}^*, r) \leq b_\theta, \quad \forall \theta \in \Theta,$$

where $\perp$ indicates a complementarity condition between the multiplier and the expenditure, that is, at least one condition should be met with equality.

The expected expenditure for a bidder of type θ over its campaign when bidders use a vector of multipliers $\boldsymbol{\mu}$ is given by $\alpha_\theta \eta s_\theta G_\theta(\boldsymbol{\mu}, r)$, because on average she faces ηs_θ auctions and participates in a fraction α_θ of them. Intuitively, the result states that, in equilibrium, advertisers of a given type may only shade their bids if their total expenditure over the course of the campaign (in expectation) is equal to their budget. If, in expectation, advertisers have excess budget at the end of a campaign, then, their multiplier is equal to zero and they should bid truthfully. This equilibrium characterization lends itself for tractable algorithms to compute FMFE, because the strategy of each advertiser type is determined by *a single number* that satisfies the complementary conditions above. See, for example, Chapter 9 of Facchinei and Pang (2003*b*) for a discussion of numerical algorithms for this kind of NCPs.

We conclude this subsection by refining the result for the case of homogeneous bidders, in which one can provide a quasi-closed form characterization for FMFE. Suppose that Θ is a singleton. Let $G_0(r) = G_\theta(0, r)$ denote the steady-state unconstrained expected expenditure-per-auction of a single bidder for a second price auction with reserve price r when all advertisers (including herself) bid their own values. Note that the expected expenditure for a bidder over its campaign when all bidders are truthful is given by $\alpha \eta s G_0(r)$. This quantity plays a key role in the FMFE characterization.

Proposition 4.1. *Suppose $\boldsymbol{\Theta}$ is a singleton. Then a Fluid Mean Field Equilibrium exists and is unique. In addition, the equilibrium may be characterized as follows: $\beta_\theta(v) = v/(1+\mu^*)$ for all $v \in [\underline{V}, \bar{V}]$, where $\mu^* = 0$ if $\alpha \eta s G_0(r) < b$, and μ^* is the unique solution to $\alpha \eta s G_0(r(1+\mu)) = b(1+\mu)$ if $\alpha \eta s G_0(r) \geq b$.*

The result provides a complete characterization of the unique FMFE. In particular, it states that if budgets are large (i.e., $\alpha \eta s G_0(r) < b$), then in equilibrium advertisers will bid truthfully. If however, budgets are tight (i.e., $\alpha \eta s G_0(r) \geq b$), then advertisers will be shading their bids in equilibrium, considering the option value of future opportunities. We also further note here that in the case in which the reserve price is equal to zero ($r = 0$), the equilibrium multiplier may be characterized in closed form by $\mu^* = (\alpha \eta s G_0(0)/b - 1)^+$.

5 Auction Design: Reserve Price Optimization

In this section we study the publisher's profit maximization problem. First, we use the framework developed in the previous sections to formulate the problem. Then, we study the resulting optimization problem and derive insights on how to account for budgets when setting the reserve price.

We model the grand game played between the publisher and advertisers as a Stackelberg game in which the publisher is the leader and the advertisers are the followers. In particular, the publisher first selects the reserve price in the second-price auction r, and then the advertisers react and play the induced dynamic game among them. In our analysis we assume that the solution concept for the game played between advertisers is that of FMFE. The publisher's objective is to maximize her long run average profit rate from the auctions, while considering the opportunity cost c of the alternative channel.

To mathematically formulate the problem we define $I(\boldsymbol{\mu}, r) = 1 - F_d(r; \boldsymbol{\mu})$ as the probability that the impression is won by some advertiser in the exchange when advertisers shade according to the profile $\boldsymbol{\mu}$ and the publisher sets a reserve price r. Using the characterization of an FMFE in Corollary 4.1, we can write the publisher's problem in terms of multipliers, and obtain the following Mathematical Program with Equilibrium Constraints (MPEC):

$$\begin{aligned} \max_{r} \quad & \lambda \sum_{\theta \in \boldsymbol{\Theta}} p_\theta \alpha_\theta \eta s_\theta G_\theta(\boldsymbol{\mu}, r) - \eta c I(\boldsymbol{\mu}, r) \\ \text{s.t.} \quad & \mu_\theta \geq 0 \quad \perp \quad \alpha_\theta \eta s_\theta G_\theta(\boldsymbol{\mu}, r) \leq b_\theta, \quad \forall \theta \in \Theta, \end{aligned} \tag{4}$$

where $p_\theta \triangleq \mathbb{P}_\Theta\{\theta\}$ is the probability that an arriving advertiser is of type θ. We denote by $\Pi(\boldsymbol{\mu}, r)$ the objective function of the MPEC. The first term in the objective is the publisher's revenue rate obtained from all bidders' types in the auctions, which is equal to the average expenditure of the advertisers. Note that the revenue rate obtained from a given type is equal to the bidders' average expenditure over their campaign times the arrival rate of bidders. The second term is the opportunity cost by unit of time, which is incurred whenever an impression is won by some advertiser in the exchange and, therefore, cannot be sold in the alternative channel.

Note that the MPEC above considers that when the publisher changes the reserve price, bidders react by playing a corresponding FMFE. By Theorem 4.1 we know that such a FMFE always exist. Further, when Assumption 4.1 holds, the FMFE is unique. In cases for which we do not know whether the assumption holds, we will assume that advertisers play the FMFE computed by our best response algorithm.[9]

[9] Assuming that the equilibrium being played is the one selected by a specific algorithm is a prevalent approach in the analysis of dynamic games for which uniqueness results are extremely rare. For example, Iyer et al. (2011) use this approach in a repeated auction setting and many of the references in Doraszelski and Pakes (2007) use it in other

5.1 Reserve Price: Homogeneous Advertisers

We first consider the case in which $\mathbf{\Theta}$ is a singleton, i.e., all advertisers have a fixed budget b, stay in the market for a deterministic time s, and share the same matching probability α and valuation parameter γ. By Proposition 4.1 we know that in this case a unique FMFE exists and we can characterize it in quasi-closed form. We leverage this result to study the publisher's decisions. Throughout this section, we drop the dependence on θ. In the following we denote by $h_v(x) = f_v(x)/\bar{F}_v(x)$ the *failure rate* of the advertisers values (who have a common distribution), and by $\xi_v(x) = xh_v(x)$ the *generalized failure rate* of the values. We assume that values possess strictly increasing generalized failure rates (IGFR). This assumption is common in the pricing and auction theory literature, and many distributions satisfy this condition (see, e.g., Myerson (1981) and Lariviere (2006)).[10]

In the absence of budgets, the auctions are not coupled and each auction is equivalent to a one-shot second-price auction with opportunity cost $c > 0$ and symmetric bidders with private values. In such a setting, it is well-known that the optimal reserve price, which we denote by r_c^*, is independent of the number of bidders and given by the unique solution of $1/h_v(x) = x - c$ (see, e.g., Laffont and Maskin (1980)). The next result establishes a counterpart for the present case with budget constraints.

Theorem 5.1. (Optimal reserve price)**.** *If $\alpha\eta s G_0(r_c^*) < b$, then r_c^* is the unique optimal reserve price. If $\alpha\eta s G_0(r_c^*) \geq b$, then the unique optimal reserve price is $\bar{r} = \sup \mathcal{R}^*$, where $\mathcal{R}^* = \{r : \alpha\eta s G_0(r) \geq b\}$. Furthermore, in the FMFE induced by the optimal reserve price, advertisers bid truthfully.*

The optimal reserve price admits a closed-form expression that highlights how it balances various effects. The expected expenditure for a bidder over its campaign when all bidders are truthful evaluated at r_c^*, $\alpha\eta s G_0(r_c^*)$, plays a key role in the result. In fact, when the budget is large in the sense that advertisers do not deplete their budget in expectation when the reserve price is r_c^* ($\alpha\eta s G_0(r_c^*) \leq b$), then it is expected that r_c^* should still be optimal in our setting. Intuitively, if the budget does not bind, the auctions decouple into independent second price auctions. When, however, $\alpha\eta s G_0(r_c^*) > b$, advertisers shade their values when the reserve price is r_c^*. In the proof, we show that in this case the optimal reserve price must be in $\mathcal{R}^*$, that is, it must induce bidders to deplete their budgets in expectations. For all such reserve prices, the revenue rate for the publisher is given by λb and this is the maximum revenue rate she can extract from advertisers. Hence, recalling the objective value (4) of the publisher, the optimal reserve price must be the value $r \in \mathcal{R}^*$ that minimizes the probability of selling an impression in the exchange, and therefore the opportunity cost. Increasing the reserve price has two effects on this probability: (1) a *direct effect*: assuming advertiser's strategies do not change, an increase of the reserve price decreases the probability of selling an impression in the exchange; and

industrial organization games.

[10]For instance the uniform, exponential, triangular, truncated normal, gamma, Weibull, and log-normal distribution have IGFR.

(2) an *indirect effect*: a change in the reserve price also alters the strategies of the advertisers through the induced FMFE. In the proof we show that the direct effect is dominant, implying that $\bar{r} = \sup \mathcal{R}^*$ is optimal since it minimizes the opportunity cost within $\mathcal{R}^*$.

We emphasize that the optimal reserve price with budget constraints is larger or equal than r_c^*, the static reserve price that does not account for budgets. In fact, the optimal reserve price is $\max\{\bar{r}, r_c^*\}$, because one can show that $\bar{r} \geq r_c^*$ if and only if $\alpha\eta s G_0(r_c^*) \geq b$. Theorem 5.1 highlights that ignoring budgets can result in a suboptimal decision. In the next section, we numerically evaluate the extent of the sub-optimality in markets with heterogeneous bidders.

Before, we note that when advertisers are highly budget-constrained, the reserve price $\bar{r}$ tends to be high, and therefore it is very unlikely that two advertisers will bid higher than $\bar{r}$. In this case, the advantage of running a second price auction becomes limited and its performance is similar to that of a fixed posted price mechanism.

We finish this subsection by comparing the result above with the studies pertaining to one-shot auctions with budget constraints. In the case of a common budget for all bidders, authors have typically found that budget constraints decrease the optimal reserve price relative to the setting without budget constraints (see Laffont and Robert (1996) and Maskin (2000)). The reason is that with budget constraints the reserve price is less effective in extracting rents of higher valuation types; hence, when trading-off higher revenues conditional on a sale taking place with an increase in the probability of a sale, the latter has more weight than in the absence of budgets. In our case, instead, the optimal reserve price with budget constraints is larger or equal than r_c^*. The difference with the one-shot auction is that the budget constraint is imposed over a large set of auctions as opposed to having a constraint per auction, leading to a different trade-off for the publisher. Indeed, when the budget constraint binds, the reserve price does not affect expected revenues, the publisher is already extracting all budgets from the bidders. Therefore, the only role of the reserve price becomes one of reducing the opportunity cost by decreasing the probability of a sale. As we saw, this is achieved by increasing the reserve price while still extracting the maximum amount of revenues.

5.2 Reserve Price: Heterogeneous Advertisers

While it was possible to obtain essentially a closed-form solution for the publisher's optimal reserve price in the case of homogeneous advertisers, it is not in general possible to derive such a result for the case of heterogeneous advertisers. However, one may always numerically analyze the impact of the publisher's decisions on the advertisers' equilibrium outcome under different scenarios by solving for the FMFE using the characterization in Corollary 4.1 for different auction parameters. We provide such a study in this section and start by describing an algorithm to compute FMFE.

5.2.1 Algorithm to Compute FMFE

For each model instance, we solve for FMFE using the following myopic best response algorithm over the space of dual multipliers. The algorithm starts from an arbitrary vector of multipliers $\boldsymbol{\mu}$:

Algorithm 1 Best Response Algorithm for FMFE

$\mu_\theta^0 := \mu_\theta,\ \forall \theta \in \Theta;\ \ i := 0$
repeat
  $\mu_\theta^{i+1} := \arg\min_{\mu' \geq 0} \Psi_\theta(\mu'; \boldsymbol{\mu}^i),\ \forall \theta \in \Theta$
  $\Delta := \|\boldsymbol{\mu}^{i+1} - \boldsymbol{\mu}^i\|_\infty;\ \ i := i + 1$
until $\Delta < \epsilon$

If the termination condition is satisfied with $\epsilon = 0$, we have a FMFE (see equation (3)). Small values of ϵ allow for small errors associated with limitations of numerical precision. While we cannot prove convergence of the algorithm, in practice, it converged in a small number of iterations. In fact, for fixed auction parameters, solving for the FMFE takes a few seconds on a laptop computer.

5.2.2 Measuring the Impact of Budgets on the Optimal Reserve Price

The analysis with homogeneous bidders highlighted that ignoring budgets can lead to suboptimal reserve prices. In this section, we measure the extent of the sub-optimality in markets with heterogeneous bidders. We believe this exercise is particularly relevant, because several papers that study online advertising in fact ignore budgets when setting optimal reserve prices in the Ad Exchange (see, e.g., McAfee et al. (2009), Balseiro et al. (2011), Chen (2011), and Celis et al. (2011)).

The setup for our numerical experiments is as follows. We consider randomly generated instances with a heterogeneous population of advertisers with five types. Budgets for each type are sampled from a discrete uniform distribution with support $\{1, 2, ..., 10\}$. Additionally, we experiment with the proportion of these types by choosing the probabilities p_θ of an arriving advertiser being of type θ uniformly from the probability simplex. Throughout the experiments we fix the matching probability $\alpha = 0.1$ and the campaign length to $s = 10$, but select the arrival rate λ uniformly in $[1, 5]$ so that the average number of matching bidders in an auction $\alpha\lambda s$ varies from 1 to 5. Advertisers have the same distribution of values, which is drawn uniformly from the set $\text{Exp}(\gamma)$, $\mathcal{N}(\gamma, 1)$, and $\text{Unif}[0, 2\gamma]$ with γ uniformly sampled from $[1, 5]$ (the supports of valuations are truncated to $[0, 10]$). From the perspective of the publisher, we study scenarios with different opportunity costs c for the alternative channel, by choosing the cost uniformly from $[1, 5]$. Additionally, we consider 10 levels for the impressions allocated to the exchange, as given by η.[11] In total, we consider 920 different scenarios.

[11] In particular, we consider 10 uniformly spaced points in the interval $[0, 1.25 \max_\theta \bar{\eta}_\theta]$ where $\bar{\eta}_\theta$ is the least rate of impressions guaranteeing that a population of type θ bidders in isolation is budget constrained when the reserve is r_c^*.

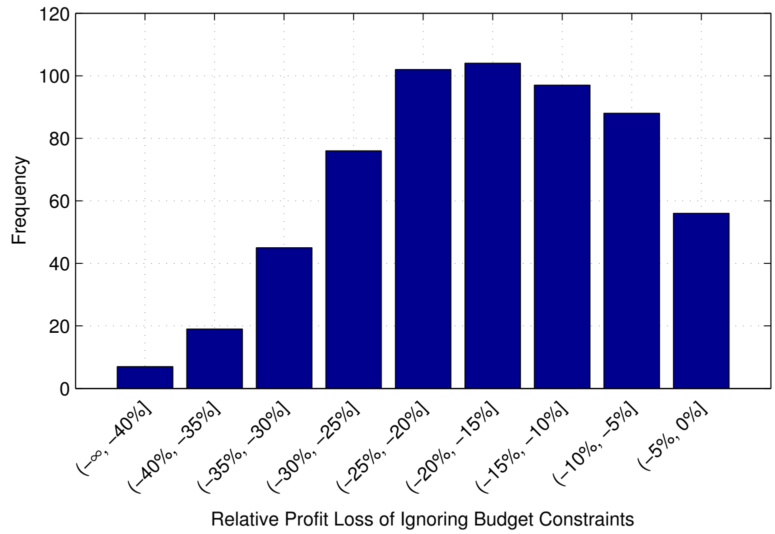

Figure 1: Histogram of the relative profit loss of ignoring budget constraints for randomly generated instances. The relative profit loss is given by $\Pi(\boldsymbol{\mu}(r_c^*), r_c^*)/\Pi(\boldsymbol{\mu}(r^*), r^*) - 1$ where $\boldsymbol{\mu}(r)$ denotes the FMFE multipliers at reserve price r. The histogram is restricted to those instances in which the mean advertiser truthful expenditure at r_c^* exceeds the mean budget, i.e., $\sum_\theta p_\theta G_\theta(0, r_c^*) \geq \sum_\theta p_\theta b_\theta$.

For each model instance we compute two reserve prices. First, the optimal static reserve price r_c^* as given in Section 5.1, which assumes advertisers always bid truthfully, and therefore, ignores budget constraints. Second, the reserve price r^* that solves optimization problem (4), and therefore, considers the rational response of budget-constrained advertisers via FMFE.

From the numerical experiments we obtain two conclusions that are robust across all model instances. First, consistent with the results from the homogenous case, the reserve price r^* is larger than r_c^*. Second, the extent of sub-optimality associated with ignoring budgets and selecting r_c^* instead of r^* can be significant with profit losses up to 40%. A histogram of the relative profit loss across the generated instances is shown in Figure 1. Overall, our results show that ignoring the rational response of budget-constrained advertisers can yield significant profit losses for the publisher.

5.2.3 Structure of the Optimal Reserve Price

In this section, we study in more detail the structure of the optimal reserve price r^* in markets with heterogeneous bidders to illustrate the trade-offs the publisher faces in these settings. For this purpose, it is useful to depict the optimal reserve price and the resulting shading multipliers as a function of the allocation of impressions to the exchange η. Figure 2 shows such dependence for a given set of parameters with two types. Notice that when the publisher prices optimally, the high-budget type always bids truthfully. However, in contrast to the homogeneous case, this is not necessarily true for the low-budget type: for some levels of supply, low-type advertisers will shade their bids under the optimal reserve price.

Focusing on the optimal reserve price, we observe that advertisers do not have a chance to deplete their budgets for low levels of supply. In this case, advertisers bid truthfully and r_c^* is the optimal reserve price. As the rate of impressions increases, the expenditures increase up to the point at which the low-type becomes budget constrained. From then on the publisher needs to balance two effects. On the one hand, since the low-type is now shading her bids, the publisher has an incentive to increase the reserve price so as to minimize the number of impressions won and the opportunity cost. The latter is achieved by $\bar{r}_1(\eta)$, the optimal reserve price if all advertisers shared the same budget b_1 (the top dashed line). On the other hand, the publisher has an incentive to price close to r_c^* to extract the surplus from the high-type advertisers, who are not depleting their budgets. The tradeoff is such that, initially, the weight of the low-budget type bidders is higher and it is optimal for the publisher to price close to $\bar{r}_1(\eta)$, and thus increasing the reserve price with the allocation of impressions. At this price, however, the expenditure of the high-budget type is well below its budget, and the publisher may be leaving money on the table. When enough impressions are allocated to the exchange this effect becomes dominant and the publisher tries to extract this surplus by pricing closer to r_c^*; thus the sudden kink and decrease in the optimal reserve price. If the publisher keeps increasing the allocation of impressions, eventually both types become budget constrained. Similarly to the homogeneous case, the publisher is now better off pricing in a way such that both types deplete their budgets, but with the high-type bidding truthfully, so that the number of impressions won by the advertisers is minimized. For this reason, at some point the optimal reserve price starts increasing again.

In our numerical experiments, a similar structure and tradeoff appears when there are more than two types of advertisers with different budgets in the population, with one new kink in the optimal reserve price for each additional type.

6 FMFE as a Near-Optimal Best Response

In this section we aim to provide further support for the concept of FMFE introduced in Section 3.2 along two dimensions. First we rigorously justify that playing an FMFE strategy when all other advertisers play the FMFE strategy is a near-optimal best response in markets of large "size", i.e., when both the number of advertisers and the number of auctions are appropriately large. Second, we aim to illustrate theoretically and numerically the main trade-offs faced by advertisers and why FMFE strategies are potentially near-optimal even when the number of advertisers is small, lending further practical support to the concept.

Preliminaries. To achieve the above goals, we focus on a simplified version of the problem, the case of *synchronous campaigns*, that is, when all campaigns start at the same time and finish simultaneously.

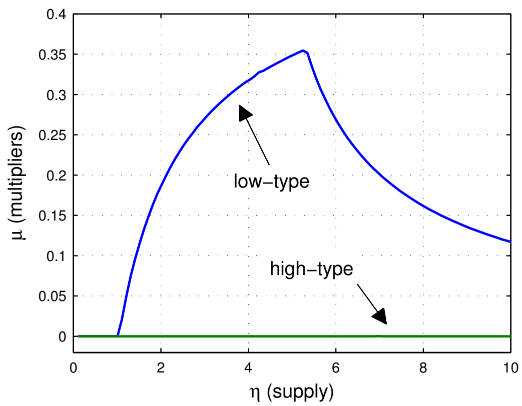

(a) Equilibrium multipliers

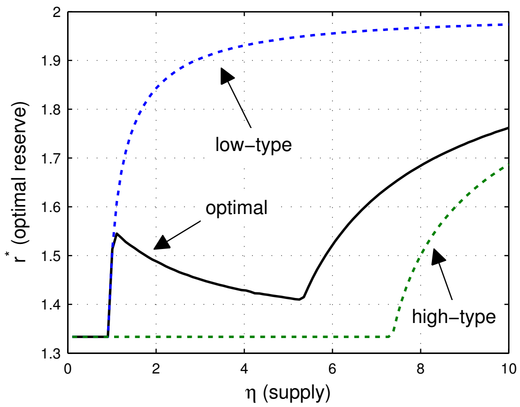

(b) Optimal reserve price

Figure 2: Equilibrium multipliers and optimal reserve price as a function of the rate of impressions for an instance with $\alpha = 0.1$, $\lambda = 1$, $s = 40$, Unif$[0,2]$ valuation distribution, $c = \frac{2}{3}$, $b = (1, 8)$, and $p = (\frac{1}{5}, \frac{4}{5})$. For illustration purposes we only consider two types and different parameters than above. (a) Equilibrium multipliers as a function of the allocation of impressions. (b) The solid line corresponds to optimal reserve price, while the dashed lines denote the optimal prices one would set for a homogeneous population with budget b_1 (low-type) or b_2 (high-type). The reserve price r_c^* is equal to $\frac{4}{3}$.

This simpler model corresponds, for example, to the case when advertisers have periodic (daily or weekly) budgets. It captures some of the key characteristics of the market, and allows to highlight the main issues at play in a relatively transparent fashion. The general case of *asynchronous campaigns* introduces a significant additional layer of complexity, and we provide an asymptotic approximation result pertaining to the latter in Balseiro et al. (2012).[12]

We next describe the synchronous model and adapt FMFE to this setting. There is a fixed number of agents in the market, which we denote by K. All campaigns start at time 0 and finish at a common time s, and neither arrivals nor departures are allowed during the time horizon $[0, s]$. Agents are indexed by $k = 1, \ldots, K$. Similarly as before, the k^{th} agent is characterized by a type vector $\theta_k = (b_k, \alpha_k, \gamma_k) \in \mathbb{R}^3$. Types are publicly known and revealed at the beginning of the horizon. While this assumption is not necessary for our analysis, we make it to simplify some arguments and notation.

Now, the expected expenditure function of the k^{th} advertiser of a single auction when advertisers shade their bids according to a vector of multipliers $\boldsymbol{\mu} \in \mathbb{R}_+^K$, denoted by $G_k(\boldsymbol{\mu}; r)$, is given as in §4 but with the maximum competing bid given by $D_{-k} = \max_{i \neq k, M_i = 1} \left\{ V_i/(1 + \mu_i) \right\} \vee r$, where we let

[12]Due to the asynchronous nature of the market, for this result we extend the propagation of chaos argument of Graham and Méléard (1994) and Iyer et al. (2011) to accommodate the additional fluid approximation and the queuing dynamics of the number of advertisers in the market, which leads to a more restrictive scaling than our result below for synchronous campaigns. An interesting technical avenue for future research is to show whether the scaling under which we obtain our asymptotic approximation result for synchronous campaigns holds in broader settings. This generalization is likely to have other applications in mean-field models beyond the one presented in this paper.

$M_k = 1$ indicate that the k^{th} agent participates in the auction and we ignored the index n to simplify the notation. A similar analysis to the one performed in the case of asynchronous campaigns yields that the vector of multipliers in the FMFE can be characterized as the solution of the following NCP:

$$\mu_k \geq 0 \quad \perp \quad \alpha_k \eta s G_k(\boldsymbol{\mu}; r) \leq b_k, \quad \forall k = 1, \ldots, K. \tag{5}$$

Moreover, similar results about the existence and uniqueness of FMFE also apply to this setting.

6.1 Asymptotic Analysis for Large Markets

We consider a sequence of markets indexed by the number of advertisers K. For each market size K, bidders' types are given by: $\left\{\theta_k^{(K)} = \left(b_k^{(K)}, \alpha_k^{(K)}, \gamma_k^{(K)}\right)\right\}_{k=1}^{K}$, where we use superscript $^{(K)}$ to denote quantities associated to market size K. Similarly, we denote by $\eta^{(K)}$ as the intensity of the arrival process of users in market K. We will prove an approximation result by considering a sequence of markets that satisfy the following set of assumptions on the primitives.

Assumption 6.1. *There exists positive bounded constants $\underline{g}, \bar{g}, \underline{z}, \bar{a}$, such that for all market sizes K:*

i.) For any advertiser k, $b_k^{(K)}/(\alpha_k^{(K)}\eta^{(K)}s) \in [\underline{g}, \bar{g}]$.

ii.) For every pair of advertisers $k \neq i$, $\alpha_k^{(K)}/\alpha_i^{(K)} \leq \bar{a}$.

iii.) For any advertiser k, $G_k^{(K)}(\mathbf{0}; r) \geq \underline{z}$.

The first assumption states that the ratio of budget to number of matching auctions is uniformly bounded from above and below across advertisers and the second one that the ratio of matching probabilities of any two advertisers is uniformly bounded across advertisers. These assumptions guarantee that no advertiser has an excessive market influence by limiting budgets and the number of matching auctions in which they participate. The third assumption ensures that, in equilibrium, all advertisers have a positive expected expenditure per auction so that no advertiser is systematically outbid in equilibrium. Thus, these assumptions simply guarantee that for every market along the sequence considered, there is no dominant or irrelevant advertiser. These assumptions do not impose further heterogeneity restrictions across advertisers.

We denote the k^{th} advertiser *history* up to time t by $h_k(t)$. The history encapsulates all available information up to time t including the advertisers' types; the realizations of her values up to that time; her bids; the budgets of all advertisers; and the result of every past auction. We define a pure strategy β for advertiser k as a mapping from histories to bids, and we denote by $\mathbb{B}^{(K)}$ the space of strategies that are non-anticipating and adaptive to the history in market K. We study the expected payoff of advertiser k when she implements a strategy $\beta^{(K)} \in \mathbb{B}^{(K)}$ and all other advertisers follow FMFE strategies $\boldsymbol{\beta}^{\text{F},(K)}$ for market size K. The latter amounts to shading bids according to the multipliers that solve the NCP (5) while bidders have remaining budgets. This expected payoff is denoted by $J_k^{(K)}(\beta^{(K)}, \boldsymbol{\beta}_{-k}^{\mathrm{F},(K)})$, where the expectation is taken over the actual market process. In this notation, $J_k^{(K)}(\beta_k^{\mathrm{F},(K)}, \boldsymbol{\beta}_{-k}^{\mathrm{F},(K)})$ measures the *actual* expected payoff of the FMFE strategy for the advertiser in the exchange, which takes into account that advertisers may run out of budget before the end of the horizon. It is obvious that $J_k^{(K)}(\beta_k^{\mathrm{F},(K)}, \boldsymbol{\beta}_{-k}^{\mathrm{F},(K)}) \le \sup_{\beta \in \mathbb{B}^{(K)}} J_k^{(K)}(\beta, \boldsymbol{\beta}_{-k}^{\mathrm{F},(K)})$. We will analyze the gap $\sup_{\beta \in \mathbb{B}^{(K)}} J_k(\beta, \boldsymbol{\beta}_{-k}^{\mathrm{F},(K)}) - J_k^{(K)}(\beta_k^{\mathrm{F},(K)}, \boldsymbol{\beta}_{-k}^{\mathrm{F},(K)})$ to bound the sub-optimality of FMFE relative to unilaterally deviating to a best response strategy. In what follows, $O(\cdot)$ stands for Landau's big-O notation as K goes to infinity.

Theorem 6.1. *Suppose that Assumption 6.1 holds. Consider a sequence of markets indexed by K in which all bidders, except the k^{th} bidder, follow FMFE strategies $\boldsymbol{\beta}^{\mathrm{F},(K)}$ in market K. Suppose that the k^{th} advertiser unilaterally deviates and implements a non-anticipating and adaptive strategy $\beta^{(K)} \in \mathbb{B}^{(K)}$ in market K. The expected payoff of these deviations compared to the FMFE strategy satisfies*

$$\frac{1}{\alpha_k^{(K)} \eta^{(K)} s} \left(J_k^{(K)}(\beta^{(K)}, \boldsymbol{\beta}_{-k}^{\mathrm{F},(K)}) - J_k^{(K)}(\beta_k^{\mathrm{F},(K)}, \boldsymbol{\beta}_{-k}^{\mathrm{F},(K)}) \right) = O\left(\alpha_k^{(K)} + (\alpha_k^{(K)} \eta^{(K)} s)^{-1/2} K^{1/2} \right).$$

The bound in Theorem 6.1 states that $1 - J_k^{(K)}(\beta_k^{\mathrm{F},(K)}, \boldsymbol{\beta}_{-k}^{\mathrm{F},(K)}) / \sup_{\beta \in \mathbb{B}^{(K)}} J_k^{(K)}(\beta, \boldsymbol{\beta}_{-k}^{\mathrm{F},(K)})$ converges to zero as K grows to infinity when: (1) the matching probabilities $\alpha_k^{(K)}$ converge to zero; and (2) $K = o(\alpha_k^{(K)} \eta^{(K)} s)$, that is, the expected number of auctions a bidder participates in grows at a faster rate than the number of advertisers. In addition, the assumption imposes that the expected number of auctions a bidder participates in and the budget $b_k^{(K)}$ grow at the same rate. Typically, the scaling will also impose that the expected number of advertisers per auction remains constant (even though the overall number of advertisers grows large). These conditions naturally represent many Ad Exchange markets in which the number of auctions a bidder participates in is typically much larger than the number of competitors, the expected expenditure per auction is typically small compared to the budget, and the number of competitors per auction is small.

The key idea of the proof of Theorem 6.1 is to bound, in some appropriate way the impact that the k^{th} advertiser may have on the competitors and based on that, bound the value that may be obtained by deviating from the FMFE strategy. To do so, we exploit: first the fundamental observation that, independently of the k^{th} advertiser's strategy, the competing advertisers bid exactly as prescribed by the FMFE while they have budgets remaining; and second, the fact that not all advertisers match the same impressions and as a result, the impact of a single advertiser on any other specific advertiser (in terms of running out of budget) is limited. In particular, we establish that all advertisers will run out of budget close to the end of their campaigns no matter which strategy the deviant advertiser implements. Hence, the competitive landscape coincides with that predicted by the FMFE for most of the campaign. Based on this, we bound the performance of an arbitrary strategy by that of a strategy with the benefit of hindsight (which has complete knowledge of the future realizations of bids and values). This yields the result.

Finally, it is worthwhile to put this result in perspective with regard to typical revenue management heuristic fluid-based prescriptions. In most such settings, the bounds obtained (see, e.g., Talluri and van Ryzin (1998)) are of order $n^{-1/2}$ where n is a proxy for the number of opportunities (akin to the number of auctions one participates in our setting). In the present context, this term is present as $(\alpha_k^{(K)}\eta^{(K)}s)^{-1/2}$, but it is multiplied by $K^{1/2}$ to control for the fact that there are K bidders that could potentially run out of budget before the length of the campaign. Moreover, the term $\alpha_k^{(K)}$ in the bound controls for the potential impact bidder k may have on any competitor, which is bounded by the expected fraction of auctions in which they compete together.

6.2 Analysis for Small Markets

Recall that the FMFE concept involves two approximations: a fluid one motivated by the fact that advertisers participate in a large number of auctions during the course of their campaigns; and a mean-field approximation motivated by the fact that in the presence of many advertisers, it may not be necessary to track the state of each individual competitor. The first approximation is natural in the setting of Ad Exchanges where indeed advertisers participate in many repeated auctions and spend a small fraction of the budget in each one of them. In addition, while in some Ad Exchange markets the number of advertisers may be large, it is also useful to study the validity of the second approximation when this is not the case and the same set of advertisers meet repeatedly in common auctions. For this reason, we next isolate the impact of the mean-field approximation, and analyze it numerically for markets with a small number of advertisers.

To do so, we propose studying the best response to other advertisers playing FMFE in a *fluid model* in which there is a continuous flow of arriving impressions at rate η, auctions occur continuously in time, payments are infinitesimal, and budgets are depleted deterministically. The fluid model can be understood as an appropriately normalized market obtained in the limit as budgets and number of impressions are simultaneously scaled to infinity, while the number of players is fixed.

6.2.1 Fluid model

We introduce a *fluid model* in which impressions arrive continuously at a rate $\eta = 1$, the time horizon has a length s, and there are K advertisers in the market running synchronous campaigns. We denote by $\boldsymbol{x}(t) \in \mathbb{R}_+^K$ the vector of budgets remaining of the advertisers at time t as the state vector of the market, and by $\boldsymbol{b}$ the vector of initial budgets. At each point in time an advertiser determines an action in the space of bidding strategies $\mathcal{B} \triangleq [\underline{V}, \bar{V}] \to \mathbb{R}$, which maps a valuation to a bid. A *control policy* $\beta : \mathbb{R}_+ \times \mathbb{R}_+^K \to \mathcal{B}$ maps a point in time and state vector to an action.

The dynamics are given by the following. Let the functional $u_k : \mathcal{B}^K \to \mathbb{R}$ denote the instantaneous rate of expected utility obtained by the advertiser k when competing advertisers bid according to a given strategy profile. When the profile is $\boldsymbol{w} \in \mathcal{B}^K$ we have that

$$u_k(\boldsymbol{w}) = \alpha_k \mathbb{E}\Big[\mathbf{1}\{D_{-k} \le w_k(V_k)\}(V_k - D_{-k})\Big] ,$$

with the expectation taken over the valuation random variable and the maximum competing bid which is given by $D_{-k} = \max_{i \neq k, M_i = 1} \big\{w_i(V_i)\big\} \vee r$, where we let $M_i = 1$ indicate that the i^{th} agent participates in the auction. Similarly we let the functional $g_k : \mathcal{B}^K \to \mathbb{R}$ denote the instantaneous rate of expected expenditure incurred by the advertiser k when advertisers bid according to a given strategy profile, which is given by

$$g_k(\boldsymbol{w}) = \alpha_k \mathbb{E}\Big[\mathbf{1}\{D_{-k} \le w_k(V_k)\} D_{-k}\Big] .$$

Best response problem. We study the benefit of a unilateral deviation to a strategy that keeps track of the full market state, when competitors implement FMFE strategies. In this setting the FMFE strategies are given by $\beta_i^{\text{F}}(t, \boldsymbol{x})(v) = v/(1+\mu_i)\mathbf{1}\{x_i > 0\}$, where the multipliers $\boldsymbol{\mu}$ solve the NCP given in (5). The problem faced by advertiser k of determining the optimal payoff of a unilateral deviation when competitors implement the FMFE strategies is given by

$$\begin{aligned} \max_{\beta_k(t,\boldsymbol{x})} \quad & \int_0^s u_k\big(\beta_k(t, \boldsymbol{x}(t)), \beta_{-k}^{\text{F}}(t, \boldsymbol{x}(t))\big)\, \mathrm{d}t \\ \text{s.t.} \quad & \frac{\mathrm{d}\boldsymbol{x}(t)}{\mathrm{d}t} = -\boldsymbol{g}\big(\beta_k(t, \boldsymbol{x}(t)), \beta_{-k}^{\text{F}}(t, \boldsymbol{x}(t))\big), \qquad t \ge 0 \\ & \boldsymbol{x}(0) = \boldsymbol{b}, \quad \boldsymbol{x}(s) \ge 0. \end{aligned} \tag{6}$$

To simplify our arguments, for the rest of this section we assume that the reserve price $r = 0$. Moreover, we assume the following tie-breaking rule: when the advertiser under focus and her competitors have run out of budget, the advertiser under focus may still bid zero and win the remaining auctions.[13]

6.2.2 Best Response Analysis

We consider the case when advertisers have equal budgets, distribution of values and matching probabilities. We do allow, however, for advertiser k to have a different budget than its competitors. Because competitors are symmetric and the dynamics in the fluid model are deterministic, the budgets of the competitors deplete at the same rate. Thus, one can simplify the state by keeping track of the budget

[13] This is without loss of generality since by not bidding in a small fraction of the campaign, the advertiser under focus can guarantee that the competitors deplete first and, by saving an infinitesimal budget, she can win all the auctions with no competition for the remaining of the campaign.

of only one competitor.

Some definitions are in order. Let (μ_k, μ_{-k}) denote the multipliers associated with an FMFE. Let V_k^{FMFE} denote the total utility obtained by advertiser k when implementing the FMFE strategy $\beta_k^{\text{F}}(t, \boldsymbol{x})$.

Next, we define an alternative strategy. Let $H : \mathcal{B} \times \mathbb{R}^2 \to \mathbb{R}$ be a functional given by

$$H(w, \boldsymbol{p}) = u_k\big(w, \boldsymbol{w}_{-k}^{\text{F}}\big) - p_k g_k\big(w, \boldsymbol{w}_{-k}^{\text{F}}\big) - p_{-k} g_{-k}\big(w, \boldsymbol{w}_{-k}^{\text{F}}\big)\,.$$

where $w \in \mathcal{B}$ is a bidding strategy, $w_i^{\text{F}}(v) = v/(1 + \mu_i)$ the FMFE bidding strategies, $g_{-k}(\cdot)$ denotes the instantaneous rate of expected expenditure incurred by one of the competitors of firm k, and $p_k, p_{-k} \in \mathbb{R}$. Consider the following problem

$$V_k^{\text{D}} \triangleq \inf_{p_k \geq 0, p_{-k}} \alpha s \mathbb{E} V + p_k b_k + p_{-k} b_{-k} \tag{7a}$$

$$\text{s.t.} \;\; \sup_{w \in \mathcal{B}} H(w, \boldsymbol{p}) \leq \alpha \mathbb{E} V\,, \tag{7b}$$

which is a convex optimization problem since the set $\mathcal{P} = \{\boldsymbol{p} \in \mathbb{R}^2 : \sup_{w\in\mathcal{B}} H(w, \boldsymbol{p}) \leq \alpha \mathbb{E} V\}$ is convex. The latter follows because the lower-level set of a convex function is convex, and the point-wise supremum of linear functions is a convex function (see, e.g., Boyd and Vandenberghe (2009)). Additionally, let V_k^{D} denote the value of (7), with the convention that it is $-\infty$ if it is unbounded; and when it is bounded, let p_k^* and p_{-k}^* denote a corresponding optimal solution. Let, assuming it is well defined, $\bar{w} \in \arg\max_{w\in\mathcal{B}} H(w, \boldsymbol{p}^*)$ be the bidding strategy that verifies the supremum in constraint (7b).

Theorem 6.2 (Best response strategy)**.** *Suppose that bidders' values possess increasing failure rates (IFR) and have bounded support, and that the reserve price is zero. Suppose that all competing advertisers use FMFE strategies. Then,*

- *i.) If $V_k^{\text{D}} \leq V_k^{\text{FMFE}}$, the FMFE strategy is the optimal control for advertiser k in problem* (6)*.*
- *ii.) If $V_k^{\text{D}} > V_k^{\text{FMFE}}$, $\bar{w}$ is well defined and the optimal strategy for advertiser k in problem* (6) *is to bid according to $\bar{w}(\cdot)$ until competitors deplete their budgets, and zero afterwards. Furthermore this strategy yields exactly V_k^{D}.*

In other words, the result states that the value of the deviant advertiser's control is the maximum of V_k^{D} and V_k^{FMFE}. Furthermore, the result provides a crisp characterization of an optimal policy: one would only need to compute two candidate strategies, the FMFE strategy and $\bar{w}(\cdot)$, to determine a best response and the associated payoff. We show in the proof that when $V_k^{\text{D}} > V_k^{\text{FMFE}}$, then the competitors will deplete their budgets before the end of the horizon under $\bar{w}$, allowing advertiser k to take advantage of the time during which she operates alone in the market. This result highlights the only type of profitable deviation that one may witness compared to FMFE: use a stationary strategy to deplete competitors faster than what the FMFE strategy does. The strategy involves bidding above ones value in some auctions, and carefully balances the lower expected net utility in the first part of the campaign with the benefit of facing no competition at the end of the campaign.[14]

Quite remarkably, one may establish that in some special cases of interest, the strategy $\bar{w}$ admits a very simple structure: in the cases of uniform and exponential distributions, one may restrict attention to affine bidding functions when searching for a best response (see Corollary D.1 in Appendix D). Furthermore, one may establish that in this fluid model the losses of playing FMFE relative to a best response are at most of order $O(\alpha_k)$, a behavior we illustrate numerically next.

6.2.3 Numerical Experiments: FMFE Sub-Optimality Gap

Intuitively, when there are multiple players in the market, depleting the budgets of the competitors becomes more costly and as a result the benefit introduced from deviating from the FMFE strategies becomes negligible. To investigate this, we compare the campaign utility of an advertiser in the fluid model under the FMFE strategy with that of the best response as the number of competitors increases for many problem instances with different parameters. We present the results of a representative instance in Figure 3.[15] Budgets and matching probabilities decrease with the number of competitors so that the average number of matching advertisers per auction remains invariant, equal to 2. We plot the relative sub-optimality gap as a function of the number of advertisers. For a given number of advertisers, we analyze the gap when all competitors have the same budget, but we allow the budget of the advertiser under analysis to change and be 75%, 100%, or 150% of the individual budgets of competitors. This allows to study the gap when the deviant advertiser has varying degrees of market influence.

We observe that as the number of player increases, the sub-optimality of playing FMFE decreases fast. As a matter of fact, for the case of identical advertisers ($b_1^{(K)} = b_2^{(K)} = \ldots = b_k^{(K)}$), the FMFE strategy yields utility within 2.5% of that obtained by a best response as soon as there are more than 6 advertisers in the market. In addition, when the deviant firm has a smaller budget, her ability to deplete its competitors decreases.

In Figure 4, we analyze the same setting as earlier except that now, we fix the matching probability to $\alpha = 1$. In other words, all advertisers participate in all auctions. In some settings, it is possible

[14] Lu and Zhu (2013) also identify similar strategies in which one advertiser tries to deplete the budget of its competitor in a stylized sponsored search auction duopoly model under complete information.

[15] All results can be obtained from the authors upon request.

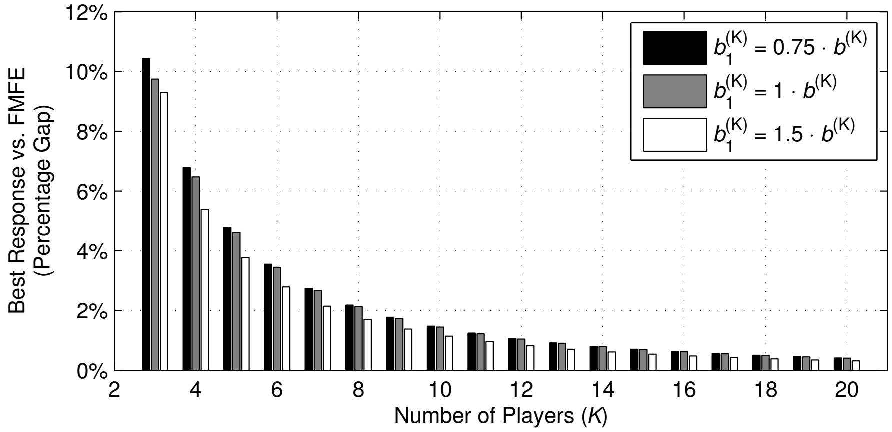

Figure 3: **FMFE versus best response.** Advertisers are homogenous with arrival rate $\eta = 1$, campaign length $s = 16$, competitors' budgets $b^{(K)} = 4/K$, matching probabilities $\alpha^{(K)} = 2/K$, and uniform valuations with support $[0, 2]$.

to imagine that a small number of advertisers would focus on the same viewer types and hence would compete more intensively. In such a setting, the sub-optimality gap of FMFE decreases fast as the number of competitors increases, dropping below 5% when there are more than 5 players in the market and getting around 2% when there are 8. We highlight here that the sub-optimality gap we estimate in these examples are conservative in that the benchmark policy has unrealistic informational requirements; in practice, bidders would not be able to perfectly monitor competitors' budgets. Hence, their ability to strategize to deplete competitors' budgets would be even more limited.

The fluid analysis and our numerical results above suggest that the value of tracking the market state is small even in the presence of few competitors. In other words, a given bidder has a limited ability to strategize and impact the market when all other competitors play a FMFE strategy. This provides further practical support to use FMFE as a solution concept to study competition in Ad Exchanges.

7 Conclusions

A framework for the analysis of the impact of different levers. In this paper, the analysis has focused on optimally setting the reserve price. However, the proposed framework based on FMFE is general and may be used to study other important auction design decisions for the publisher.

In fact, it is possible to show that the framework proposed allows to quantify the impact of increasing the allocation of impressions sent to the exchange vis-à-vis collecting the opportunity cost upfront on the bidding behavior of advertisers, and optimize this allocation while accounting for budgets.

We also show how one may optimize other dimensions that may be under the control of the publisher

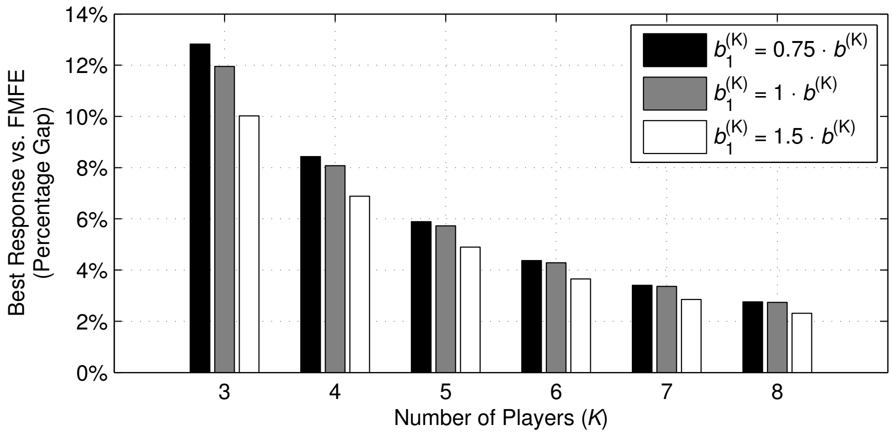

Figure 4: **FMFE versus best response.** Advertisers are homogenous with arrival rate $\eta = 1$, campaign length $s = 16$, budgets $b^{(K)} = 4/K$, matching probabilities $\alpha^{(K)} = 1$, and uniform valuations with support $[0, 2]$.

such as the extent of user information to disclose to the advertisers. On the one hand, more information enables advertisers to improve targeting, which results in higher bids conditional on participating in an auction. On the other hand, as more information is provided, fewer advertisers match with each user, resulting in thinner markets, which could decrease the publisher's profit.[16] We show that given any mapping from user information to advertiser valuation distribution, one may apply our framework to quantify the impact of budgets on the key trade-offs at play. In particular, we demonstrate this through a stylized model for information disclosure with homogenous bidders.

These results, available in the working paper version of this work, complement the ones in the paper and reinforce the importance of reserve price optimization. In particular, we show that proper adjustment of the reserve price is key in (1) making profitable for the publisher to try selling all impressions in the exchange before utilizing the alternative channel; and (2) compensating for the thinner markets created by greater disclosure of viewers' information.

Building on the framework. Overall, our results provide a new approach to study Ad Exchange markets and the publishers' decisions. The techniques developed build on two fairly distinct streams of literature, revenue management and mean-field models and are likely to have additional applications. The sharp results regarding the publisher's decisions could inform how these markets are designed in practice. At the same time, our framework opens up the door to study a range of other relevant issues in this space. For example, one interesting avenue for future work may be to study the impact of Ad networks, that aggregate bids from different advertisers and bid on their behalf, on the resulting

[16]This trade-off is discussed in Levin and Milgrom (2010). Fu et al. (2012) studies this problem in the context of a static auction with out budget constraints and shows that if the auctioneer implements the optimal mechanism, then additional data leads to additional revenue.

competitive landscape and auction design decisions. Similarly, another interesting direction to pursue is to incorporate common advertisers' values and analyze the impact of cherry-picking and adverse selection. Finally, our framework and its potential extensions can provide a possible structural model for bidding behavior in exchanges, and open the door to pursue an econometric study using transactional data in exchanges.

Acknowledgments

The authors thank Awi Federgruen, Vahab Mirrokni, Hamid Nazerzadeh, Nicolas Stier-Moses, and John Turner for helpful insights as well as the participants at the 2012 MSOM Conference, ISMP 2012, INFORMS 2012, INFORMS Revenue Management and Pricing Conference 2013, NET Institute Conference 2013, ACM Conference on Electronic Commerce 2013, and the audiences at various universities for their feedback. They also thank the AE and three anonymous referees for valuable input that significantly improved the paper. Financial support support from the NET Institute (www.NETinst.org) and the Deming Center (www.gsb.columbia.edu/deming) is greatly acknowledged.

References

Abraham, I., Athey, S., Babaioff, M. and Grubb, M. (2012), Peaches, lemons, and cookies: Designing auction markets with dispersed information. Working paper.

Adlakha, S., Johari, R. and Weintraub, G. Y. (2011), Equilibria of dynamic games with many players: Existence, approximation, and market structure. Working paper.

Alaei, S., Arcaute, E., Khuller, S., Ma, W., Malekian, A. and Tomlin, J. (2009), Online allocation of display advertisements subject to advanced sales contracts, *in* 'Proceedings of the Third International Workshop on Data Mining and Audience Intelligence for Advertising', ADKDD '09, ACM, pp. 69–77.

Araman, V. and Fridgeirsdottir, K. (2011), Cost-per-impression pricing and campaign delivery for online display advertising. Working Paper.

Araman, V. and Popescu, I. (2010), 'Media revenue management with audience uncertainty: Balancing upfront and spot market sales', *Manufacturing & Service Operations Management* **12**(2), 190–212.

Arnon, A. and Mansour, Y. (2011), Repeated budgeted second price ad auction, *in* G. Persiano, ed., 'Algorithmic Game Theory', Vol. 6982 of *Lecture Notes in Computer Science*, Springer Berlin Heidelberg, pp. 7–18.

Asmussen, S. (2003), *Applied Probability and Queues*, Stochastic Modelling and Applied Probability, Vol. 51, 2 edn, Springer.

Aven, T. (1985), 'Upper (lower) bounds on the mean of the maximum (minimum) of a number of random variables', *Journal of Applied Probability* **22**(3), 723–728.

Balseiro, S., Feldman, J., Mirrokni, V. S. and Muthukrishnan, S. (2011), Yield optimization of display advertising with ad exchange, *in* 'ACM Conference on Electronic Commerce', pp. 27–28.

Balseiro, S. R., Besbes, O. and Weintraub, G. Y. (2012), Auctions for online display advertising exchanges: Approximation result. Technical Report.

Bergemann, D. and Bonatti, A. (2011), Targeting in advertising markets: Implications for offline vs. online media. working paper.

Bergemann, D. and Said, M. (2010), Dynamic auctions: A survey. Working Paper.

Board, S. (2009), 'Revealing information in auctions: the allocation effect', *Economic Theory* **38**, 125–135.

Boyd, S. and Vandenberghe, L. (2009), *Convex Optimization*, Cambridge University Press.

Celis, L. E., Lewis, G., Mobius, M. M. and Nazerzadeh, H. (2011), Buy-it-now or take-a-chance: a simple sequential screening mechanism, *in* 'Proceedings of the 20th International Conference on World Wide Web', pp. 147–156.

Che, Y.-K. and Gale, I. (1998), 'Standard auctions with financially constrained bidders', *The Review of Economic Studies* **65**(1), 1–21.

Che, Y.-K. and Gale, I. (2000), 'The optimal mechanism for selling to a budget-constrained buyer', *Journal of Economic Theory* **92**(2), 198 – 233.

Chen, Y.-Y. (2011), Optimal dynamic auctions for display advertising. Working paper.

David, H. A. and Nagaraja, H. N. (2003), *Order statistics*, thrid edn, John Wiley and Sons.

de Albéniz, V. M. and Talluri, K. (2011), 'Dynamic price competition with fixed capacities', *Management Science* **57**, 1078–1093.

Doraszelski, U. and Pakes, A. (2007), A framework for applied dynamic analysis in IO, *in* 'Handbook of Industrial Organization', Vol. 3, North-Holland, Amsterdam.

Dudey, M. (1992), 'Dynamic edgeworth-bertrand competition', *Quarterly Journal of Economics* **107**, 1461–1477.

Eick, S. G., Massey, W. A. and Whitt, W. (1993), 'The physics of the $M_t/G/\infty$ queue', *Operations Research* **41**(4), 731–742.

Facchinei, F. and Pang, J.-S. (2003*a*), *Finite-Dimensional Variational Inequalities and Complementarity Problems, Volume 1*, Springer Series in Operations Research and Financial Engineering, Springer.

Facchinei, F. and Pang, J.-S. (2003*b*), *Finite-Dimensional Variational Inequalities and Complementarity Problems, Volume 2*, Springer Series in Operations Research and Financial Engineering, Springer.

Farias, V., Monsch, M. and Perakis, G. (2011), Tractable notions of equilibrium in revenue management, *in* 'MSOM Annual Conference'.

Fridgeirsdottir, K. and Najafi, S. (2010), Cost-per-click pricing for display advertising. Working Paper.

Fu, H., Jordan, P., Mahdian, M., Nadav, U., Talgam-Cohen, I. and Vassilvitskii, S. (2012), Ad auctions with data, *in* 'Proceedings of the 5th international conference on Algorithmic Game Theory', SAGT'12, Springer-Verlag, pp. 168–179.

Gallego, G. and Hu, M. (2011), Dynamic pricing of perishable assets under competition. Working paper.

Gallego, G. and van Ryzin, G. (1994), 'Optimal dynamic pricing of inventories with stochastic demand over finite horizons', *Management Science* **40**(8), 999–1020.

Ghosh, A., Rubinstein, B. I., Vassilvitskii, S. and Zinkevich, M. (2009), Adaptive bidding for display advertising, *in* 'Proceedings of the 18th international conference on World wide web', WWW '09, ACM, pp. 251–260.

Graham, C. and Méléard, S. (1994), 'Chaos hypothesis for a system interacting through shared resources', *Probability Theory and Related Fields* **100**(2), 157–174.

Gummadi, R., Key, P. and Proutiere, A. (2012), Repeated auctions under budget constraints: Optimal bidding strategies and equilibria. Working paper.

Horn, R. A. and Johnson, C. R. (1991), *Topics in Matrix Analysis*, Cambridge University Press.

Internet Advertising Bureau (2012), Internet advertising revenue report, 2011 full year results, Technical report, PricewaterhouseCoopers.

Iyer, K., Johari, R. and Sundararajan, M. (2011), Mean Field Equilibria of Dynamic Auctions with Learning. SSRN eLibrary.

Jofre-Bonet, M. and Pesendorfer, M. (2003), 'Estimation of a dynamic auction game', *Econometrica* **71**(5), 1443–1489.

Laffont, J.-J. and Maskin, E. (1980), 'Optimal reservation price in the vickrey auction', *Economics Letters* **6**(4), 309 – 313.

Laffont, J.-J. and Robert, J. (1996), 'Optimal auction with financially constrained buyers', *Economics Letters* **52**(2), 181 – 186.

Lariviere, M. A. (2006), 'A note on probability distributions with increasing generalized failure rates', *Oper. Res.* **54**, 602–604.

Levin, J. and Milgrom, P. (2010), 'Online advertising: Heterogeneity and conflation in market design', *American Economic Review: Papers & Proceedings* **100**, 603–607.

Liu, G. and Li, W. V. (2009), 'Moment bounds for truncated random variables', *Statistics & Probability Letters* **79**(18), 1951 – 1956.

Lu, S. and Zhu, Y. (2013), An analysis of the impact of advertising budget in position auctions. Working Paper.

Mas-Colell, A., Whinston, M. D. and Green, J. R. (1995), *Microeconomic Theory*, Oxford University Press.

Maskin, E. S. (2000), 'Auctions, development, and privatization: Efficient auctions with liquidity-constrained buyers', *European Economic Review* **44**(46), 667 – 681.

McAfee, P., Ghosh, A., Papineni, K. and Vassilvitskii, S. (2009), Bidding for representative allocations for display advertising, *in* 'WINE', pp. 208–219.

Milgrom, P. (1999), The envelope theorems. Working Paper.

Muthukrishnan, S. (2009), Ad exchanges: Research issues., *in* 'Internet and Network Economics, Lecture Notes in Computer Science', pp. 1–12.

Myerson, R. (1981), 'Optimal auction design', *Mathematics of Operations Research* **6**(1), 58–73.

Netmining (2011), 'Mapping the display landscape: A fundamental guide to data, targeting and the future of display'.

Pai, M. M. and Vohra, R. (2011), Optimal auctions with financially constrained bidders. Working Paper.

Roels, G. and Fridgeirsdottir, K. (2009), 'Dynamic revenue management for online display advertising', *Journal of Revenue and Pricing Management* **8**, 452–466.

Ross, S. M. (1996), *Stochastic Processes*, John Wiley and Sons.

Talluri, K. and van Ryzin, G. (1998), 'An analysis of bid-price controls for network revenue management', *Management Science* **44**(11), 1577–1593.

Turner, J. (2012), 'The planning of guaranteed targeted display advertising', *Operations Research* .

Vulcano, G., van Ryzin, G. and Maglaras, C. (2002), 'Optimal dynamic auctions for revenue management', *Management Science* **48**(11), 1388–1407.

Wall Street Journal (2013), 'Automated ad buying surges online'.

Weintraub, G. Y., Benkard, C. L. and Roy, B. V. (2008), 'Markov perfect industry dynamics with many firms', *Econometrica* **76**, 1375 – 1411.

Yang, J., Vee, E., Vassilvitskii, S., Tomlin, J., Shanmugasundaram, J., Anastasakos, T. and Kennedy, O. (2010), Inventory allocation for online graphical display advertising. Yahoo Technical Report.

A Selected Proofs

A.1 Proof of Proposition 3.1

We prove the result in three steps. First, we derive the dual of the primal problem by introducing a lagrange multiplier for the budget constraint. Second, we determine the optimal dual solution through first-order conditions. Third, we show that complementary slackness holds and that there is no duality gap. To simplify notation we drop the dependence on F_d when clear from the context.

Step 1. We introduce a lagrange multiplier $\mu \geq 0$ for the budget constraint and let

$$\mathcal{L}_\theta(w,\mu) = \alpha\eta s \mathbb{E}\left[\mathbf{1}\{D \leq w(V)\}\Big(V-(1+\mu)D\Big)\right] + \mu b$$

denote the Lagrangian for type θ (for simplicity we omit the subindex θ for other quantities). The dual problem is given by

$$\begin{aligned}\inf_{\mu\geq 0}\sup_{w(\cdot)} \mathcal{L}_\theta(w,\mu) &= \inf_{\mu\geq 0}\left\{\alpha\eta s \sup_{w(\cdot)}\left\{\mathbb{E}\left[\mathbf{1}\{D\leq w(V)\}\Big(V-(1+\mu)D\Big)\right]\right\} + \mu b\right\}\\ &= \inf_{\mu\geq 0}\left\{\alpha\eta s\mathbb{E}\left[\mathbf{1}\{(1+\mu)D \leq V\}\Big(V-(1+\mu)D\Big)\right] + \mu b\right\}\\ &= \inf_{\mu\geq 0}\left\{\alpha\eta s\mathbb{E}\Big[V-(1+\mu)D\Big]^+ + \mu b\right\},\end{aligned}$$

where the second equality follows from observing that the inner optimization problem is similar to the problem faced by a bidder with value $\frac{v}{1+\mu}$ seeking to maximize its expected utility in a second-price auction, in which case it is optimal to bid truthfully. Let $\Psi_\theta(\mu) = \alpha\eta s\mathbb{E}\Big[V-(1+\mu)D\Big]^+ + \mu b$. Notice that the term within the expectation is convex in μ; given that expectation preserves convexity, the dual problem is convex. As a consequence of the previous analysis one obtains for any given multiplier $\mu \geq 0$, the policy $w(v) = \frac{v}{1+\mu}$ maximizes the Lagrangian.

Step 2. In order to characterize the optimal multiplier we shall analyze the first-order conditions of the dual problem. The integrability of D, in conjunction with the dominated convergence theorem, yield that Ψ_θ is differentiable w.r.t. μ. The derivative is given by $\frac{d}{d\mu}\Psi_\theta = b - \alpha\eta s\mathbb{E}\left[\mathbf{1}\left\{D \leq \frac{V}{1+\mu}\right\}D\right]$, which is equal to the expected remaining budget by the end of the campaign when the optimal bid function is employed.

Suppose $\alpha\eta s\mathbb{E}\left[\mathbf{1}\left\{D \le V\right\} D\right] \le b$, i.e., Ψ_θ admits a non-negative derivative at $\mu = 0$. Since Ψ_θ is convex, the optimal multiplier is $\mu^* = 0$. Suppose now $\alpha\eta s\mathbb{E}\left[\mathbf{1}\left\{D \le V\right\} D\right] > b$. The derivative of Ψ_θ is continuous (by another application of the dominated convergence theorem) and converges to b as $\mu \to \infty$. We deduce that the equation $\alpha\eta s\mathbb{E}\left[\mathbf{1}\left\{D \le \frac{V}{1+\mu}\right\} D\right] = b$, admits a solution and the optimal multiplier μ^* solves the latter.

Step 3. Combining both cases, one obtains that the optimal multiplier μ^* and the corresponding bid function $\beta^{\mathrm{F}}_\theta(v) = v/(1+\mu^*)$ satisfy $\mu^*\Big(b - \alpha\eta s\mathbb{E}\left[\mathbf{1}\left\{D \le \beta^{\mathrm{F}}_\theta(V)\right\} D\right]\Big) = 0$, and thus the complementary slackness conditions hold. Additionally from the first-order conditions of the dual, we get that the bid function $\beta^{\mathrm{F}}_\theta(\cdot)$ is primal feasible. We conclude by showing that the primal objective of the proposed bid function attains the dual objective. That is,

$$\begin{aligned}\alpha\eta s\mathbb{E}\Big[\mathbf{1}\{D \le \beta^{\mathrm{F}}_\theta(V)\}\Big(V - D\Big)\Big] &= \mathcal{L}_\theta(\beta^{\mathrm{F}}_\theta, \mu^*) + \mu^*\Big(b - \alpha\eta s\mathbb{E}\left[\mathbf{1}\left\{D \le \beta^{\mathrm{F}}_\theta(V)\right\} D\right]\Big)\\ &= \mathcal{L}_\theta(\beta^{\mathrm{F}}_\theta, \mu^*) = \Psi_\theta(\mu^*),\end{aligned}$$

where the second equality follows from the complementary slackness conditions and the last from the fact that $\Psi_\theta(\mu^*) = \sup_{w(\cdot)} \mathcal{L}_\theta(w, \mu^*)$, and the fact $\beta^{\mathrm{F}}_\theta$ is the optimal bid function.

A.2 Proof of Theorem 6.1

We prove the result in two steps. First, we lower bound the expected performance of the k^{th} advertiser when all advertisers (including herself) implement the FMFE strategy in terms of the objective value of the fluid problem (1). Second, we upper bound the expected payoff of *any* strategy the k^{th} advertiser may implement when the remaining implement the FMFE strategies via a hindsight bound.

Proposition A.1 (Lower Bound)**.** *Suppose that Assumption 6.1 holds and all advertisers implement FMFE strategies $\boldsymbol{\beta}^{\mathrm{F}}$. The expected payoff of the k^{th} advertiser is lower bounded by*

$$\frac{1}{\alpha_k\eta s} J_k(\beta^{\mathrm{F}}_k, \boldsymbol{\beta}^{\mathrm{F}}_{-k}) \ge \bar{J}^{\mathrm{F}}_k - O\left((\alpha_k\eta s)^{-1/2} K^{1/2}\right),$$

where $\bar{J}^{\mathrm{F}}_k \triangleq J^{\mathrm{F}}_k/(\alpha_k\eta s)$ is the normalized objective value of the fluid problem (1)*.*

The performance metric $J_k(\beta^{\mathrm{F}}_k, \boldsymbol{\beta}^{\mathrm{F}}_{-k})$ may differ from the FMFE value function, given by the objective value of the approximation problem J^{F}_k, since the former takes into account that bidders may run out of budget before the end of their campaigns. The proof is based on the fundamental observation that advertisers bid exactly as prescribed by the FMFE while they have budgets remaining. This allows one to consider an alternative system where advertisers are allowed to bid (i) when they have no budget, and (ii) after the end of their campaigns; and in which the expected performance *exactly* coincides with that of the approximation problem J^{F}_k. Using a coupling argument the proof shows that the expected performance in the original and alternate systems coincide until the first time *some* advertiser runs out of budget, which in turn is shown to be close to the end of the horizon via a martingale argument.

Proposition A.2 (Upper Bound)**.** *Suppose that Assumption 6.1 holds and all advertisers implement FMFE strategies $\boldsymbol{\beta}^{\text{F}}$, and the k^{th} advertiser implements an alternative strategy $\beta \in \mathbb{B}$. The expected payoff of the k^{th} advertiser is upper bounded by*

$$\frac{1}{\alpha_k \eta s} J_k(\beta, \boldsymbol{\beta}^{\text{F}}_{-k}) \leq \bar{J}^{\text{F}}_k + O\left(\alpha_k + (\alpha_k \eta s)^{-1/2} K^{1/2}\right).$$

To prove the result, we first upper bound the performance of an arbitrary strategy by that of a strategy with the benefit of hindsight (which has complete knowledge of the future realizations of bids and values). This is akin to what is typically done in revenue management settings (see, e.g., Talluri and van Ryzin (1998)), with the exception that here, the competitive environment (which is the counterpart of the demand environment in RM settings) is endogenous and determined through the FMFE consistency requirement. As a result, the optimal hindsight policy may force competitors to run out of budget so as to reduce competition. To facilitate the analysis of the expected performance of the hindsight policy the proof considers the same alternative system in which competitors bid regardless of the budget; in which the hindsight policy can be analyzed simply via linear programming duality theory. Since the original and alternative system coincide until some advertiser runs out of budget, we are left again with the problem of showing that advertisers run out of budget close to the end of the campaign.

The proof concludes by showing the k^{th} advertiser has a limited impact on the system, in the sense that competitors run out of budget -in expectation- close to the end of their campaigns no matter which strategy the advertiser implements. To this end the proof exploits that any two advertisers compete a limited number of times during their campaigns to bound the potential impact the k^{th} advertiser may have on her competitors. This result relies heavily on the matching probability decreasing with the scaling.

A.3 Proof of Proposition A.1

Consider an alternate system in which advertisers are allowed to bid (i) when they have no budget, and (ii) after the end of their campaigns. The argument revolves around the fact that the performance of the tagged advertiser in the real and alternate coincide until the first time some advertiser runs out of budget. This follows from the fact that advertisers bid exactly as prescribed by the FMFE while they have budgets remaining.

In order to study the performance on the alternate system we shall consider the sequence $\{(Z_{n,k}, U_{n,k})\}_{n\geq 1}$ of realized expenditures and utilities of the k^{th} in the alternate system. In view of our mean-field assumption this sequence is i.i.d. and independent of the impressions' inter-arrival times. The k^{th} advertiser's expenditure in the n^{th} auction is $Z_{n,k} = M_{n,k} \mathbf{1}\{D_{n,-k} \leq \beta^{\text{F}}_k(V_k)\} D_{n,-k}$ and her corresponding utility is $U_{n,k} = M_{n,k} \mathbf{1}\{D_{n,-k} \leq \beta^{\text{F}}_k(V_{n,k})\}(V_{n,k} - D_{n,-k})$. Additionally, let $b'_k(t) = b_k - \sum_{n=1}^{N(t)} Z_{n,k}$ be the evolution of the k^{th} advertiser's budget in this alternate system, where we denote by $N(t)$ the number of impressions arrived by time t.

The following stopping time will play a key role in the proof. Let $\tilde{N}_k$ be the first auction in which advertiser k^{th} runs out of budget, that is, $\tilde{N}_k = \inf\{n \geq 1 : b'_k(t_n) < 0\}$. This stopping time is relative to all auctions in the market and not restricted to the auctions in which the k^{th} advertiser participates. Similarly, let $\tilde{N}$ as the first auction in which some advertiser runs out of budget, that is, $\tilde{N} = \min_k \tilde{N}_k$.

Next, we lower bound the performance of the k^{th} advertiser. Denoting by I_k the number of auctions that advertiser k^{th} participates during his campaign, that is, $I_k = \sum_{n=1}^{N(s)} M_{n,k}$; and by $\tilde{I}_k$ the number of auctions that advertiser k^{th} participates until some agent runs out of budget, that is, $\tilde{I}_k = \sum_{n=1}^{\tilde{N}} M_{n,k}$; one obtains by using a coupling argument that the performance of both systems coincide until time $\tilde{N}$ and as result

$$\begin{aligned} J_k(\beta^{\text{F}}, \boldsymbol{\beta}^{\text{F}}_{-k}) &\geq \mathbb{E}\left[\sum_{n=1}^{\tilde{N}\wedge N(s)} U_{n,k}\right] \geq \mathbb{E}\left[\sum_{n=1}^{N(s)} U_{n,k}\right] - \bar{V}\mathbb{E}\left[\sum_{n=1}^{N(s)} M_{n,k} - \sum_{n=1}^{\tilde{N}} M_{n,k}\right]^+ \\ &= \mathbb{E}\left[\sum_{n=1}^{N(s)} U_{n,k}\right] - \bar{V}\mathbb{E}[I_k - \tilde{I}_k]^+ \\ &\geq \mathbb{E}\left[\sum_{n=1}^{N(s)} U_{n,k}\right] - \bar{V}\mathbb{E}[I_k - \alpha_k \eta s]^+ - \mathbb{E}[\alpha_k \eta s - \tilde{I}_k]^+ \end{aligned}$$

where the first inequality follows from discarding all auctions after the time some advertiser runs out of budget; the second from the fact that $0 \leq U_{n,k} \leq M_{n,k}\bar{V}$; and the third from the fact that for every $a, b, c \in \mathbb{R}$ we have that $(a-c)^+ \leq (a-b)^+ + (b-c)^+$. In the remainder of the proof we address one term at a time.

Term 1. Notice that the in the alternate system the number of matching impressions in the campaign is independent of the utility, and thus we have that

$$\mathbb{E}\left[\sum_{n=1}^{N(s)} U_{n,k}\right] = \alpha_k \eta s \mathbb{E}[U_{1,k}] = \Psi_k(\mu_k; F_d) + \mu_k(G_k(\boldsymbol{\mu}) - \beta_k) = J_k^{\text{F}},$$

where the second equality follows from the fact that $\beta_k^{\text{F}}(x) = x/(1+\mu_k)$ and $U_{n,k} = (V_{n,k} - (1+\mu_k)D_{n,k})^+ + \mu_k Z_{n,k}$, and the last from complementarity slackness and the optimality of the FMFE multipliers.

Term 2. Note that for any random variable X and constant x, we have that $\mathbb{E}(X-x)^+ \leq (\mathbb{E}X - x)^+ + \sqrt{\text{Var}(X)/2}$, by the upper bound on the maximum of random variables given in Aven (1985). Because the agent participates in each auction with probability α_k, we have that I_k is a Poisson random variable with mean $\alpha_k \eta s$ and one obtains that

$$\frac{1}{\alpha_k \eta s}\mathbb{E}[I_k - \alpha_k \eta s]^+ \leq (2\alpha_k \eta s)^{-1/2} = O\left((\alpha_k \eta s)^{-1/2}\right).$$

Term 3. Define $\tilde{I}_{k,i}$ as the number of auctions that advertiser k^{th} participates until agent i^{th} runs out of budget, that is, $\tilde{I}_{k,i} = \sum_{n=1}^{\tilde{N}_i} M_{n,k}$. Using this notation we obtain that the number of auctions the k^{th} advertiser participates until someone runs out of budget can be alternatively written as $\tilde{I}_k = \sum_{n=1}^{\min_i \tilde{N}_i} M_{n,k} = \min_i \sum_{n=1}^{\tilde{N}_i} M_{n,k} = \min_i \tilde{I}_{k,i}$. Using this identity we obtain that

$$\begin{aligned}\mathbb{E}\left[\alpha_k \eta s - \tilde{I}_k\right]^+ &= \mathbb{E}\left[\alpha_k \eta s - \min_i \tilde{I}_{k,i}\right]^+ = \mathbb{E}\left[\max_i \{\alpha_k \eta s - \tilde{I}_{k,i}\}^+\right] \\ &\leq \max_i \left\{\alpha_k \eta s - \mathbb{E}\tilde{I}_{k,i}\right\}^+ + \sqrt{\sum_i \operatorname{Var}[\tilde{I}_{k,i}]},\end{aligned}$$

where the inequality follows from the upper bound on the maximum of random variables given in Aven (1985), that is, for any sequence of random variables $\{X_i\}_{i=1}^n$ we have that $\mathbb{E}[\max_i X_i] \leq \max_i \mathbb{E}X_i + \sqrt{\frac{n-1}{n}\sum_i \operatorname{Var}(X_i)}$. Dividing by the expected number of impressions in the horizon and using the bounds on the mean and variance of the stopping times of Lemma B.3 we get that

$$\begin{aligned}\frac{1}{\alpha_k \eta s}\mathbb{E}[\alpha_k \eta s - \tilde{I}_k]^+ &\leq \max_i \left\{1 - \frac{b_i}{\alpha_i \eta s G_i(\boldsymbol{\mu})}\right\}^+ + \frac{1}{\alpha_k \eta s}\sqrt{\sum_{i=1}^K O(b_i)} \\ &= O\left((\alpha_k \eta s)^{-1} K^{1/2} \bar{b}^{1/2}\right) = O\left((\alpha_k \eta s)^{-1/2} K^{1/2}\right),\end{aligned}$$

where the second inequality follows from the fact that the expected expenditure in the FMFE never exceeds the budget, that is, $\alpha_i \eta s G_i(\boldsymbol{\mu}) \leq b_i$, and by setting $\bar{b} = \max_i b_i$; and the last because $\alpha_k \eta s = O(\bar{b})$ from Assumption 6.1.

Supplementary Appendix

B Additional Proofs

We start by providing characterizations of the distribution of the maximum bid and the expenditure function that are used throughout the results.

Lemma B.1. *i.) The distribution of the maximum competing bid when $x \geq r$ is given by*

$$F_d(x;\boldsymbol{\mu}) = \exp\left\{-\mathbb{E}[\alpha_\Theta \lambda s_\Theta] \sum_\theta \mathbb{P}_{\hat{\Theta}}\{\theta\} \bar{F}_{v_\theta}((1+\mu_\theta)x)\right\},$$

where $F_{v_\theta}(\cdot) \triangleq F_v(\cdot;\gamma_\theta)$ is the distribution of values for type θ.

ii.) The expenditure function for type θ can be characterized by

$$G_\theta(\boldsymbol{\mu}, r) = r\bar{F}_{v_\theta}((1+\mu_\theta)r)F_d(r;\boldsymbol{\mu}) + \int_r^{\bar{V}} x\bar{F}_{v_\theta}((1+\mu_\theta)x)\, dF_d(x;\boldsymbol{\mu}).$$

Proof of Lemma B.1. We prove each item at a time.

i.) Let $F_w(\cdot;\boldsymbol{\mu})$ be the cumulative distribution function of the bid from a single matching advertiser when bidders implement the fluid-based strategy with a profile of multipliers $\boldsymbol{\mu}$, which is given by the random variable $\hat{W} = V_{\hat{\Theta}}/(1+\mu_{\hat{\Theta}})$. Since valuations are i.i.d., one can write the c.d.f. of bids as $F_w(x;\boldsymbol{\mu}) = \mathbb{E}\left[F_{v_{\hat{\Theta}}}\left(x(1+\mu_{\hat{\Theta}})\right)\right]$, where the expectation is taken over the steady-state distribution of types $\hat{\Theta}$. As a consequence, the maximum competing bid is given by $D = \max\left(\hat{W}_{1:\hat{M}}, r\right)$, where $\hat{W}_{1:\hat{M}}$ is the first order statistic of $\hat{M}$ i.i.d. samples of $\hat{W}$. Its distribution when $x \geq r$ is

$$F_d(x;\boldsymbol{\mu}) = \mathbb{E}\left[F_w(x;\boldsymbol{\mu})^{\hat{M}}\right] = \exp\left\{-\mathbb{E}[\alpha_\Theta\lambda s_\Theta]\bar{F}_w(x;\boldsymbol{\mu})\right\},$$

where we used the fact that bids are independent, that $\hat{M}$ is Poisson with mean $\mathbb{E}[\alpha_\Theta\lambda s_\Theta]$, and the Poisson probability generating function. The result follows by replacing the expression for F_w in the equation above.

ii.) The expenditure function can be written as

$$\begin{aligned} G_\theta(\boldsymbol{\mu}, r) &= \mathbb{E}\left[\mathbf{1}\{(1+\mu_\theta)D \leq V_\theta\}D\right] = \mathbb{E}\left[D\bar{F}_{v_\theta}((1+\mu_\theta)D)\right] \\ &= r\bar{F}_{v_\theta}((1+\mu_\theta)r)F_d(r;\boldsymbol{\mu}) + \int_r^{\bar{V}} x\bar{F}_{v_\theta}((1+\mu_\theta)x)\,\mathrm{d}F_d(x;\boldsymbol{\mu}), \end{aligned}$$

where the second equation follows by the independence of V_θ and D, and the third by recognizing that D is the maximum between the largest bid from advertisers and the reserve price r.

□

Using the previous characterizations we state a set of useful properties of the expenditure function.

Lemma B.2. *i.) For any $\boldsymbol{\mu}$, the maximum bid $D \sim F_d(\boldsymbol{\mu})$ is integrable, that is, $\mathbb{E}[D] < \infty$.*

ii.) For any $\theta \in \Theta$, the expenditure function $G_\theta(\boldsymbol{\mu}, r)$ is differentiable with respect to $\boldsymbol{\mu}$ and r.

iii.) For any $\theta \in \Theta$ and $r \in [\underline{V}, \bar{V}]$, $\partial G_\theta(\boldsymbol{\mu}, r)/\partial \mu_\theta < 0$.

iv.) For any vector of multipliers $\boldsymbol{\mu} \in \mathbb{R}_+^{|\Theta|}$, $\lim_{r\to\infty} G_\theta(\boldsymbol{\mu}, r) = 0$.

v.) For any $r \geq 0$ and vector of multipliers $\boldsymbol{\mu}_{-\theta} \in \mathbb{R}_+^{|\Theta|-1}$, $\lim_{\mu_\theta\to\infty} G_\theta(\boldsymbol{\mu}, r) = 0$.

Proof of Lemma B.2. We prove each item at a time.

i.) Note that $D = \max(\hat{W}_{1:\hat{M}}, r) \leq r + \sum_{k=1}^{\hat{M}} \hat{W}_k$, and that advertisers shade their bids, i.e, $\hat{W}_\theta \leq V_\theta$. Thus,

$$\mathbb{E}[D] \leq r + \mathbb{E}\left[\sum_{k=1}^{\hat{M}} V_{\hat{\Theta}_k}\right] = r + \mathbb{E}[\hat{M}]\mathbb{E}[V_{\hat{\Theta}}] < \infty,$$

where the equality follows from conditioning on the number of matching bidders and using that bids are independent; and the last inequality because $\hat{M}$ is Poisson with mean $\mathbb{E}[\alpha_\Theta \lambda s_\Theta] < \infty$, and the expected valuation satisfies $\mathbb{E}[V_{\hat{\Theta}}] = \sum_\theta \mathbb{P}_{\hat{\Theta}}\{\theta\}\mathbb{E}[V_\theta] < \infty$.

ii.) By Lemma B.1(i), the distribution of the maximum competing bid when $x \geq r$ is given by $F_d(x; \boldsymbol{\mu}) = \exp\left\{-\mathbb{E}[\alpha_\Theta \lambda s_\Theta] \sum_\theta \hat{p}_\theta \bar{F}_{v_\theta}((1+\mu_\theta)x)\right\}$, where $\hat{p}_\theta = \mathbb{P}_{\hat{\Theta}}\{\theta\}$. Since the cumulative distribution of values is differentiable, the distribution of the maximum bid is differentiable w.r.t. x and $\boldsymbol{\mu}$. Indeed, its partial derivatives are given by $\partial F_d/\partial \mu_\theta = F_d(x; \boldsymbol{\mu})\mathbb{E}[\alpha_\Theta \lambda s_\Theta]\hat{p}_\theta x f_{v_\theta}((1+\mu_\theta)x)$, and $\partial F_d/\partial x = F_d(x; \boldsymbol{\mu})\mathbb{E}[\alpha_\Theta \lambda s_\Theta] \sum_\theta \hat{p}_\theta (1+\mu_\theta) f_{v_\theta}((1+\mu_\theta)x)$. Moreover, the second derivatives of the distribution of the maximum bid are continuous because densities $f_{v_\theta}(\cdot)$ are continuously differentiable.

By Lemma B.1(ii), the expenditure function can be written as $G_\theta(\boldsymbol{\mu}, r) = r\bar{F}_{v_\theta}((1+\mu_\theta)r)F_d(r; \boldsymbol{\mu}) + \int_r^{\bar{V}} x\bar{F}_{v_\theta}((1+\mu_\theta)x)\,\mathrm{d}F_d(x; \boldsymbol{\mu})$, which is clearly differentiable in r. Moreover, for any $\theta' \in \boldsymbol{\Theta}$ the first term is differentiable w.r.t. $\mu_{\theta'}$, while the integrand is continuously differentiable. We conclude by an application of Leibniz's integral rule, which holds because $[\underline{V}, \bar{V}] \times U$ is bounded.

iii.) The partial derivative of one first type's expenditure w.r.t. her multiplier is

$$\frac{\partial G_\theta}{\partial \mu_\theta} = (I) + (II)$$

where

$$
\begin{aligned}
(I) &= \frac{\partial}{\partial \mu_\theta}\left(r\bar{F}_{v_\theta}((1+\mu_\theta)r)F_d(r;\boldsymbol{\mu})\right) \\
&= -r^2 f_{v_\theta}((1+\mu_\theta)r)F_d(r;\boldsymbol{\mu}) + r\bar{F}_{v_\theta}((1+\mu_\theta)r)\frac{\partial F_d}{\partial \mu_\theta}(r;\boldsymbol{\mu})
\end{aligned}
$$

and

$$
\begin{aligned}
(II) &= \frac{\partial}{\partial \mu_\theta}\int_r^{\bar{V}} x\bar{F}_{v_\theta}((1+\mu_\theta)x)\frac{\partial F_d}{\partial x}\,\mathrm{d}x \\
&= -\int_r^{\bar{V}} x^2 f_{v_\theta}((1+\mu_\theta)x)\frac{\partial F_d}{\partial x}\,\mathrm{d}x + \int_r^{\bar{V}} x\bar{F}_{v_\theta}((1+\mu_\theta)x)\frac{\partial^2 F_d}{\partial \mu_\theta \partial x}\,\mathrm{d}x \\
&= -\int_r^{\bar{V}} x^2 f_{v_\theta}((1+\mu_\theta)x)\frac{\partial F_d}{\partial x}\,\mathrm{d}x - r\bar{F}_{v_\theta}((1+\mu_\theta)r)\frac{\partial F_d}{\partial \mu_\theta}(r;\boldsymbol{\mu}) \\
&\quad - \int_r^{\bar{V}} \frac{\partial}{\partial x}\left(x\bar{F}_{v_\theta}((1+\mu_\theta)x)\right)\frac{\partial F_d}{\partial \mu_\theta}\,\mathrm{d}x,
\end{aligned}
$$

where the second equality follows from exchanging integration and differentiation, which is valid from item (ii); the third from exchanging partial derivatives by Clairaut's theorem, which holds because the second partial derivatives are continuous almost everywhere; and the last from integrating the second term by parts and using the fact that $\bar{F}_{v_\theta}((1+\mu_\theta)\bar{V}) = 0$. Note that increasing μ_θ decreases the bidder under consideration own bids, but also its competitors' bids of the same type through D. In what follows, we show that these effects are such that the expected expenditure decreases.

In order to simplify the notation, we denote by $f_\theta(x) \triangleq x f_{v_\theta}((1+\mu_\theta)x)$, $\bar{F}_\theta(x) \triangleq \bar{F}_{v_\theta}((1+\mu_\theta)x)$, and by $\langle u, v\rangle \triangleq \int_0^\infty u(x)v(x)w(x)\,dx$ the inner product of two functions u and v with respect to the weight $w(x) \triangleq \mathbb{E}[\alpha_\Theta \lambda s_\Theta]F_d(x;\boldsymbol{\mu})$. Using this new notation and canceling terms we can write the partial derivative as

$$
\frac{\partial G_\theta}{\partial \mu_\theta} = -\sum_{\theta' \neq \theta}(1+\mu_{\theta'})\hat{p}_{\theta'}\langle f_\theta, f_{\theta'}\rangle - \hat{p}_\theta\langle f_\theta, \bar{F}_\theta\rangle - r f_\theta(r)F_d(r;\boldsymbol{\mu}), \tag{8}
$$

which is strictly negative.

iv.) The result follows by noting that $\bar{F}_{v_\theta}((1+\mu_\theta)x) = 0$, for sufficiently large x.

v.) In the homogeneous case we have that $G(\mu, r) = G(0, (1+\mu)r)/(1+\mu)$ and the result follows directly from (iv). In the heterogeneous case when $r > 0$ the result also follows directly. When

$r = 0$ we have that

$$\begin{aligned} G_\theta(\boldsymbol{\mu}, r) &= \mathbb{E}\left[D\bar{F}_{v_\theta}((1+\mu_\theta)D)\mathbf{1}\{\text{only one or more } \theta \text{ type bidders match}\}\right] \\ &\quad + \mathbb{E}\left[D\bar{F}_{v_\theta}((1+\mu_\theta)D)\mathbf{1}\{\text{another type } \Theta' \text{ matches}, D > x\}\right] \\ &\quad + \mathbb{E}\left[D\bar{F}_{v_\theta}((1+\mu_\theta)D)\mathbf{1}\{\text{another type } \Theta' \text{ matches}, D \le x\}\right] \\ &\le \frac{1}{1+\mu_\theta}\mathbb{E}\left[(V_\theta)_{1:\hat{M}}\bar{F}_{v_\theta}((V_\theta)_{1:\hat{M}})\right] + \frac{\mathbb{E}[V_\theta^2]}{x(1+\mu_\theta)^2} + x\mathbb{E}\left[\bar{F}_{v_\theta}\left(\frac{1+\mu_\theta}{1+\mu_{\Theta'}}V_{\Theta'}\right)\right], \end{aligned}$$

where in first term we used that $D = (V_\theta)_{1:\hat{M}}/(1+\mu_\theta)$; the second term follows by Markov's inequality; and the third term because $D \ge V_{\Theta'}/(1+\mu_{\Theta'})$ and $\bar{F}_{v_\theta}(\cdot)$ is non-increasing. The first two terms trivially converge to zero as $\mu_\theta \to \infty$. The third term converges to zero from Dominated Convergence Theorem because $\bar{F}_{v_\theta}(\cdot) \le 1$, and $\lim_{x\to\infty} \bar{F}_{v_\theta}(x) = 0$. □

B.1 Proof of Theorem 4.1

We prove the result in three steps. First, we show that the best-response correspondence can be restricted to a compact set. Second, we prove that the dual objective function is jointly continuous in its arguments. We conclude in the third step.

Step 1. Let $\bar{s} = \max_{\theta\in\Theta} s_\theta$ be the largest possible campaign length, $\bar{\alpha} = \max_{\theta\in\Theta} \alpha_\theta$ be the largest matching probability, $\underline{b} = \min_{\theta\in\Theta} b_\theta$ be the smallest possible budget, and note that $\bar{s}, \bar{\alpha}, \underline{b}$ are positive. We establish that selecting a multiplier outside of $U \triangleq [0, \bar{\mu}]$ with $\bar{\mu} \triangleq \bar{\alpha}\eta\bar{s}\overline{V}/\underline{b}$ is a dominated strategy. To see this notice that for every $\mu > \bar{\mu}$ we have that

$$\Psi_\theta(\mu; \boldsymbol{\mu}) \ge \mu b_\theta > \bar{\mu}\underline{b} = \bar{\alpha}\eta\bar{s}\overline{V} \ge \alpha_\theta\eta s_\theta\overline{V} \ge \Psi_\theta(0; \boldsymbol{\mu}),$$

and thus every $\mu > \bar{\mu}$ in the dual problem is dominated by $\mu = 0$.

Consider the best-response correspondence restricted to U, $\boldsymbol{M} : U^{|\Theta|} \to \mathcal{P}(U^{|\Theta|})$ defined for each type $\theta \in \Theta$ as $M_\theta(\boldsymbol{\mu}) = \arg\min_{\mu\in U} \Psi_\theta(\mu; \boldsymbol{\mu})$. By the above, to establish the existence of a FMFE, it is sufficient to show that $\boldsymbol{M}$ admits a fixed-point, that is, there is some profile of multipliers $\boldsymbol{\mu}^* \in U^{|\Theta|}$ such that $\boldsymbol{\mu}^* \in \boldsymbol{M}(\boldsymbol{\mu}^*)$.

Step 2. Next, we show that for each type $\theta \in \Theta$ the objective function $\Psi_\theta(\mu; \boldsymbol{\mu})$ is jointly continuous in μ and $\boldsymbol{\mu}$. Consider a sequence $(\mu^n, \boldsymbol{\mu}^n) \in U \times U^{|\Theta|}$ converging as $n \to \infty$ to some $(\mu, \boldsymbol{\mu})$ in the set. Notice that under the discreteness of the type space we can write the distribution of bids as $F_w(x; \boldsymbol{\mu}) = \sum_{\theta\in\Theta} \mathbb{P}\{\hat{\Theta} = \theta\} F_{v_\theta}(x(1+\mu_\theta))$. Because the sum is finite and $F_{v_\theta}(\cdot)$ is continuous; we have that $F_w(x; \boldsymbol{\mu}^n) \to F_w(x; \boldsymbol{\mu})$ as $n \to \infty$ for all x. Furthermore, because the distribution F_d of the maximum bid is a continuous function of F_w (cf. Lemma B.1(i)), we get that the same holds for the maximum bid. Denoting by D^n the maximum bid random variable associated to $\boldsymbol{\mu}^n$, by D the maximum bid random variable associated to $\boldsymbol{\mu}$; the previous argument implies that D^n converges in distribution to D. Additionally, by Slutsky's Theorem we get that $(1+\mu^n)D^n \Rightarrow (1+\mu)D$.

Consider the function $\ell(x) = \mathbb{E}[V - x]^+ = \int_x^\infty \bar{F}_v(y)\,\mathrm{d}y$. The function ℓ is bounded by $\mathbb{E}V$ and continuous. Using the fact that valuations are independent and conditioning on the maximum bid, we may write the dual objective as $\Psi_\theta(\mu;\boldsymbol{\mu}) = \alpha\eta s\mathbb{E}\left[\ell\big((1+\mu)D\big)\right] + \mu b$. By portmanteau theorem we have that $\mathbb{E}\left[\ell\big((1+\mu^n)D^n\big)\right] \to \mathbb{E}\left[\ell\big((1+\mu)D\big)\right]$, and thus Ψ is jointly continuous in $(\mu, \boldsymbol{\mu})$.

Step 3. Because the domain is compact, Ψ is jointly continuous in $(\mu, \boldsymbol{\mu})$, and convex in μ for fixed $\boldsymbol{\mu}$ (cf. proof of Proposition 3.1), an FMFE is guaranteed to exist by Proposition 8.D.3 in Mas-Colell et al. (1995).

B.2 Proof of Theorem 4.2

Exploiting the fact that the dual objective is convex and differentiable, one may write the equilibrium condition (3) as a Nonlinear Complementarity Problem (NCP). From the optimality conditions of the dual, it should be the case that for each type $\theta \in \Theta$, one of the following alternatives holds

$$\mu^*_\theta = 0, \frac{\partial \Psi_\theta}{\partial \mu}(\mu^*_\theta, \boldsymbol{\mu}^*) \geq 0,$$
$$\mu^*_\theta > 0, \frac{\partial \Psi_\theta}{\partial \mu}(\mu^*_\theta, \boldsymbol{\mu}^*) = 0.$$

Recall that the derivative of the dual is $\frac{\partial \Psi_\theta}{\partial \mu} = b_\theta - \alpha_\theta \eta s_\theta G_\theta(\boldsymbol{\mu}, r)$, where $\mathbf{G} : \mathbb{R}_+^{|\Theta|} \times \mathbb{R}_+ \to \mathbb{R}_+^{|\Theta|}$ denotes the vector-valued function that maps a profile of multipliers and reserve price to the expected expenditures of each bidder type. Thus, we have that a vector of multipliers $\boldsymbol{\mu}^*$ constitutes a FMFE if it solves the NCP

$$\mu^*_\theta \geq 0 \quad \perp \quad \alpha_\theta \eta s_\theta G_\theta(\boldsymbol{\mu}^*, r) \leq b_\theta, \qquad \forall \theta \in \Theta, \tag{9}$$

where $\perp$ indicates a complementarity condition between the multiplier and the expenditure, that is, at least one condition should be met with equality. From item (ii) of Lemma B.2 we have that the mapping $\mathbf{G}$ is differentiable. The latter, together with the P-matrix assumption, allows one to invoke (Facchinei and Pang, 2003*a*, Proposition 3.5.10) and conclude that the NCP (9) has at most one solution.

B.3 Proof of Proposition 4.1

Fix $r \geq 0$. The existence of the equilibrium follows from Theorem 4.1. The uniqueness follows from the fact that Assumption 4.1 is automatically satisfied in the present case from item *iii.*) of Lemma B.2. We next derive the characterization of the FMFE.

Suppose first that $\alpha\eta s G_0(r) < b$. By Lemma B.2 *iii.*), increasing the multiplier cannot increase the expenditure, and no solution to the NCP with $\mu > 0$ exists. Thus $\mu^* = 0$ is the unique equilibrium multiplier. Suppose now that $\alpha\eta s G_0(r) \geq b$, then advertisers need to shade their bids by picking a non-negative equilibrium multiplier μ^* that solves for $\alpha\eta s G(\mu^*, r) = b$ (and such a solution exists by the proof of Theorem 4.1). Noting that $(1+\mu^*)G(\mu^*, r) = G_0((1+\mu^*)r)$ concludes the proof.

B.4 Proof of Theorem 5.1

The proof proceeds as follows. We first state and prove some basic properties of the publisher's profit function. Second, we characterize the optimal reserve price in the two cases described in the statement of the theorem.

Using (4), the publisher's profit as a function of the reserve price and the multiplier can be written as $\Pi(\mu, r) = \alpha\lambda\eta s G(\mu, r) - \eta c I(\mu, r)$, with $I(\mu, r)$ the probability that the impression is won by some advertiser in the exchange when advertisers employ a multiplier μ and the publisher sets a reserve price r. Note that $I(\mu, r) = I_0((1+\mu)r)$, where by Lemma B.1(i), $I_0(r) = 1 - e^{-\alpha\lambda s\bar{F}_v(r)}$ is the probability that the impression is won in the exchange by truthful advertisers . The publisher's problem amounts to solving $\max_{r\geq 0} \Pi(\mu(r), r)$, where $\mu(r)$ is the unique equilibrium multiplier for price r.

It is simple to show that $r_c^* \geq c$, and that r_c^* (the optimal reserve price of the one-shot auction) is increasing in c; that is, when the opportunity cost increases, the publisher is more inclined to keep the impression, and thus she increases the reserve price. Let $g = b/(\alpha\eta s)$ be the maximum target expenditure per auction of a bidder. We show the following preliminary results:

(i) *The function $\Pi(0, r)$ is quasi-concave in r on $[\underline{V},\ \bar{V}]$, and the maximum is obtained at $r = r_c^*$.* When $\mu = 0$, all advertisers bid truthfully and the auctions decouple; the result then essentially follows by the optimality of r_c^* in a second-price auction with the only caveat that in our setting the number of bidders is random. Formally, one may write the derivative of the profit w.r.t. the reserve price as

$$\frac{\partial \Pi}{\partial r}(0, r) = \alpha\lambda\eta s \left(G_0'(r) + c e^{-\alpha\lambda s\bar{F}_v(r)} f_v(r) \right) = \alpha\lambda\eta s \bar{F}_v(r) e^{-\alpha\lambda s\bar{F}_v(r)} \left(1 - \frac{r-c}{r}\xi(r) \right), \quad (10)$$

where the second equation follows by Lemma B.1(ii). The previous expression vanishes at r_c^*. Notice that the leading terms in the derivative are non-negative, and by the IGFR assumption, it follows that the derivative is non-negative for $r < r_c^*$ and non-positive for $r > r_c^*$. Thus, $\Pi(0, r)$ is strictly quasi-concave on $[\underline{V},\ \bar{V}]$.

(ii) *Then the set $\mathcal{R}^*$ is a closed bounded interval.* The proof follows by noticing that setting $c = 0$ in (10) implies that $G_0(r)$ is strictly quasi-concave in r (in the interval $[\underline{V},\ \bar{V}]$). Since $\mathcal{R}^*$ is an upper-level set of G_0, and G_0 is continuous we get that $\mathcal{R}^*$ is a closed interval. The boundedness of $\mathcal{R}^*$ follows from Lemma B.2 *iv.*).

(iii) *The equilibrium multiplier verifies $\mu > 0$ for r in the interior of $\mathcal{R}^*$, and zero otherwise.* That $\mu = 0$ outside the interior of $\mathcal{R}^*$ follows directly from the statement of Proposition 4.1. By the strict quasi-concavity of $G_0(r)$ in r, $\alpha\eta s G_0(r) > b$ for r in the interior of $\mathcal{R}^*$, so by Proposition 4.1, $\mu > 0$ for r in this set.

(iv) *When $r \in \mathcal{R}^*$ the probability that the impression is won $I(\mu(r), r)$ is decreasing in r.* Write the total derivative of the probability that the impression is won as

$$I'(\mu(r),r) \;=\; I'_\mu(\mu(r),r)\mu'(r) + I'_r(\mu(r),r) = -\alpha\lambda s e^{-\alpha\lambda s \bar{F}_v((1+\mu)r)} f_v((1+\mu)r)\left(\mu' r + 1 + \mu\right),$$

where to simplify the notation we dropped the dependence of r in μ in the second equation. Hence, it suffices to show that $\mu' r + 1 + \mu \geq 0$ to conclude that $I(\mu(r), r)$ is decreasing in r. Since $r \in \mathcal{R}^*$ we have that $G(\mu, r) = g$, and by the implicit function theorem the derivative of the multiplier w.r.t. r is given by $\mu' = -G'_r(\mu, r)/G'_\mu(\mu, r)$. Thus,

$$\mu' r + 1 + \mu = -\frac{G'_r(\mu,r)r - (1+\mu)G'_\mu(\mu,r)}{G'_\mu(\mu,r)} = -\frac{G(\mu,r)}{G'_\mu(\mu,r)} \geq 0,$$

where the second equation follows from the fact that $G'_r(\mu, r) = G'_0((1 + \mu)r)$ and $G'_\mu(\mu, r) = G'_0((1 + \mu)r)r/(1 + \mu) - G(\mu, r)/(1 + \mu)$, and the inequality from the fact that $G(\mu, r) = g \geq 0$, and that $G(\mu, r)$ is decreasing in μ for fixed r by Lemma B.2 *iii.*).

Now, we study the two cases.

Case 1. Suppose that the expenditure at r_c^* does not exceed the budget-per-auction g (i.e., $G_0(r_c^*) < g$), we should show that r_c^* is optimal. If the set $\mathcal{R}^*$ is empty (which occurs when $G_0(r_0^*) < g$, because r_0^* maximizes G_0), then by property (iii) the equilibrium multiplier is $\mu(r) = 0$ for all r, so bidders are truthful for all r. Hence, r_c^* is the optimal reserve price by (i).

Next, assume that the set $\mathcal{R}^*$ is non-empty. By property (ii), the set is compact and thus $\bar{r} = \sup \mathcal{R}^* < \infty$. Moreover, $G_0(\bar{r}) = g$, because $\mathcal{R}^*$ is closed. For prices $r \in \mathcal{R}^*$ we have that $\Pi(\mu(r), r) \leq \Pi(0, \bar{r}) \leq \Pi(0, r_c^*)$. The first inequality follows by the following observation: bidders exhaust their budgets for $r \in \mathcal{R}^*$ (and spend g per auction). Therefore, the reserve price in $\mathcal{R}^*$ that maximizes profits is the one that minimizes the probability of selling an impression. Note that decreasing the reserve price from $\bar{r}$ has two effects: (1) the probability of a sell increases because of the direct effect; and (2) the probability of a sell decreases because of the indirect effect that bidders start shading their equilibrium bids. Property (iv) shows that the direct effect is dominant, and therefore, $\bar{r}$ minimizes the probability of selling an impression within $\mathcal{R}^*$. The second inequality follows from the fact that $\mu(\bar{r}) = 0$ by (iii). Every reserve price $r \notin \mathcal{R}^*$ is dominated by r_c^*. Since in both cases the multipliers are zero and advertisers are truthful, r_c^* is optimal by property (i).

Case 2. Suppose that the expenditure at r_c^* exceeds the maximum expenditure g (i.e., $G_0(r_c^*) \geq g$). Bidders are budget constrained at r_c^* and $r_c^* \in \mathcal{R}^*$. Take any price $r \in \mathcal{R}^*$. As in case 1, property (iv) implies that the profit for any price in that set is dominated by that of $\bar{r}$. Now consider prices strictly greater than those in $\mathcal{R}^*$, that is, those satisfying $r > \bar{r}$, which have $\mu(r) = 0$. From property (i), we have that $\Pi(0, r)$ is non-increasing to the right of r_c^*. Because $r_c^* \leq \bar{r} \leq r$, we have that $\Pi(0, \bar{r}) \geq \Pi(0, r)$. Hence, every reserve price $r > \bar{r}$ is dominated by $\bar{r}$. A similar argument holds for prices strictly less than those in $\mathcal{R}^*$ and the optimality of $\bar{r}$ follows.

B.5 Proof of Proposition A.2

Fix an arbitrary policy β. The result is proven in two steps. First, we upper bound the performance of the policy β by the performance of a policy with the benefit of hindsight, denoted by β^{H}, which assumes complete knowledge of the future realizations of bids and values. Second, we upper bound the performance of β^{H} by the dual objective function.

Let $J_k(\beta^{\mathrm{H}}, \boldsymbol{\beta}^{\mathrm{F}}_{-k})$ denote the expected payoff under perfect hindsight, which is obtained by looking at the optimal expected payoff when the realization of the number of impressions, the matching indicators and the values of all advertisers for the whole horizon are revealed up-front. No strategy can perform better than the perfect hindsight strategy β^{H} and we have that

$$J_k(\beta, \boldsymbol{\beta}^{\mathrm{F}}_{-k}) \leq J_k(\beta^{\mathrm{H}}; \boldsymbol{\beta}^{\mathrm{F}}_{-k}).$$

Let $\tilde{N}$ be the first auction in which some advertiser runs out of budget when the k^{th} advertiser implements the hindsight policy, and $J^{\mathrm{I}}_k(\beta^{\mathrm{I,H}}, \boldsymbol{\beta}^{\mathrm{F}}_{-k})$ denote the expected payoff under perfect hindsight in an alternate system (I) in which advertisers are allowed to bid (i) when they have no budget, and (ii) after the end of their campaigns. Note that the hindsight policy in the alternate system (I), denoted by $\beta^{\mathrm{I,H}}$ is potentially different to the hindsight policy for the original one. We can bound the performance of the hindsight policy by

$$\begin{aligned} J_k(\beta^{\mathrm{H}}, \boldsymbol{\beta}^{\mathrm{F}}_{-k}) &= \mathbb{E}\left[\sum_{n=1}^{N(s)} U_{n,k}(\beta^{\mathrm{H}})\right] \leq \mathbb{E}\left[\sum_{n=1}^{N(s)\wedge \tilde{N}} U_{n,k}(\beta^{\mathrm{H}})\right] + \bar{V}\mathbb{E}\left[\sum_{n=1}^{N(s)} m_{n,k} - \sum_{n=1}^{\tilde{N}} m_{n,k}\right]^{+} \\ &\leq \mathbb{E}\left[\sum_{n=1}^{N(s)} U^{\mathrm{I}}_{n,k}(\beta^{\mathrm{H}})\right] + \bar{V}\mathbb{E}[I_k - \tilde{I}_k]^{+} \\ &\leq J^{\mathrm{I}}_k(\beta^{\mathrm{I,H}}, \boldsymbol{\beta}^{\mathrm{F}}_{-k}) + \bar{V}\mathbb{E}[I_k - \alpha_k \eta s]^{+} + \mathbb{E}[\alpha_k \eta s - \tilde{I}_k]^{+} \end{aligned}$$

where $U_{n,k}(\beta^{\mathrm{H}})$ and $U^{\mathrm{I}}_{n,k}(\beta^{\mathrm{H}})$ denote the realized utility under the hindsight policy in the original and alternate system (I), respectively; and I_k denotes the number of auctions that advertiser k^{th} participates during his campaign, that is, $I_k = \sum_{n=1}^{N(s)} M_{n,k}$; and $\tilde{I}_k$ denotes the number of auctions that advertiser k^{th} participates until some agent runs out of budget, that is, $\tilde{I}_k = \sum_{n=1}^{\tilde{N}} M_{n,k}$. The first inequality follows from the fact that $0 \leq U_{n,k}(\beta^{\mathrm{H}}) \leq M_{n,k}\bar{V}$. The second from the fact the alternate system (I) and the original one coincide until the $\tilde{N}$-th auction, and adding the utility on the alternate system (I) obtained after the $\tilde{N}$-th auction only increases the right-hand side. The third from the fact $\beta^{\mathrm{I,H}}$ is the optimal policy in the alternate system (I) and that for every $a, b, c \in \mathbb{R}$ we have that $(a-c)^{+} \leq (a-b)^{+} + (b-c)^{+}$. In the remainder of the proof we address one term at a time.

Term 1. We proceed to bound the performance of the policy $\beta^{\mathrm{I,H}}$ in the alternate system (I). Note that in this system all advertisers bid regardless of the budget. Hence the k^{th} advertiser can not strategize to deplete the budgets of her competitors. Given a sample path ω, which determines the number of impressions $N(s)(\omega) = N$, the matching indicators $\{M_{n,k}(\omega)\}_{n=1}^{N(s)(\omega)} = \{m_{n,k}\}_{n=1}^{N}$, and the realization of the competing bids and values $\{(D_{n,-k}(\omega), V_{n,k}(\omega))\}_{n=1}^{N(s)(\omega)} = \{(d_{n,-k}, v_{n,k})\}_{n=1}^{N}$; the advertiser only needs to determine which auctions to win (since bidding an amount $\epsilon > 0$ larger than the maximum bid guarantees her winning the auction). Let the decision variable $x_n \in \{0,1\}$ indicate whether the k^{th} advertiser decides to wins the auction or not. In hindsight, the zeroth advertiser needs to solve, for each realization ω, the following knapsack problem

$$J_k^{\text{I,H}}(\omega) = \max_{x_n \in \{0,1\}} \sum_{n=1}^{N} x_n (v_{n,k} - d_{n,-k}) \tag{11a}$$

$$\text{s.t. } \sum_{n=1}^{n_\theta} x_n d_{n,-k} \leq b_k, \tag{11b}$$

$$x_n \leq m_{n,k}. \tag{11c}$$

The perfect hindsight bound is obtained by averaging over all possible realizations consistently with the strategy of the other bidders, or equivalently $J_k^{\text{I}}(\beta^{\text{I,H}}, \boldsymbol{\beta}^{\text{F}}_{-k}) = \mathbb{E}_\omega \left[J_k^{\text{I,H}}(\omega) \right]$.

Consider the continuous relaxation of the hindsight program (11) in which we replace the integrality constraints by $0 \leq x_n \leq m_{n,k}$. Let μ_k be the equilibrium multiplier of the FMFE for k^{th} advertiser. Introducing dual variables $\mu \geq 0$ for the budget constraint and $z_n \geq 0$ for the constraints $x_n \leq m_{n,k}$, we get by weak duality that

$$\begin{aligned} J_k^{\text{I,H}}(\omega) &\leq \min_{\mu \geq 0, z_n \geq 0} \left\{ \sum_{n=1}^{N} m_{n,k} z_n + \mu b_k \text{ s.t. } z_n \geq v_{n,k} - (1+\mu) d_{n,-k}, \forall n = 1, \ldots, N \right\} \\ &= \min_{\mu \geq 0} \left\{ \sum_{n=1}^{N} m_{n,k} [v_{n,k} - (1+\mu) d_{n,-k}]^+ + \mu b_k \right\} \\ &\leq \sum_{n=1}^{N} m_{n,k} [v_{n,k} - (1+\mu_k) d_{n,-k}]^+ + \mu_k b_k \end{aligned}$$

where the equality follows from the fact that in the optimal solution of the dual problem it is either the case that $z_n = 0$ or $z_n = v_{n,k} - (1+\mu) d_{n,-k}$, and the second inequality from the fact that μ_k is not necessarily optimal for the hindsight program. Taking expectations and using the fact that the number of matching impressions is Poisson with mean $\alpha_k \eta s$ independently of values and competing bids, we get that

$$J_k^{\text{I}}(\beta^{\text{I,H}}, \boldsymbol{\beta}^{\text{F}}_{-k}) \leq J_k^{\text{F}}.$$

Term 2. Using the same argument that in the proof of Proposition A.1 we obtain that

$$\frac{1}{\alpha_k \eta s} \mathbb{E}[I_k - \alpha_k \eta s]^+ \leq (2 \alpha_k \eta s)^{-1/2} = O\left((\alpha_k \eta s)^{-1/2} \right).$$

Term 3. In order to bound the term $\mathbb{E}[\alpha_k \eta s - \tilde{I}_k]^+$ we shall consider a second alternate system (II) in which all advertisers (including k^{th}) implement the FMFE strategies and the initial budgets for every advertiser are discounted in advance to take into account the potential impact that the k^{th} advertiser may have on its competitors. Since a competitor can spend at most $\bar{V}$ in an auction, this potential impact can be upper bounded by $\bar{V}$ times the number of auctions in which she competes against the k^{th} advertiser. That is, we set the budgets to $b_i^{\text{II}} = \big(b_i - \bar{V} T_{k,i}\big)^+$ for all $i \neq k$ where $T_{k,i} = \sum_{n=1}^{N(s)} M_{n,k} M_{n,i}$ is the number of auctions in which k and i compete together. Defining by $\tilde{N}^{\text{II}}$ be the first auction in which some advertiser runs out of budget in the alternate system (II), we obtain using a coupling argument that

$$\tilde{N}^{\text{II}} \leq \tilde{N} \quad \text{(a.s.)}. \tag{12}$$

Next we proceed to bound the number of auctions in the left-over regime after time $\tilde{N}$ using the alternate system (II).

Let $\tilde{I}_{k,i}^{\text{II}}$ be the number of auctions that advertiser k^{th} participates until agent i^{th} runs out of budget, and $\tilde{I}_k^{\text{II}}$ be the number of auctions the k^{th} advertiser participates until someone runs out of budget. Equation (12) implies that $\tilde{I}_k^{\text{II}} \leq \tilde{I}_k$ almost surely, which implies using the steps in the proof of Proposition A.1 that

$$\begin{aligned} \mathbb{E}\left[\alpha_k \eta s - \tilde{I}_k\right]^+ \leq \mathbb{E}\left[\alpha_k \eta s - \tilde{I}_k^{\text{II}}\right]^+ &\leq \mathbb{E}\left[\max_i \{\alpha_k \eta s - \tilde{I}_{k,i}^{\text{II}}\}^+\right] \\ &\leq \max_i \left\{\alpha_k \eta s - \mathbb{E}\tilde{I}_{k,i}^{\text{II}}\right\}^+ + \sqrt{\sum_i \operatorname{Var}[\tilde{I}_{k,i}^{\text{II}}]}, \end{aligned} \tag{13}$$

where the inequality follows (again) from the upper bound on the maximum of random variables given in Aven (1985). We now proceed to bound the mean and variance of the stopping times $\tilde{I}_{k,i}^{\text{II}}$ by conditioning on the initial budgets.

For the mean we obtain that

$$\begin{aligned} \mathbb{E}[\tilde{I}_{k,i}^{\text{II}}] = \mathbb{E}\Big[\mathbb{E}[\tilde{I}_{k,i}^{\text{II}} \mid b_i^{\text{II}}]\Big] &\geq \frac{\alpha_k}{\alpha_i G_i(\boldsymbol{\mu})} \mathbb{E}[b_i^{\text{II}}] = \frac{\alpha_k}{\alpha_i G_i(\boldsymbol{\mu})} \left(b_i - \bar{V}\mathbb{E}[T_{k,i}]\right) \\ &= \frac{\alpha_k}{\alpha_i G_i(\boldsymbol{\mu})} \left(b_i - \bar{V}\alpha_k \alpha_i \eta s\right), \end{aligned}$$

where the inequality follows from property (v) of Lemma B.3, and the last equality from the fact that $T_{k,i}$ is Poisson with mean $\alpha_k \alpha_i \eta s$.

For the variance we employ the conditional variance formula to obtain that

$$\operatorname{Var}[\tilde{I}_{k,i}^{\text{II}}] = \mathbb{E}\Big[\operatorname{Var}[\tilde{I}_{k,i}^{\text{II}} \mid b_i^{\text{II}}]\Big] + \operatorname{Var}\Big[\mathbb{E}[\tilde{I}_{k,i}^{\text{II}} \mid b_i^{\text{II}}]\Big].$$

For the first term we use that from property (vi) of Lemma B.3 there exists non-negative constants

C_0, C_1 such that $\mathrm{Var}[\tilde{I}^{\mathrm{II}}_{k,i} \mid b^{\mathrm{II}}_i] \le C_0 + C_1 b^{\mathrm{II}}_i$, together with the fact that $b^{\mathrm{II}}_i \le b_i$ to obtain that

$$\mathbb{E}\Big[\mathrm{Var}[\tilde{I}^{\mathrm{II}}_{k,i} \mid b^{\mathrm{II}}_i]\Big] \le \mathbb{E}[C_0 + C_1 b^{\mathrm{II}}_i] = C_0 + C_1 \mathbb{E}[b^{\mathrm{II}}_i] \le C_0 + C_1 b_i = O(b_i).$$

For the second term we combine the upper and lower bounds on property (v) of Lemma B.3 to obtain that there exists some different non-negative constants C_0, C_1 such that $\big|\mathbb{E}[\tilde{I}^{\mathrm{II}}_{k,i} \mid b^{\mathrm{II}}_i] - C_1 b^{\mathrm{II}}_i\big| \le C_0$. Together with Lemma B.4 we obtain that

$$\begin{aligned}
\sqrt{\mathrm{Var}\Big[\mathbb{E}[\tilde{I}^{\mathrm{II}}_{k,i} \mid b^{\mathrm{II}}_i]\Big]} &= \Big\|\mathbb{E}[\tilde{I}^{\mathrm{II}}_{k,i} \mid b^{\mathrm{II}}_i] - \mathbb{E}[\tilde{I}^{\mathrm{II}}_{k,i}]\Big\|_2 \\
&\le \Big\|\mathbb{E}[\tilde{I}^{\mathrm{II}}_{k,i} \mid b^{\mathrm{II}}_i] - C_1 b^{\mathrm{II}}_i\Big\|_2 + \Big\|C_1 b^{\mathrm{II}}_i - C_1 \mathbb{E}[b^{\mathrm{II}}_i]\Big\|_2 \\
&\le C_0 + C_1 \sqrt{\mathrm{Var}[b^{\mathrm{II}}_i]} \\
&\le C_0 + C_1 \sqrt{\mathrm{Var}[T_{k,i}]} = C_0 + C_1 \sqrt{\alpha_k \alpha_i \eta s} = O(\sqrt{b_i}),
\end{aligned}$$

where the third inequality follows from the fact that truncation reduces variance, that is, for any random variable X and constant x we have that $\mathrm{Var}(x - X)^+ \le \mathrm{Var} X$ (see, e.g., Liu and Li (2009)); and last inequality follows from the fact that $\alpha_k \le 1$ and $\alpha_i \eta s = O(b_i)$ from Assumption 6.1. Combining the bounds for the first and second terms we get that $\mathrm{Var}[\tilde{I}^{\mathrm{II}}_{k,i}] = O(b_i)$.

We put everything together by plugging in our bounds for the mean and variance of the stopping times in the main bound (13) and dividing by the expected number of impressions in the horizon to obtain that

$$\begin{aligned}
\frac{1}{\alpha_k \eta s}\mathbb{E}[\alpha_k \eta s - \tilde{I}_k]^+ &\le \max_i \left\{1 - \frac{b_i - \bar{V}\alpha_k \alpha_i \eta s}{\alpha_i \eta s G_i(\boldsymbol{\mu})}\right\}^+ + \frac{1}{\alpha_k \eta s}\sqrt{\sum_{i=1}^K O(b_i)} \\
&\le \max_i \left\{1 - \frac{b_i}{\alpha_i \eta s G_i(\boldsymbol{\mu})}\right\}^+ + \max_i \left\{\frac{\bar{V}\alpha_k}{G_i(\boldsymbol{\mu})}\right\} + \frac{O\left(\sqrt{K\bar{b}}\right)}{\alpha_k \eta s} \\
&= O\left(\alpha_k + (\alpha_k \eta s)^{-1} K^{1/2} \bar{b}^{1/2}\right) = O\left(\alpha_k + (\alpha_k \eta s)^{-1/2} K^{1/2}\right),
\end{aligned}$$

where the second inequality follows from the fact that the maximum of a sum is dominated by the sum of the maximums and by setting $\bar{b} = \max_i b_i$, the third inequality because expected expenditure in the FMFE never exceeds the budget, that is, $\alpha_i \eta s G_i(\boldsymbol{\mu}) \le b_i$, and because the second term is $O(\alpha_k)$ because the expected expenditure is bounded from below from Assumption 6.1; and the last because $\alpha_k \eta s = O(\bar{b})$ from Assumption 6.1 too.

B.6 Additional Results

Lemma B.3 (Identities and Bounds for Stopping Times)**.** *Suppose that Assumption 6.1 holds. We have that* *(i)* $\frac{b_k}{G_k(\boldsymbol{\mu})} \leq \mathbb{E}[\tilde{I}_{k,k}] \leq \frac{b_k+\bar{V}}{G_k(\boldsymbol{\mu})}$ *for all advertiser* k,

(ii) $Var[\tilde{I}_{k,k}] = O(b_k)$ *for all advertiser* k,

(iii) $\mathbb{E}[\tilde{N}_k] = \alpha_k^{-1}\mathbb{E}[\tilde{I}_{k,k}]$ *for all advertiser* k,

(iv) $Var[\tilde{N}_k] = O(\alpha_k^{-2}b_k)$ *for all advertiser* k,

(v) $\frac{\alpha_k}{\alpha_i}\frac{b_i}{G_i(\boldsymbol{\mu})} \leq \mathbb{E}[\tilde{I}_{k,i}] \leq \frac{\alpha_k}{\alpha_i}\frac{b_i+\bar{V}}{G_i(\boldsymbol{\mu})}$ *for all pair of advertisers* $k \neq i$, *and*

(vi) $Var[\tilde{I}_{k,i}] = O(b_i)$ *for all pair of advertisers* $k \neq i$.

(vii) The expected expenditure per auction in the FMFE is uniformly bounded from below across advertisers, i.e., for all advertiser k *we have that* $G_k(\boldsymbol{\mu}) \geq \underline{z}'$ *for some* z'.

Proof. In order to study the hitting time we consider the sequence $\{Z'_{n,k}\}_{n\geq 1}$ of expenditures of the k^{th} advertiser for the auctions she *participates in* (here we are restricting ourselves to the auctions in which $m_{n,k} = 1$). In view of our mean-field assumption the sequence of expenditures is i.i.d. and independent of the impressions' inter-arrival times. Let $C_{n,k} = \sum_{j=1}^{n} Z'_{j,k}$ denote the cumulative expenditure incurred by advertiser k after the n^{th} auction she *participates in.*

Item (i). Since expenditures are bounded, $Z_{n,k} \leq \bar{V} < \infty$ a.s., the cumulative expenditure at the stopping time can be bounded from below and above by

$$b_k \leq C_{\tilde{I}_{k,k},k} \leq b_k + \bar{V}.$$

Note that from Item (vii) with positive probability the advertiser spends a positive amount and thus $\mathbb{E}\tilde{I}_{k,k} < \infty$. Hence, we may employ Wald's identities to bound the mean and variance of the stopping time $\tilde{I}_{k,k}$. In particular, Wald's first identity implies that $\mathbb{E}[C_{\tilde{I}_{k,k},k}] = \mathbb{E}\tilde{I}_{k,k}\mathbb{E}Z'_k$ with Z'_k in shorthand for $Z'_{1,k}$. Using the fact that $C_{\tilde{I}_{k,k},k} \geq b_k$, one obtains that the mean is bounded from below by $\mathbb{E}[\tilde{I}_{k,k}] \geq b_k/\mathbb{E}[Z'_k]$. Using the fact that $C_{\tilde{I}_{k,k},k} \leq b_k + \bar{V}$, one may also bound the mean from above by $\mathbb{E}[\tilde{I}_{k,k}] \leq (b_k + \bar{V})/\mathbb{E}[Z'_k]$. The result follows from the fact that $\mathbb{E}[Z'_k] = \mathbb{E}[Z_{1,k}] = G_k(\boldsymbol{\mu})$.

Item (ii). The variance is bounded from above by $\text{Var}(\tilde{I}_{k,k}) \leq (b_k+\bar{V})\text{Var}(Z'_k)/\mathbb{E}[Z'_k]^3+\bar{V}^2/\mathbb{E}[Z'_k]^2$ (use Wald's second identity to get $\mathbb{E}[C_{\tilde{I}_{k,k},k} - \tilde{I}_{k,k}\mathbb{E}Z'_k]^2 = \text{Var}(Z'_k)\mathbb{E}\tilde{I}_{k,k}$). The result follows because expenditures are bounded from above by V and because expected expenditures are bounded from below by Assumption 6.1.

Items (iii) and (iv). Recall that $\tilde{N}_k$ is a sum of a random number $\tilde{I}_{k,k}$ of independent geometric random variables with success probability α_k. Thus, we obtain by taking conditional expectations that $\mathbb{E}[\tilde{N}_k] = \alpha_k^{-1}\mathbb{E}[\tilde{I}_{k,k}]$, and $\text{Var}[\tilde{N}_k] = (1-\alpha_k)\alpha_k^{-2}\mathbb{E}[\tilde{I}_{k,k}] + \alpha_k^{-2}\text{Var}[\tilde{I}_{k,k}]$ (see, e.g., Ross (1996, pp.22)).

Items (v) and (vi). Recall that $\tilde{I}_{k,i} = \sum_{n=1}^{\tilde{N}_i} m_{n,k}$ is the number of auctions that advertiser k^{th} participates until agent i^{th} runs out of budget. For the bound on the mean we use Wald's Inequality to obtain that $\mathbb{E}[\tilde{I}_{k,i}] = \alpha_k\mathbb{E}[\tilde{N}_i]$, and the result follows from properties (i) and (iii) of this lemma.

For the bound on the variance we use Wald's Inequality to obtain that $\mathbb{E}[\tilde{I}_{k,i}] = \alpha_k \mathbb{E}[\tilde{N}_i]$ and denote by $\|X\|_2 = \sqrt{\mathbb{E}[X^2]}$ the L_2 norm to obtain that

$$\begin{aligned}
\sqrt{\mathrm{Var}[\tilde{I}_{k,i}]} &= \left\|\tilde{I}_{k,i} - \alpha_k \mathbb{E}[\tilde{N}_i]\right\|_2 = \left\|\tilde{I}_{k,i} - \alpha_k \tilde{N}_i + \alpha_k \tilde{N}_i - \alpha_k \mathbb{E}[\tilde{N}_i]\right\|_2 \\
&\leq \left\|\tilde{I}_{k,i} - \alpha_k \tilde{N}_i\right\|_2 + \left\|\alpha_k \tilde{N}_i - \alpha_k \mathbb{E}[\tilde{N}_i]\right\|_2 = \sqrt{\alpha_k(1-\alpha_k)\mathbb{E}[\tilde{N}_i]} + \alpha_k \sqrt{\mathrm{Var}[\tilde{N}_i]} \\
&= O(\sqrt{\alpha_k \alpha_i^{-1} b_i}) + O(\sqrt{\alpha_k^2 \alpha_i^{-2} b_i}) = O(\sqrt{b_i}),
\end{aligned}$$

where the first inequality follows from Minkowski's inequality, the third equality follows from Wald's second identity ($\mathbb{E}[\tilde{I}_{k,i} - \alpha_k \tilde{N}_i]^2 = \alpha_k(1-\alpha_k)\mathbb{E}\tilde{N}_i$) and the definition of variance; and the last bounds from items (iii) and (iv) from this lemma and Assumption 6.1's restriction of matching probabilities.

Item (vii). Note that when $\mu_k > 0$ we have by the FMFE characterization that the advertiser is budget constrained and thus $G_k(\boldsymbol{\mu}) = b_k/(\alpha_k \eta s) \geq \underline{g}$ by Assumption 6.1. Next, we show that the expenditure is lower bounded when the advertiser is not shading her bids. Recall from the proof of Theorem 4.1 that FMFE multipliers are upper bounded by $\mu_k \leq \alpha_k \eta s \overline{V}/b_k \leq \underline{g}\overline{V}$ with the second inequality by Assumption 6.1. Let $D_{-k}(\boldsymbol{\nu}_{-k}) = \max_{i \neq k, M_i = 1}(V_i/(1+\nu_i)) \vee r$ be the maximum competing bid observed by the k^{th} advertiser when competitors shade their bids according to $\boldsymbol{\nu}_{-k}$. The expected expenditure of the k^{th} advertiser is bounded from below by

$$\begin{aligned}
G_k(\boldsymbol{\mu}) &= \mathbb{E}\left[\mathbf{1}\{D_{-k}(\boldsymbol{\mu}_{-k}) \leq V_k\} D_{-k}(\boldsymbol{\mu}_{-k})\right] \geq \frac{1}{1+\underline{g}\overline{V}} \mathbb{E}\left[\mathbf{1}\{D_{-k}(\boldsymbol{\mu}_{-k}) \leq V_k\} D_{-k}(\mathbf{0})\right] \\
&\geq \frac{1}{1+\underline{g}\overline{V}} \mathbb{E}\left[\mathbf{1}\{D_{-k}(\mathbf{0}) \leq V_k\} D_{-k}(\mathbf{0})\right] = \frac{G_k(\mathbf{0})}{1+\underline{g}\overline{V}} \geq \frac{\underline{z}}{1+\underline{g}\overline{V}},
\end{aligned}$$

where the first inequality follows from $D_{-k}(\boldsymbol{\mu}_{-k}) \geq D_{-k}(\mathbf{0})/(1+\underline{g}\overline{V})$, and the second because the probability that the advertiser wins is lower when competitors do not shade their bids. □

Lemma B.4. *Let X and Y be two random variables, then $\|X - \mathbb{E}X\|_2 \leq \|X - Y\|_2 + \|Y - \mathbb{E}Y\|_2$.*

Proof. By adding and subtracting the difference $Y - \mathbb{E}Y$ we obtain that

$$\begin{aligned}
\mathbb{E}[X - \mathbb{E}X]^2 &= \mathbb{E}[(X - Y) + (Y - \mathbb{E}Y) + (\mathbb{E}Y - \mathbb{E}X)]^2 \\
&= \mathbb{E}[X - Y]^2 + \mathbb{E}[Y - \mathbb{E}Y]^2 - (\mathbb{E}Y - \mathbb{E}X)^2 + 2\mathbb{E}[(X - Y)(Y - \mathbb{E}Y)] \\
&\leq \mathbb{E}[X - Y]^2 + \mathbb{E}[Y - \mathbb{E}Y]^2 + 2\sqrt{\mathbb{E}[X - Y]^2 \mathbb{E}[Y - \mathbb{E}Y]^2} \\
&= (\|X - Y\|_2 + \|Y - \mathbb{E}Y\|_2)^2,
\end{aligned}$$

where the second equality follows from taking expectations and canceling terms, first inequality by Cauchy-Schwarz and dropping the negative term, and the last equality from completing squares. □

C Sufficient Conditions for P-matrix Assumption to Hold

We establish here sufficient conditions for Assumption 4.1, that was required for uniqueness of a FMFE, to hold.

Proposition C.1. *The P-matrix condition (Assumption 4.1) holds in either of the following cases.*

i.) $\mathbf{\Theta}$ *is a singleton.*

ii.) $\mathbf{\Theta}$ *contains two types, and these have a common value distribution with positively homogeneous failure rate.*

The *positively homogeneous* condition in *ii.)* imposes that there is some $n \geq 0$ such that $h_v(ax) = a^n h_v(x)$ for all $x \in \text{dom}(V)$ and $a > 0$. This property is satisfied by distributions whose failure rates are power functions; such as the exponential, Weibull, and Rayleigh distributions. Additionally, it is not difficult to show from first principles that, for the case of two types with common value distribution, Assumption 4.1 holds when values are uniformly distributed with support $[0, \bar{V}]$.

Proof of Proposition C.1. We denote by $J_{\boldsymbol{H}}$ the Jacobian of vector-valued function $\boldsymbol{H} : \mathbb{R}^{|\Theta|} \to \mathbb{R}^{|\Theta|}$. A matrix $A \in \mathbb{R}^{|\Theta| \times |\Theta|}$ is a *P-matrix* if the determinant of all its principals minors is positive, i.e., $\det(A|_T) > 0$ for all $T \subseteq \Theta$, where $A|_T$ denotes the submatrix of A restricted to the indices in T.

i.) In this case $J_{\mathbf{G}} = \partial G(\mu, r)/\partial\mu$, and the result follows directly from item *iii.)* of Lemma B.2.

ii.) We prove the result in two steps. First, we characterize the entries of the Jacobian $J_{\mathbf{G}}$. Second, we show that the Jacobian $J_{-\mathbf{G}} = -J_{\mathbf{G}}$ is a P-matrix.

Step 1. In the proof of item *iii.)* from Lemma B.2 we characterized the diagonal entries of the Jacobian, that is, $\partial G_\theta(\boldsymbol{\mu}, r)/\partial\mu_\theta$. Using a similar notation, we characterize the off-diagonal entries as follows.

We have that the partial derivative of the type θ expenditure w.r.t. the multiplier of type θ' is

$$\begin{aligned}
\frac{\partial G_\theta}{\partial \mu_{\theta'}} &= r\bar{F}_v((1+\mu_\theta)r)\frac{\partial F_d}{\partial \mu_{\theta'}} + \frac{\partial}{\partial \mu_{\theta'}} \int_r^{\bar{V}} x\bar{F}_v((1+\mu_\theta)x)\frac{\partial F_d}{\partial x}\,\mathrm{d}x \\
&= r\bar{F}_v((1+\mu_\theta)r)\frac{\partial F_d}{\partial \mu_{\theta'}} + \int_r^{\bar{V}} x\bar{F}_v((1+\mu_\theta)x)\frac{\partial^2 F_d}{\partial \mu_{\theta'}\partial x}\,\mathrm{d}x \\
&= -\int_r^{\bar{V}} \frac{\partial}{\partial x}\left(x\bar{F}_v((1+\mu_\theta)x)\right)\frac{\partial F_d}{\partial \mu_{\theta'}}\,\mathrm{d}x \\
&= (1+\mu_\theta)\hat{p}_{\theta'}\langle f_\theta f_{\theta'}\rangle - \hat{p}_{\theta'}\langle f_{\theta'}\bar{F}_\theta\rangle,
\end{aligned} \tag{14}$$

where the second equality follows from exchanging integration and differentiation; and the third from exchanging partial derivatives by Clairaut's theorem, integrating by parts, and canceling terms.

Step 2. Next, we show that the Jacobian matrix of $-\mathbf{G}$ is a P-matrix. We denote by 1 the low-type and by 2 the high-type. The Jacobian of $\mathbf{G}$ is given by

$$J_{\mathbf{G}} = \begin{pmatrix} \frac{\partial G_1}{\partial \mu_1} & \frac{\partial G_1}{\partial \mu_2} \\ \frac{\partial G_2}{\partial \mu_1} & \frac{\partial G_2}{\partial \mu_2} \end{pmatrix}.$$

From item *iii.)* of Lemma B.2 one concludes that the principal minors $J|_{\{1\}}$ and $J|_{\{2\}}$ are negative (they are, in fact, negative scalars), so the corresponding principal minors of $-\mathbf{G}$ are positive. The determinant of the remaining minor $J|_{\{1,2\}}$ is that of the whole Jacobian, which is given by

$$\begin{aligned}
\det(J) &= \frac{\partial G_1}{\partial \mu_1}\frac{\partial G_2}{\partial \mu_2} - \frac{\partial G_1}{\partial \mu_2}\frac{\partial G_2}{\partial \mu_1} \\
&= (1+\mu_1)\hat{p}_1^2\langle f_1 f_2\rangle\langle f_1 \bar{F}_1\rangle + (1+\mu_1)\hat{p}_1\hat{p}_2\langle f_1 f_2\rangle\langle f_1 \bar{F}_2\rangle \\
&+ (1+\mu_2)\hat{p}_1\hat{p}_2\langle f_1 f_2\rangle\langle f_2 \bar{F}_1\rangle + (1+\mu_2)\hat{p}_2^2\langle f_1 f_2\rangle\langle f_2 \bar{F}_2\rangle \\
&+ \hat{p}_1\hat{p}_2\langle f_1 \bar{F}_1\rangle\langle f_2 \bar{F}_2\rangle - \hat{p}_1\hat{p}_2\langle f_1 \bar{F}_2\rangle\langle f_2 \bar{F}_1\rangle,
\end{aligned}$$

where the third equation follows from substituting the expressions for the partial derivatives and canceling two terms (here we assumed, without loss of generality, that $r = 0$ since the sum of a positive diagonal matrix with a P-matrix is a P-matrix). Notice that all terms are positive with the exception of the last one. We conclude that the determinant is positive by showing that the fifth term dominates the last one. From positively homogeneous assumption we can write $f_i(x) = x f_v((1+\mu_i)x) = x h_v((1+\mu_i)x)\bar{F}_v((1+\mu_i)x) = (1+\mu_i)^n x h_v(x)\bar{F}_i(x)$. Defining a new weight function $\tilde{w}(x) = x h_v(x) w(x)$ and using Cauchy-Schwartz inequality one gets that

$$\begin{aligned}
\langle f_1 \bar{F}_1\rangle\langle f_2 \bar{F}_2\rangle &= (1+\mu_1)^n(1+\mu_2)^n\langle \bar{F}_1\bar{F}_1\rangle_{\tilde{w}}\langle \bar{F}_2\bar{F}_2\rangle_{\tilde{w}} \\
&\geq (1+\mu_1)^n(1+\mu_2)^n\langle \bar{F}_1\bar{F}_2\rangle_{\tilde{w}}\langle \bar{F}_1\bar{F}_2\rangle_{\tilde{w}} = \langle f_1 \bar{F}_2\rangle\langle f_2 \bar{F}_1\rangle.
\end{aligned}$$

Hence, the corresponding principal minor of $-\mathbf{G}$ is also positive and the result follows. □

D Fluid Model

In this section, we prove Theorem 6.2. We first state a corollary for exponential and uniform valuations.

Corollary D.1. *Suppose that the reserve price is zero and that $V_k^{\mathrm{D}} > V_k^{\mathrm{FMFE}}$. Then*

i.) *If valuations are exponentially distributed, then the optimal strategy $\bar{w}(v)$ is given by*

$$\bar{w}(v) = \frac{v - \mathbb{E}V(K-1)^{-1}(1+\mu_{-k})^{-1}p_{-k}^*}{1 + p_k^* - (K-1)^{-1}p_{-k}^*}.$$

ii.) If valuations are uniform $[0, \bar{V}]$, then the optimal strategy $\bar{w}(v)$ is given by

$$\bar{w}(v) = \frac{v - \bar{V}(K-1)^{-1}(1+\mu_{-k})^{-1}p^*_{-k}}{1 + p^*_k - 2(K-1)^{-1}p^*_{-k}} \,.$$

Hence, it follows, that in the cases of uniform and exponential distributions, one may restrict attention to affine bidding functions when searching for a best response.[17] Now, we provide the proofs of both results.

Proof of Theorem 6.2. Let $\mathbf{B}$ be the space of control policies defined at the beginning of this section. We will divide the space of controls based on the depletion times they induce on the competitors. For this purpose, let $\bar{\tau} = \inf\{t \in [0, s] : x_{-k}(t) = 0\}$ where we let $\bar{\tau} = \infty$ if the set $\{t \in [0, s] : x_{-k}(t) = 0\}$ is empty. $\bar{\tau}$ is the first at which competitors run out of budget. We define the set of controls $\mathbf{B}_1 \subseteq \mathbf{B}$, such that for all $\beta \in \mathbf{B}_1$, we have $\bar{\tau} \geq s$. We will show that any control in this set is dominated in terms of payoff by the FMFE strategy. Let $\mathbf{B}_2 \subseteq \mathbf{B}$, such that for all $\beta \in \mathbf{B}_2$, we have $\bar{\tau} \leq s$. We will show that any control within this set is dominated by the strategy characterized by $\bar{w}(v)$. Hence, the result follows by comparing the value of the FMFE strategy with that of the strategy characterized by $\bar{w}(v)$.

Now, we consider the two cases mentioned above.

Case 1 ($\bar{\tau} \geq s$). In this case we consider controls in the set $\mathbf{B}_1$ for which competitors do not deplete their budgets before the end of the horizon. We shall show that any control $\beta \in \mathbf{B}_1$ is weakly dominated in terms of payoff by the FMFE strategy, given by $w^{\mathrm{F}}_f(v) = v/(1+\mu_k)$.

Any such control is a feasible solution of a variation of the FMFE problem (1) in which all competitors bid according to the FMFE and never run out of budget (and we allow for potentially non-stationary strategies):

$$\begin{aligned} \max_{\beta_k(t,\boldsymbol{x}),\tau} \ & \int_0^{\tau} u_k\big(\beta_k(t,\boldsymbol{x}(t)), \boldsymbol{w}^{\mathrm{F}}_{-k}\big)\,\mathrm{d}t \\ \text{s.t. } & \frac{\mathrm{d}\boldsymbol{x}(t)}{\mathrm{d}t} = -\boldsymbol{g}\big(\beta_k(t,\boldsymbol{x}(t)), \boldsymbol{w}^{\mathrm{F}}_{-k}\big), \\ & \boldsymbol{x}(0) = \boldsymbol{b}, \quad x_k(\tau) \geq 0, \quad \tau \leq s. \end{aligned} \tag{15}$$

In this problem the terminal time (at which bidder k stops participating in the auctions) $\tau \leq s$ is free, and it satisfies that $x_k(\tau) \geq 0$. Note that in problem (15), the budget constraints of the competitors are not taken into account; hence a feasible solution to (15) may potentially violate these constraints in the original problem.

After integrating the advertiser's budget constraint, we introduce $p_k \geq 0$ as the Lagrange multiplier

[17] The proof of the first case of exponential distributions requires a simple additional step relative to the proof of the theorem to deal with the unbounded support of the value distribution.

for the budget constraint to obtain that (15) is upper bounded by

$$
\begin{aligned}
&\sup_{\beta_k(t),\tau\le s} \int_0^\tau \big(u_k\big(\beta_k(t),\boldsymbol{w}^{\mathrm{F}}_{-k}\big) - p_k g_k\big(\beta_k(t),\boldsymbol{w}^{\mathrm{F}}_{-k}\big)\big)\,\mathrm{d}t + p_k b_k \\
&\overset{(a)}{=} \sup_{\tau\le s} \int_0^\tau \left(\sup_{w\in\mathcal{B}} u_k\big(w,\boldsymbol{w}^{\mathrm{F}}_{-k}\big) - p_k g_k\big(w,\boldsymbol{w}^{\mathrm{F}}_{-k}\big)\right)\mathrm{d}t + p_k b_k \\
&\overset{(b)}{=} \sup_{\tau\le s}\ \tau \sup_{w\in\mathcal{B}} \alpha\mathbb{E}\Big[\mathbf{1}\{D_{-k}\le w(V)\}(V-(1+p_k)D_{-k})\Big] + p_k b_k \\
&\overset{(c)}{=} \sup_{\tau\le s}\ \tau\alpha\mathbb{E}\Big[V-(1+p_k)D_{-k}\Big]^+ + p_k b_k \\
&\overset{(d)}{=} \alpha s\mathbb{E}\Big[V-(1+p_k)D_{-k}\Big]^+ + p_k b_k \overset{(e)}{=} J^{\mathrm{F}}_k
\end{aligned}
$$

where (a) follows from optimizing point-wise for each time; (b) follows since the integrand is constant in time; (c) is a consequence of the optimal control being equal to $w^{\mathrm{F}}_k(v) = v/(1+p_k)$; for (d) we use that the term in the supremum is non-negative, and for (e), we use the fact that there is no duality gap in the FMFE problem according to Proposition 3.1 and by choosing p_k to be the FMFE multiplier μ_k.

Hence, we have established that any policy in $\mathbf{B}_1$ may not achieve higher rewards than J^{F}_k, and the latter is achieved by the FMFE strategy. Moreover, by its construction, the FMFE strategy is feasible for problem (15) and competitors' budget constraints are not violated when firm k implements it.

Case 2 ($\bar\tau \le s$). In this case we consider controls $\beta \in \mathbf{B}_2$ such that the competitors deplete before the end of the horizon. Any such control is a solution of a control problem in which we impose that at a free time $\tau \le s$ the competitors have no budget, and after that time the advertiser wins all the auctions for free obtaining a total profit of $(s-\tau)\alpha\mathbb{E}V$. We do allow, however, that the advertiser runs out of budget before her competitors and still wins the auctions when the competitors run out of budget (as we will see, this relaxation will still be dominated by a solution in which the advertiser never runs out of budget before the competitors). The optimization problem is:

$$
\begin{aligned}
\max_{\beta_k(t,\boldsymbol{x}),\tau} \quad & \int_0^\tau u_k\big(\beta_k(t,\boldsymbol{x}(t)),\boldsymbol{w}^{\mathrm{F}}_{-k}\big)\,\mathrm{d}t + (s-\tau)\alpha\mathbb{E}V \qquad (16)\\
\text{s.t.} \quad & \frac{\mathrm{d}\boldsymbol{x}(t)}{\mathrm{d}t} = -\boldsymbol{g}\big(\beta_k(t,\boldsymbol{x}(t)),\boldsymbol{w}^{\mathrm{F}}_{-k}\big),\\
& \boldsymbol{x}(0)=\boldsymbol{b}, \quad x_{-k}(\tau)=0, \quad x_k(\tau)\ge 0, \quad \tau\le s.
\end{aligned}
$$

Note that one may remove the constraint that $\tau \le s$ because for any solution with $\tau^* > s$, one can construct a solution with $\tau \le s$ while weakly increasing rewards. To see this let $\beta \in \mathbf{B}_2$ be any such solution and consider a deviation $\hat\beta$ in which the advertiser bids according to β during time $[0,s]$ and zero afterwards. This deviation is clearly feasible for the original control problem (6) and attains a profit

$$
\int_0^s u_k\big(\beta_k(t,\boldsymbol{x}(t)),\boldsymbol{w}^{\mathrm{F}}_{-k}\big)\,\mathrm{d}t \ge \int_0^{\tau^*} u_k\big(\beta_k(t,\boldsymbol{x}(t)),\boldsymbol{w}^{\mathrm{F}}_{-k}\big)\,\mathrm{d}t + (s-\tau^*)\alpha\mathbb{E}V\,,
$$

because $u_k(\cdot, \boldsymbol{w}^{\mathrm{F}}_{-k}) \leq \alpha \mathbb{E} V$.

In the remainder of the proof we use a duality argument to construct an optimal solution for problem (16).

After integrating both budget constraints, we introduce Lagrange multipliers $p_k \geq 0$ for the decision-maker's budget constraint and p_{-k} for the competitors' budget (recall, we only keep track of one of the competitors' budget since all of them follow the same evolution) to obtain that (16) is upper bounded by

$$\begin{aligned}
&\sup_{\beta_k(t), \tau \geq 0} \int_0^\tau H\Big(\beta_k(t), \boldsymbol{p}\Big)\, \mathrm{d}t + (s-\tau)\alpha \mathbb{E} V + p_k b_k + p_{-k} b_{-k} \\
&= \sup_{\tau \geq 0} \int_0^\tau \sup_{w \in \mathcal{B}} H(w, \boldsymbol{p})\, \mathrm{d}t + (s-\tau)\alpha \mathbb{E} V + p_k b_k + p_{-k} b_{-k} \\
&= \sup_{\tau \geq 0} \tau \left(\sup_{w \in \mathcal{B}} H(w, \boldsymbol{p}) - \alpha \mathbb{E} V \right) + s \alpha \mathbb{E} V + p_k b_k + p_{-k} b_{-k}
\end{aligned}$$

where the second equation follows from optimizing point-wise for each time, and the third because the integrand is constant in time. The dual problem is obtained by minimizing over the dual variables, and using that the dual function is unbounded if the Hamiltonian H is greater than $\alpha \mathbb{E} V$. We thus obtain that the dual problem is exactly given by Problem (7).

In the remainder of the proof we assume that the dual problem (7) is bounded from below. If this is not the case, then (16) is infeasible and the set $\mathbf{B}_2$ is empty.

The rest of the argument is divided in the following five steps.

(i) We show that at an optimal solution, constraint (7b) is binding, and $p_{-k} \leq 0$.

(ii) Let $H(y, \boldsymbol{p}; v)$ be the value of the Hamiltonian when the decision-maker's value is v, her bid is y, and the dual variables are $\boldsymbol{p}$. We establish that for each value v, the problem $\sup_y H(y, \boldsymbol{p}; v)$ admits a unique solution. Furthermore, we deduce that the solution bid function $\bar{w}(\boldsymbol{p}) \in \mathcal{B}$ amounts to bidding the *unique* root $y \in \mathbb{R}$ of

$$v = \left(1 + p_k - \frac{p_{-k}}{K-1}\right) y + \frac{p_{-k}}{h_v\big((1+\mu_{-k})y\big)(K-1)(1+\mu_{-k})} \tag{17}$$

when $\underline{v} \leq v \leq \bar{v}$, bidding $\bar{V}/(1+\mu_{-k})$ when $v > \bar{v}$, and bidding $\underline{V}/(1+\mu_{-k})$ when $v \leq \underline{v}$; where $\bar{v} = \bar{V}(1 + p_k - \frac{p_{-k}}{K-1})/(1+\mu_{-k})$ and $\underline{v} = \underline{V}(1 + p_k - \frac{p_{-k}}{K-1})/(1+\mu_{-k}) + p_{-k}/((K-1)(1+\mu_{-k}) f_v(\underline{V}))$.

(iii) We show that the function $H(\boldsymbol{p}) \triangleq \sup_{w \in \mathcal{B}} H(w, \boldsymbol{p})$ is differentiable in $\boldsymbol{p}$, and its gradient is

$$\nabla_{\boldsymbol{p}} H(p) = \left(-g_k(\bar{w}(\boldsymbol{p}), \boldsymbol{w}^{\mathrm{F}}_{-k}), -g_{-k}(\bar{w}(\boldsymbol{p}), \boldsymbol{w}^{\mathrm{F}}_{-k})\right). \tag{18}$$

(iv) We construct a candidate primal feasible solution for (16) using the first-order conditions of (7).

(v) We show that the candidate solution is primal optimal by showing that it achieves the value of the dual objective (7).

Step (i). Suppose that the constraint (7b) is not binding. Recall that $\sup_{w\in\mathcal{B}} H(w,\boldsymbol{p})$ is a convex function, and therefore, continuous in p_{-k}. Because the constraints are linear and the coefficients are positive, one could improve on the solution by decreasing the value of p_{-k} while keeping feasibility. A contradiction because the dual problem is bounded. Hence it must be that (7b) is binding at an optimal solution.

Next, we show that $p_{-k} \le 0$ by contradiction. Suppose otherwise. Discarding the third term of the hamiltonian, which is strictly positive (since $\mathbf{B}_2$ is assumed non-empty, competitors eventually run out of budget), one obtains that

$$H(w,\boldsymbol{p}) < u_k\big(w,\boldsymbol{w}^{\mathrm{F}}_{-k}\big) - p_k g_k\big(w,\boldsymbol{w}^{\mathrm{F}}_{-k}\big) = \alpha\mathbb{E}\Big[\mathbf{1}\{D_{-k}\le w(V)\}(V-(1+p_k)D_{-k})\Big] \le \alpha\mathbb{E}V$$

where the second inequality follows from $D_{-k}\ge 0$. This contradicts the fact that the hamiltonian is equal to $\alpha\mathbb{E}V$. Hence it must be the case that $p_{-k}\le 0$.

Step (ii). Let $H(y,\boldsymbol{p};v)$ be the value of the Hamiltonian when the decision-maker's value is v, her bid is y, and the dual variables are $\boldsymbol{p}$. Note that here we denote y as the bid submitted for a given value v, that is, in this context y is a number and not a function. For the remainder of this step we remove the dependence on $\boldsymbol{p}$ to simplify the notation. Point-wise for each value v the bid y should maximize the expression

$$H(y;v) = \alpha\mathbb{E}\Big[\mathbf{1}\{D_{-k}\le y\}(v-(1+p_k)D_{-k})\Big] - p_{-k}\alpha\mathbb{E}\Big[\mathbf{1}\{D_k \le V_{-k}/(1+\mu_{-k})\}D_k\Big]\,,$$

where $D_{-k} = \tilde{D}^K$ is the maximum competing bid observed by advertiser k, $D_k = \tilde{D}^{K-1}\vee M_k y$ is the maximum competing bid observed by the competing advertisers, and $\tilde{D}^K = \max\Big\{M_j V_j/(1+\mu_{-k})\Big\}_{j=1}^{K-1}$ is the maximum of $K-1$ matching bids from the competitors. Setting $q(x)=\bar{F}_v((1+\mu_{-k})x)x$ the previous expression can be written as

$$H(y;v) = \alpha F_d^K(y)v - (1+p_k)\alpha\int_{\underline{D}}^{y} x\,\mathrm{d}F_d^K(x) - p_{-k}\alpha\int_{\underline{D}}^{\bar{D}}\left(\alpha q(y\vee x)+(1-\alpha)q(x)\right)\,\mathrm{d}F_d^{K-1}(x)\,.$$

Taking derivatives with respect to the bid y we obtain that

$$\begin{aligned}\frac{\partial H(y;v)}{\partial y} &= \alpha f_d^K(y)v - (1+p_k)\alpha y f_d^K(y) - p_{-k}\alpha^2 q'(y)F_d^{K-1}(y)\\ &= \alpha^2(K-1)F_d^{K-1}(y)\Big((v-(1+p_k-\frac{p_{-k}}{K-1})y)f_v\big((1+\mu_{-k})y\big)(1+\mu_{-k}) - \frac{p_{-k}}{K-1}\bar{F}_v\big((1+\mu_{-k})y\big)\Big)\,,\end{aligned}$$

where the latter follows from $f_d^K(x)=\alpha(K-1)f_v((1+\mu_{-k})x)(1+\mu_{-k})F_d^{K-1}(x)$, and $q'(x)=\bar{F}_v((1+\mu_{-k})x) - f_v((1+\mu_{-k})x)(1+\mu_{-k})x$. Equating the last expression to zero we get that the control policy should satisfy

$$(v-(1+p_k-\frac{p_{-k}}{K-1})y)h_v\big((1+\mu_{-k})y\big)(1+\mu_{-k})-\frac{p_{-k}}{K-1}=0\,.$$

By the IFR assumption and the fact that $p_{-k}\le 0$, we have that the left-hand side is decreasing at every root, and thus Poincare-Hopf Theorem implies that a root is unique whenever it exists. As a result, the function $H(y;v)$ is strictly quasi-concave in y. By evaluating the derivative at the extreme points $\underline{V}/(1+\mu_{-k})$ and $\bar{V}/(1+\mu_{-k})$, we obtain that $\underline{v}\le v\le\bar{v}$ implies that the derivative crosses the zero axis, the root exists, and equation (17) follows. When $v\ge\bar{v}$ the optimal bid is $\bar{V}/(1+\mu_{-k})$ since the latter implies that $H(y;v)$ is non-decreasing in y. A similar argument follows when $v\le\underline{v}$. Hence, $\bar{w}(v)$ is well-defined and characterized as just described.

Step (iii). Note that $H(\boldsymbol{p})=\sup_{w\in\mathcal{B}}H(w,\boldsymbol{p})=\mathbb{E}_V[\sup_y H(y,\boldsymbol{p};V)]$. So we first show that the envelope theorem holds for $\sup_y H(y,\boldsymbol{p};v)$ point-wise in v, and then invoke Leibniz's rule to show that the expectation is differentiable.

Fix the value v. Because the distribution of values has continuous densities, we have that $H(y,\boldsymbol{p};v)$ is continuous in y, and the gradient $\nabla_{\boldsymbol{p}}H(y,\boldsymbol{p};v)$ is jointly continuous in y and $\boldsymbol{p}$. Because the solution $\bar{w}(\boldsymbol{p})(v)$ is unique, by Corollary 2 of Milgrom (1999) we get that the envelope theorem holds and

$$\nabla_{\boldsymbol{p}}\sup_y H(y,\boldsymbol{p};v)=\big(-g_k(\bar{w}(\boldsymbol{p})(v),\boldsymbol{w}^{\mathrm{F}}_{-k};v),-g_{-k}(\bar{w}(\boldsymbol{p})(v),\boldsymbol{w}^{\mathrm{F}}_{-k};v)\big)\ ,$$

where $g_k(y,\boldsymbol{w}^{\mathrm{F}}_{-k};v),g_{-k}(y,\boldsymbol{w}^{\mathrm{F}}_{-k};v)$ are the expenditures when the value is v and the advertiser bids y. Note that the expenditure for the advertiser is bounded by $g_k(y,\boldsymbol{w}^{\mathrm{F}}_{-k};v)\le\alpha\mathbb{E}\tilde{D}^K\le\alpha K\mathbb{E}V/(1+\mu_{-k})\le K\mathbb{E}V<\infty$ because for $\boldsymbol{x}\in\mathbb{R}^K$ we have that $\max\{x_k\}_{k=1}^K\le\sum_k^K|x_i|$. Similarly, we have that $g_{-k}(y,\boldsymbol{w}^{\mathrm{F}}_{-k};v)\le\alpha(\mathbb{E}\tilde{D}^{K-1}+y)\le[(K-2)\mathbb{E}V+\bar{V}]<\infty$. Thus $\nabla_{\boldsymbol{p}}\sup_y H(y,\boldsymbol{p};v)$ is bounded by an integrable function, and thus we can invoke Leibniz's rule to show that the expectation is differentiable, and the gradient expression (18) follows.

Step (iv). Because the dual problem (7) is convex and its primitives differentiable, one may employ the Karush-Kuhn-Tucker (KKT) conditions to characterize its optimal solution. Additionally, from the previous steps we know that $\bar{w}(\boldsymbol{p})$ achieves the supremum of $H(w,\boldsymbol{p})$ for given $\boldsymbol{p}$, and that we can use the envelope theorem to differentiate $H(\bar{w}(\boldsymbol{p}),\boldsymbol{p})$ with respect to $\boldsymbol{p}$. The Karush-Kuhn-Tucker (KKT) conditions for the dual problem are

$$\begin{aligned}
b_k-\gamma_1 g_k\big(\bar{w}(\boldsymbol{p}),\boldsymbol{w}^{\mathrm{F}}_{-k}\big)-\gamma_2 &= 0\,, &(19)\\
b_{-k}-\gamma_1 g_{-k}\big(\bar{w}(\boldsymbol{p}),\boldsymbol{w}^{\mathrm{F}}_{-k}\big) &= 0\,, &(20)\\
H(\bar{w}(\boldsymbol{p}),\boldsymbol{p}) &= \alpha EV\,, &\\
p_k\ge 0 &\perp \gamma_2\ge 0\,, &(21)\\
\gamma_1 &\ge 0\,, &(22)
\end{aligned}$$

where γ_1 is the Lagrange multiplier corresponding to the constraint (7b) and γ_2 is the one corresponding to $p_k\ge 0$, and we used the fact that the constraint on the hamiltonian binds by step (i).

We next construct a primal feasible solution. Set $\bar{\tau} \triangleq \gamma_1$, and set the bidding strategy to be $\beta_k(t, \boldsymbol{x})(\cdot) = \bar{w}(\cdot)$ from 0 until $\bar{\tau}$ and 0 afterwards. Then it follows from the KKT conditions that this solution is feasible for the relaxed primal problem (16) without the constraint $\bar{\tau} \leq s$.

Step (v). We conclude by showing that the candidate primal feasible solution achieves the dual optimal objective. Indeed, the primal objective value is

$$\begin{aligned}
&\bar{\tau} u_k(\bar{w}(\boldsymbol{p}), \boldsymbol{w}^{\mathrm{F}}_{-k}) + (s - \bar{\tau})\alpha \mathbb{E}V \\
&= \alpha s \mathbb{E}V + \bar{\tau}\left(H(\bar{w}(\boldsymbol{p}), \boldsymbol{p}) - \alpha \mathbb{E}V + p_k g_k(\bar{w}(\boldsymbol{p}), \boldsymbol{w}^{\mathrm{F}}_{-k}) + p_{-k} g_{-k}(\bar{w}(\boldsymbol{p}), \boldsymbol{w}^{\mathrm{F}}_{-k})\right) \\
&= \alpha s \mathbb{E}V + p_k \bar{\tau} g_k(\bar{w}(\boldsymbol{p}), \boldsymbol{w}^{\mathrm{F}}_{-k}) + p_{-k} \bar{\tau} g_{-k}(\bar{w}(\boldsymbol{p}), \boldsymbol{w}^{\mathrm{F}}_{-k}) \\
&= \alpha s \mathbb{E}V + p_k b_k + p_{-k} b_{-k} \,,
\end{aligned}$$

where the second equation follows from rearranging terms, the third and last follow from the KKT conditions (19), (20), and (21) above. □

Proof of Corollary D.1. In the case when the valuations are exponential distributed one can use the fact $h_v(x) = 1/\mathbb{E}V$ and exploit (17) to obtain that the optimal strategy is

$$\bar{w}(v) = \frac{v - \mathbb{E}V(K-1)^{-1}(1+\mu_{-k})^{-1} p_{-k}}{1 + p_k - (K-1)^{-1} p_{-k}} \,.$$

Note that in the proof we only used the fact that valuations were bounded to show that $g_{-k}(y, \boldsymbol{w}^{\mathrm{F}}_{-k}; v)$ is bounded by an integrable function. This claim holds in this setting too by exploiting the linearity of the bidding strategy.

Similarly when valuations are uniform $[0, \bar{V}]$, the failure rate is $h_v(x) = \left(\bar{V} - x\right)^{-1}$ and one obtains that the optimal strategy is

$$\bar{w}(v) = \frac{v - \bar{V}(K-1)^{-1}(1+\mu_{-k})^{-1} p_{-k}}{1 + p_k - 2(K-1)^{-1} p_{-k}} \,. \qquad \square$$

E Other Auction Design Decisions

The previous analysis focused on optimally setting the reserve price. Similarly, one can use our framework to study other important auction design decisions for the publisher. In this section we briefly discuss results for two such decisions.

These results complement the ones in the previous sections and show that proper adjustment of the reserve price is key in (1) making profitable for the publisher to try selling all impressions in the exchange before utilizing the alternative channel; and (2) compensating for the thinner markets created by greater disclosure of viewers' information.

E.1 Optimal Allocation of Impressions

Up to this point we assumed that all users visiting the web-site were shared with the exchange. However, the publisher may have an incentive to allocate only a fraction of the web-site's traffic and sell the rest through an alternative channel. In the following we study, for a fixed reserve price r and information disclosure ι, the publisher's selection of an optimal allocation of impressions to the exchange $\eta \in [0, \bar{\eta}]$, where $\bar{\eta} > 0$ denotes the total number of users per unit of time visiting the website. Here, $I_0(r)$ denotes the probability that the impression is won by some advertiser in the exchange when advertisers bid truthfully and the publisher sets a reserve price r. The following result characterizes the optimal rate of the impressions in the presence of an opportunity cost c.

Theorem E.1. (Optimal allocation of impressions). *Suppose that r and ι are fixed. If $cI_0(r) \geq \alpha\lambda s G_0(r)$, then the publisher is better off not participating in the exchange, and the optimal rate of impressions is $\eta^* = 0$. If $cI_0(r) < \alpha\lambda s G_0(r)$, then the publisher stands to gain from participating in the exchange, and the optimal allocation of impressions to the exchange is $\eta^*(r) = \min\{\eta^0(r), \bar{\eta}\}$ where $\eta^0(r) = b/(\alpha s G_0(r))$.*

The first condition corresponds to the expected opportunity cost being greater or equal than the average revenue per impression when bidders are truthful and in such a case, it is natural to expect that the publisher should not allocate any impressions to the exchange. Interestingly, when the publisher benefits from participating in the exchange, he need not always allocate all the impressions to the exchange. While it may seem at first sight that the exchange is a "free option" to test, it is not so due to the presence of budget constraints as we now explain.

Initially, when the supply is sufficiently small, bidders do not deplete their budget and hence are truthful (cf. Proposition 4.1). In such a region, increasing the allocation of impressions yields larger revenues for the publisher, which is in line with intuition. However, when the rate of impressions is high enough, all advertisers deplete their budgets by the end of their campaign and no further revenue may be extracted by allocating more impressions to the exchange, which corresponds to the market being "saturated". The smallest rate at which saturation takes place is exactly given by $\eta^0(r) = b/(\alpha s G_0(r))$. At that rate, advertisers are truthful; beyond that rate, bidders start shading their bids. Allocating further impressions to the exchange does not yield additional revenues since advertisers are already spending all their budget and hence the key resides in understanding the impact of an increase in supply on the opportunity cost. When the publisher increases the impression rate above market saturation, there are again two effects to consider; (1) the direct effect: sending more impressions to the exchange directly increases the publisher's opportunity cost if these impressions are won; and (2) the indirect effect: as more impressions are available, advertisers shade their bids more and in the presence of a reserve price, the probability of a sell and the opportunity cost decrease. In the proof we show that the direct effect dominates and increasing the rate above market saturation is suboptimal. Thus, the optimal rate of impression is the minimum of $\eta^0(r)$ and $\bar{\eta}$.

Next we characterize the optimal decision of the publisher when she optimizes over both the allocation of impressions and the reserve price. In contrast to Theorem E.1, when jointly optimizing over the reserve price and the rate of impressions, the publisher is always better off allocating $\bar{\eta}$ impressions to the exchange. In this case, because the reserve price optimization considers the alternative channel, the exchange becomes a "free option" that is always worth testing.

Corollary E.1. (Joint allocation and reserve price optimization)**.** *Suppose that ι is fixed. The optimal decision for the publisher is to send all impressions to the exchange, and set the reserve price according to Theorem 5.1. That is, the unique optimal rate of impressions is $\eta^* = \bar{\eta}$, and the optimal reserve price is equal to* $\max\{r_c^*, \bar{r}(\bar{\eta})\}$*, where $\bar{r}(\eta) = \sup \mathcal{R}^*(\eta)$ and $\mathcal{R}^*(\eta) = \{r : \alpha\eta s G_0(r) \geq b\}$.*

We study the joint optimization problem by partitioning it in two stages. In the first stage the publisher looks for the allocation of impressions that maximizes the optimal value of the second-stage problem, obtained from optimizing over reserve prices. Exploiting Theorem 5.1 to characterize the maximum profit over reserve prices, we show that the second-stage objective value is increasing with the rate of impressions. Therefore, when jointly optimizing over reserve prices and the rate of impressions, the publisher is better off allocating $\bar{\eta}$ impressions to the exchange.

Indeed, when $\eta < \eta^0(r_c^*)$, advertisers bid truthfully and the auctions decouple. Since the optimal reserve price r_c^* is larger than the opportunity cost, any given impression will potentially raise more revenues in the exchange than in the alternative channel, and the publisher is better off increasing the supply to the exchange. When $\eta \geq \eta^0(r_c^*)$, the publisher sets the reserve price in a way that the advertisers deplete their budgets in expectation while bidding truthfully. In this case the publisher's revenue is constant and does not increase as she increases the supply to the exchange, so we focus on the probability of selling an impression in the exchange. Note that for $\eta \geq \eta^0(r_c^*)$, there is no indirect effect as in the previous cases; for all such values of η, when the reserve price is optimally set, advertisers bid truthfully in equilibrium. In the proof we show, however, that the direct effect decreases the opportunity cost as the allocation of impressions increase. As more impressions are allocated, the reserve price is increased in a way that the advertisers spend the same amount on average, but pay a higher price per impression and receive fewer impressions. As a result, the publisher is better off increasing the allocation to the exchange.

E.2 Optimal Disclosure of Information

When the publisher posts an impression in the exchange she can decide which user information (if any) to disclose to the advertisers. The publisher may decide to disclose, e.g., the content of the web-page, user geographical location, user demographics, or cookie-based behavioral information (which allows bidders to track the user's past activity in the web). On the one hand, each additional level of targeting may reduce the probability that an advertiser matches with a given user, because more criteria need to be satisfied to do so. This may lead to *thinner* markets and could decrease the publisher's revenue. On the other hand, more information provides better targeting that results in higher valuations and higher bids, conditional on a match. Our FMFE framework can be used to analyze (numerically or analytically) different settings regarding the impact of information disclosure on the publisher's profit.

Below, we illustrate this through a particular stylized model for information disclosure.[18]

We assume that information disclosure ι is continuous, and that there is a one-to-one decreasing mapping between information and the matching probability; which, allows one to parameterize information disclosure by α. As a consequence, the publisher can indirectly choose an $\alpha \in [0,1]$. Fix a distribution of values $F_v(\cdot)$. We assume that, conditional on the choice of α, for some $\sigma(\alpha) > 0$, the distribution of values of the advertisers is such that

$$\begin{aligned} F_{v(\alpha)}(x) &= F_v(x/\sigma(\alpha)), \\ \alpha\sigma(\alpha) &= 1, \text{ for all } \alpha \in (0,1]. \end{aligned} \tag{23}$$

The first condition corresponds to the values being scaled by a deterministic factor $\sigma(\alpha)$ (i.e., under this scaling the new values $V(\alpha) \triangleq \sigma(\alpha)V$). In other words, the model is one in which information impacts the scale but not the shape of distribution of valuations. The second assumption governs the relationship between the matching probability and the scaling factor. Ensuring that $\alpha\sigma(\alpha)$ is constant guarantees that the ex-ante mean valuation is independent of the level of information disclosure (the ex-ante mean value is given by $\alpha\mathbb{E}[V(\alpha)] = \alpha\sigma(\alpha)\mathbb{E}V$). That is, under such a model, when the matching probability is halved, advertisers participate in half the number of auctions on average, but in each auction their values are doubled.

With some abuse of notation, let $\Pi(\mu, r, \alpha)$ be the publisher's long-run average profit as a function of the reserve price r and the matching probability α, when advertisers employ a multiplier μ, with values scaled as above. For a fixed matching probability α, the publisher's objective is to maximize profit by choosing a reserve price. The publisher's maximum profit is given by $\Pi(\alpha) = \max_{r\geq 0} \Pi(\mu(r,\alpha), r, \alpha)$, where $\mu(r,\alpha)$ denotes the equilibrium multiplier for the given auction parameters.

Theorem E.2. (Joint information disclosure and reserve price optimization)**.** *Suppose that η is fixed and that advertisers' valuations follow* (23)*. When the publisher reacts to thinner markets by setting an appropriate reserve price, then disclosing more information improves the profit, that is, the publisher's profit* $\Pi(\alpha) = \max_{r\geq 0} \Pi(\mu(r,\alpha), r, \alpha)$ *is non-increasing in α.*

We provide the main ideas behind the argument. First, in view of Theorem 5.1, advertisers bid truthfully at the optimal reserve price, and there is no need to take into account the shading of bids. Hence, when changing α, there is no *indirect effect*, and it is enough to show that the *direct effect* of decreasing α (corresponding to the impact of having thinner markets but larger valuations) increases the publisher's profits. Now, there are two cases to consider. When the expenditure at $r_c^*(\alpha)$ does not exceed the budget, we have that $r_c^*(\alpha)$ is the optimal reserve price. In the proof, we show that in this case profits increase as α decreases. When the expenditure at $r_c^*(\alpha)$ exceeds the budget, then the publisher prices at $\bar{r}(\alpha) = \sup\{r \geq 0 : \alpha\eta s G_0(r,\alpha) \geq b\}$, and advertisers deplete their budgets

[18] Previous papers have analyzed the trade-off introduced by targeting between increasing valuations by improving the match and reducing revenues by creating thinner markets. Bergemann and Bonatti (2011) does so in a market with a continuum of advertisers and a continuum of consumers and Board (2009) in a static auction setting with a fixed reserve price.

in expectation. Here, we show that as the matching probability decreases, the optimal reserve price $\bar{r}(\alpha)$ changes, resulting in more impressions returned to the publisher, and a lower opportunity cost, increasing publishers' profits.

A key piece in the previous result is that the publisher reacts to changes in the distribution of values by adjusting the reserve price. In this case, the publisher can extract advertisers' surplus even if markets are thin. However, failing to properly adjust the reserve price may prevent the publisher from extracting the surplus generated by targeting. In fact, the publisher's revenue may deteriorate when disclosing more information if the reserve price is not properly adjusted as we now explain.

Suppose that the publisher is disclosing an initial level of information that attains a matching probability α_0, she is pricing at the optimal reserve price $r_c^*(\alpha_0)$, and the advertiser's expenditure does not exceed the budget. Consider the publisher's profit as a function of the matching probability when the reserve price is *not* adjusted, which is given by $\Pi(r_c^*(\alpha_0), \alpha)$ (we dropped the dependence on $\mu(r, \alpha)$ to simplify the notation). One can show that $\Pi(r_c^*(\alpha_0), \alpha)$ is locally non-increasing near α_0, that is, a small increment in the disclosure of information actually increases profits.[19] Nonetheless, it is possible to show that disclosing more information and further decreasing α may cause profits to decrease.

E.3 Proof of Results

E.3.1 Proof of Theorem E.1

We use the following lemma to prove the theorem.

Lemma E.1. *Let Y be a non-negative continuous random variable with increasing generalized failure rate. Then for all $y > 0$*

$$\mathbb{P}\{Y \geq y\} \geq \frac{\xi_Y(y) - 1}{\xi_Y(y) y} \mathbb{E}[Y \mathbf{1}\{Y \geq y\}],$$

where $\xi_Y(y)$ is the generalized failure rate of Y.

Proof. Notice that the bound is trivial when $\xi_Y(y) \leq 1$. We prove the equivalent bound $\mathbb{E}[Y | Y \geq y] \leq y \frac{\xi_Y(y)}{\xi_Y(y) - 1}$ when $\xi_Y(y) > 1$. Let $Y_y \triangleq Y | Y \geq y$ be the random variable Y conditional on Y being larger that y. Clearly, the generalized failure rates $\xi_Y(x)$ and $\xi_{Y_y}(x)$ coincide whenever $x \geq y$. By the IGFR assumption we have that the failure rate of the conditional random variable is larger than that of a Pareto random variable with scale y and shape $\xi_Y(y)$, which we denote by P_y. Indeed,

$$h_{Y_y}(x) = \frac{\xi_{Y_y}(x)}{x} = \frac{\xi_Y(x)}{x} \geq \frac{\xi_Y(y)}{x} = h_{P_y}(x).$$

Thus, we have that the random variable P_y dominates Y_y in the failure rate order, which in turns

[19] This follows from the fact that $\Pi(\alpha)$ is the envelope of $\Pi(r, \alpha)$ over reserve prices, $r^*(\alpha_0)$ is optimal at α_0, and that $\Pi(\alpha)$ is non-increasing.

implies that P_y first-order stochastically dominates Y_y (see, e.g., Ross (1996)). Thus,

$$\mathbb{E}[Y|Y \geq y] = \mathbb{E}[Y_y] \leq \mathbb{E}[P_y] = y\frac{\xi_Y(y)}{\xi_Y(y)-1}. \quad \square$$

Proof of Theorem E.1. Fix $r \geq 0$ and let $\Pi(\mu, \eta)$ be the publisher's profit as a function of the rate of impressions, and the equilibrium multiplier, respectively. The publisher's problem amounts to solving $\max_{0\leq\eta\leq\bar{\eta}} \Pi(\mu(\eta), \eta)$. We use Proposition 4.1 to analyze the dependence of the FMFE multiplier on the rate of impressions, $\mu(\eta)$. When $\eta < \eta^0$ advertisers bid truthfully and the equilibrium multiplier is $\mu(\eta) = 0$. When $\eta \geq \eta^0$ advertisers shade their bids so as to deplete their budgets in expectation and the multiplier is the unique solution of the equation $\alpha\eta s G_0\Big((1+\mu)r\Big) = (1+\mu)b$. We deduce that

$$\Pi(\eta) = \begin{cases} \eta\Big(\alpha\lambda s G_0(r) - cI_0(r)\Big), & \text{if } \eta < \eta^0, \\ \lambda b - \eta c I_0\Big((1+\mu(\eta))r\Big), & \text{if } \eta \geq \eta^0. \end{cases}$$

Notice that $\Pi(\eta)$ is continuous in η, and that the first piece is linear in η.

When the opportunity cost is greater or equal to the average revenue per impression (i.e., $cI_0(r) \geq \alpha\lambda s G_0(r)$), the revenue function $\Pi(\eta)$ is decreasing in its domain, and the optimal rate of impressions is $\eta^* = 0$. When the opportunity cost is less than the average revenue per impression (i.e., $cI_0(r) < \alpha\lambda s G_0(r)$), the slope of the first piece is positive and the publisher is better off allocating more impressions.

In the remainder of the proof we prove the claim that $\Pi(\eta)$ is decreasing for $\eta \geq \eta^0$, and thus the optimal rate of impressions is $\min\{\eta^0, \bar{\eta}\}$. Note that in that set, revenues are fixed equal to λb, so it suffices to study the impact of η on the probability of selling an impression in the exchange. Taking derivatives w.r.t. η we obtain that

$$\frac{\mathrm{d}\Pi}{\mathrm{d}\eta} = -cI_0\Big((1+\mu)r\Big) - \eta c I_0'\Big((1+\mu)r\Big) r \frac{\mathrm{d}\mu}{\mathrm{d}\eta},$$

where we dropped the dependence of μ on η. Once again, the impact of increasing the rate of impressions can be separated in a direct and an indirect effect. The first term above corresponds to the direct effect (the impact of increasing the supply, assuming advertisers' strategies are fixed), and the second to the indirect effect (the impact of the change of advertisers' strategies). Invoking the Implicit Function Theorem we may write the derivative of the equilibrium multiplier w.r.t. the rate of impressions as

$$\frac{\mathrm{d}\mu}{\mathrm{d}\eta} = -\frac{G(\mu, r)}{\eta G_\mu'(\mu, r)} = \frac{(1+\mu)b}{\eta(b - \alpha\eta s r G_0'((1+\mu)r))},$$

where the second equation follows from writing $G(\mu, r) = G_(1+\mu)r)/(1+\mu)$, and using the fact that $\alpha\eta s G(\mu, r) = b$. Note that from Lemma B.2 point *iii.*) one gets that $G_\mu'(\mu, r) < 0$, which allows one to conclude that the multiplier is increasing with the rate of impressions. In the remainder of the proof we show that the direct effect dominates the indirect effect.

Combining terms and using the facts that $I_0'(y) = -\alpha\lambda s f_v(y)(1 - I_0(y))$, and $G_0'(y) = (\bar{F}_v(y) - f_v(y)y)(1 - I_0(y))$ one obtains

$$\begin{aligned}\frac{\mathrm{d}\Pi}{\mathrm{d}\eta} &= -cI_0 + c\alpha\lambda s(1+\mu) r f_v (1 - I_0)\frac{b}{b - \alpha\eta s r G_0'} \\ &= \frac{c}{\lambda b - \alpha\lambda\eta s r G_0'}\left(\Big(\underbrace{\lambda b - \eta r I_0}_{(A)}\Big)\alpha\lambda s(1+\mu) r f_v (1 - I_0) - \Big(\underbrace{\lambda b - \alpha\lambda\eta s r \bar{F}_v(1 - I_0)}_{(B)}\Big) I_0\right).\end{aligned}$$

Next, we consider each term in parenthesis at a time.

For the first term in parenthesis, use the fact that the expenditure of the advertisers is equal to the revenue of the publisher and that the probability that the impression is won as $\mathbb{P}\{\hat{W}_{1:\hat{M}} \geq r\} = I_0$ to write

$$\begin{aligned}\lambda b - \eta r I_0 &= \eta\mathbb{E}\left[\mathbf{1}\{\hat{W}_{1:\hat{M}} \geq r\}\left(\max\{\hat{W}_{2:\hat{M}}, r\}\right)\right] - \eta r \mathbb{P}\{\hat{W}_{1:\hat{M}} \geq r\} \\ &= \eta\mathbb{E}\left[\mathbf{1}\{\hat{W}_{1:\hat{M}} \geq r\}\left(\hat{W}_{2:\hat{M}} - r\right)^+\right] = \eta\mathbb{E}\left[\mathbf{1}\{\hat{W}_{2:\hat{M}} \geq r\}\left(\hat{W}_{2:\hat{M}} - r\right)\right] \\ &= \eta\mathbb{E}\left[\mathbf{1}\{\hat{W}_{2:\hat{M}} \geq r\}\hat{W}_{2:\hat{M}}\right] - \eta r\mathbb{P}\{\hat{W}_{2:\hat{M}} \geq r\},\end{aligned} \tag{24}$$

where the second equation follows from writing the maximum as $\max\{x, y\} = x + (y - x)^+$. Notice that this expression is equivalent to the expected publisher's revenue in excess of the reserve price.

We next bound the first term from above. Using an expression for the distribution of the second-highest bid (see, e.g., David and Nagaraja (2003)) for the first equation, and the probability generating function for the Poisson random variable $\hat{M}$ with mean $\alpha\lambda s$ for the second equation, we may write

$$F_{w_{2:M}}(x) = \mathbb{E}\left[F_w(x)^{\hat{M}} + \hat{M}F_w(x)^{\hat{M}-1}\bar{F}_w(x)\right] = (1 + \alpha\lambda s\bar{F}_w(x))e^{-\alpha\lambda s\bar{F}_w(x)},$$

where $F_w(x) = F_v((1+\mu)x)$ is the shaded distribution of values. Similarly, the p.d.f. is given by $f_{w_{2:M}}(x) = (\alpha\lambda s)^2 f_w(x)\bar{F}_w(x)e^{-\alpha\lambda s\bar{F}_w(x)}$. Note that for every multiplier μ, the resulting distribution of the second-highest bid has IGFR whenever the distribution of valuations exhibits IGFR. Indeed, letting $\xi_{w_{2:M}}(x) = x f_{w_{2:M}}(x)/\bar{F}_{w_{2:M}}(x)$ we have that $\xi_{w_{2:M}}(x) = \xi_w(x)\psi\big(\alpha\lambda s\bar{F}_w(x)\big)$, with $\psi(x) = x^2/(e^x - 1 - x)$ positive and decreasing. Since, $\xi_w(x)$ is increasing and $\bar{F}_w(x)$ decreasing, we conclude that $\xi_{w_{2:M}}(x)$ is increasing.

Using Lemma E.1, one may bound from above term (A) above

$$\lambda b - \eta r I_0 \leq \eta\frac{1}{\xi_{w_{2:M}}(r)}\mathbb{E}\left[\mathbf{1}\{\hat{W}_{2:\hat{M}} \geq r\}\hat{W}_{2:\hat{M}}\right].$$

For the second term in parenthesis, we proceed in a similar fashion. Using the joint distribution of the highest and second-highest bid (see, e.g., David and Nagaraja (2003)) we have that the probability that the impression is won and the reserve price is paid is given by $\mathbb{P}\{\hat{W}_{1:\hat{M}} \geq r, W_{2:\hat{M}} < r\} =$

$(\alpha\lambda s)\bar{F}_v(1-I_0)$. Thus, we obtain that

$$\begin{aligned}\lambda b - \eta r(\alpha\lambda s)\bar{F}_v(1-I_0) &= \eta\mathbb{E}\left[\mathbf{1}\{\hat{W}_{1:\hat{M}} \geq r\}\left(\max\{\hat{W}_{2:\hat{M}}, r\}\right)\right] - \eta r\mathbb{P}\{\hat{W}_{1:\hat{M}} \geq r, \hat{W}_{2:\hat{M}} < r\} \\ &= \eta\mathbb{E}\left[\mathbf{1}\{\hat{W}_{2:\hat{M}} \geq r\}\hat{W}_{2:\hat{M}}\right].\end{aligned}$$

Thus, the second term is equal to the expected publisher's revenue when the second-highest bid is above the reserve price.

Putting it all together, one obtains

$$\begin{aligned}\frac{\mathrm{d}\Pi}{\mathrm{d}\eta} &\leq \frac{c\eta\mathbb{E}\left[\mathbf{1}\{\hat{W}_{2:\hat{M}} \geq r\}\hat{W}_{2:\hat{M}}\right]}{\lambda b - \alpha\lambda\eta s r G_0'}\left(\frac{1}{\xi_{w_{2:M}}(r)}\alpha\lambda s(1+\mu) r f_v(1-I_0) - I_0\right) \\ &= \frac{c\eta\mathbb{E}\left[\mathbf{1}\{\hat{W}_{2:\hat{M}} \geq r\}\hat{W}_{2:\hat{M}}\right]}{\lambda b - \alpha\lambda\eta s r G_0'}\left(\frac{\alpha\lambda s\bar{F}_v}{\psi(\alpha\lambda s\bar{F}_v)}e^{-\alpha\lambda s\bar{F}_v} - (1-e^{-\alpha\lambda s\bar{F}_v})\right) \\ &= \frac{c\eta\mathbb{E}\left[\mathbf{1}\{\hat{W}_{2:\hat{M}} \geq r\}\hat{W}_{2:\hat{M}}\right]}{\lambda b - \alpha\lambda\eta s r G_0'}\left(\phi\left(\alpha\lambda s\bar{F}_v\right) - 1\right) \leq 0\end{aligned}$$

with $\phi(x) = (1-e^{-x})/x \leq 1$ for all $x \geq 0$.

E.3.2 Proof of Corollary E.1

Let $\Pi(\mu, r, \eta)$ be the publisher's profit as a function of the equilibrium multiplier, the rate of impressions, and the reserve price, respectively. The publisher's problem amounts to solving $\max_{r\geq 0, 0\leq\eta\leq\bar{\eta}} \Pi(\mu(r,\eta), r, \eta)$, where $\mu(r,\eta)$ is the equilibrium multiplier for the given auction parameters. We prove the result by partitioning the publisher's problem in two stages: in the inner stage, the optimization is conducted over r, while in the outer stage over η.

Let $\Pi(\eta) = \max_{r\geq 0}\Pi(\mu(r,\eta), r, \eta)$ be the objective of the inner optimization. By Theorem 5.1 we have that

$$\Pi(\eta) = \begin{cases}\Pi(0, r_c^*, \eta), & \text{if } \eta \leq \eta^0(r_c^*), \\ \Pi(0, \bar{r}(\eta), \eta), & \text{if } \eta > \eta^0(r_c^*).\end{cases}$$

Notice that $\Pi(\eta)$ is continuous in η since $\bar{r}(\eta_0(r_c^*)) = r_c^*$. Also note that for all values of η, once the reserve price is set optimally, advertisers bid truthfully. In that sense, changing η does not have an *indirect effect* of changing the equilibrium strategies. We next show that $\Pi(\eta)$ in increasing in η.

For the first piece, we have that $\Pi(0, r_c^*, \eta) = \alpha\lambda\eta s G_0(r_c^*) - \eta c I_0(r_c^*)$, which is linear and increasing in η. For the second piece, the objective is $\Pi(0, \bar{r}(\eta), \eta) = \lambda b - \eta c I_0(\bar{r}(\eta))$. Revenues are fixed and equal to λb, and we focus on the opportunity cost. Taking the derivative w.r.t. η, one obtains that

$$\frac{\mathrm{d}\Pi}{\mathrm{d}\eta} = -cI_0(\bar{r}) + c\Big(1 - I_0(\bar{r})\Big)f_v(\bar{r})\alpha\lambda\eta s\frac{\mathrm{d}\bar{r}}{\mathrm{d}\eta},$$

where we dropped the dependence of $\bar{r}$ on η. Since $\alpha\eta s G_0(\bar{r}) = b$, one may invoke the Implicit Function Theorem to write $\mathrm{d}\bar{r}/\mathrm{d}\eta = -b/(\alpha\eta^2 s G_0'(\bar{r}))$. Note that $G_0'(\bar{r}) < 0$ because $\bar{r} > r_0^*$, and thus the optimal reserve price is non-decreasing with the rate of impressions. Combining expressions and using that $G_0'(\bar{r}) = (1 - I_0(\bar{r}))(\bar{F}_v(\bar{r}) - \bar{r} f_v(\bar{r}))$ by Lemma B.1(ii), one obtains

$$\frac{\mathrm{d}\Pi}{\mathrm{d}\eta} = \frac{c\Big(1 - I_0(\bar{r})\Big)}{-\eta G_0'(\bar{r})} \Big(\eta I_0(\bar{r})\bar{F}_v(\bar{r}) + \Big(\lambda b - \eta I_0(\bar{r})\bar{r}\Big) f_v(\bar{r})\Big) .$$

Note that the publisher's revenue (λb) is lower bounded by $\eta \bar{r} I_0(\bar{r})$ since advertisers pay at least the reserve price of the auction. Hence the derivative above is positive and the proof is complete.

E.3.3 Proof of Theorem E.2

Fix α in $(0, 1]$. In view of Theorem 5.1, advertisers bid truthfully at the optimal reserve price. Note that the generalized failure rate of the value distribution (23) is $\xi_{v(\alpha)}(x) = \xi_v(x/\sigma(\alpha))$, and the failure rate is $h_{v(\alpha)}(x) = h_v(x/\sigma(\alpha))/\sigma(\alpha)$. Let $\Pi_0(r, \alpha)$ denote the publisher's profit when advertisers bid truthfully, which after integrating by parts is given by

$$\begin{aligned}
\Pi_0(r, \alpha) &= \alpha\lambda\eta s \int_r^\infty \bar{F}_{v(\alpha)}(x)\Big(\xi_{v(\alpha)}(x) - 1\Big) e^{-\alpha\lambda s \bar{F}_{v(\alpha)}(x)}\,\mathrm{d}x - c\eta\left(1 - e^{-\alpha\lambda s \bar{F}_{v(\alpha)}(r)}\right) \\
&= \alpha\sigma(\alpha)\lambda\eta s \int_{r/\sigma(\alpha)}^\infty \bar{F}_v(x)\Big(\xi_v(x) - 1\Big) e^{-\alpha\lambda s \bar{F}_v(x)}\,\mathrm{d}x - c\eta\left(1 - e^{-\alpha\lambda s \bar{F}_v(r/\sigma(\alpha))}\right) \\
&= \lambda\eta s \int_{\alpha r}^\infty \bar{F}_v(x)\Big(\xi_v(x) - 1\Big) e^{-\alpha\lambda s \bar{F}_v(x)}\,\mathrm{d}x - c\eta\left(1 - e^{-\alpha\lambda s \bar{F}_v(\alpha r)}\right),
\end{aligned}$$

where the second equation follows from our scaling of values and changing the integration variable, and the last from $\alpha\sigma(\alpha) = 1$. Notice that the profit depends on the reserve price exclusively through αr. Hence to simplify the analysis we perform the change of variables $y = \alpha r$, and define the scaled profit as $\Pi_y(y, \alpha) = \Pi_0(y/\alpha, \alpha)$.

For any given α, by Theorem 5.1, the optimal reserve price is unique, bidders bid truthfully at the optimal reserve, and the optimal profit is given by $\Pi_0(\max\{r_c^*(\alpha), \bar{r}(\alpha)\}, \alpha)$ (with some abuse of notation, we make the dependence on α explicit). The result follows by separately analyzing the two possible cases: (1) $r_c^*(\alpha)$ is the optimal reserve price; and (2) $\bar{r}(\alpha)$ is the optimal reserve price. With some abuse of notation, let $G_0(r, \alpha)$ denote the expected expenditure-per-auction in the absence of budget constraints when advertisers bid truthfully.

Case 1. Suppose that $\alpha\eta s G_0(r_c^*(\alpha), \alpha) < b$, i.e., the expenditure at $r_c^*(\alpha)$ does not exceed the budget. Then $r_c^*(\alpha)$ is the optimal reserve price. First, we study the dependence of the optimal reserve value of the one-shot second-price auction on values. Let $r_c^*(\alpha)$ be the optimal reserve price under information α and opportunity cost c. Since, the optimal reserve price solves for $1/h_{v(\alpha)}(x) = x - c$, we get that $r_c^*(\alpha) = \sigma(\alpha) r^*_{c/\sigma(\alpha)}$, where r_c^* is the reserve price at $\alpha = 1$ and $\sigma(1) = 1$.

We need to show that $\Pi_0(r_c^*(\alpha), \alpha) = \max_{r \geq 0} \Pi_0(r, \alpha)$ is non-increasing in α. Or alternatively, by using our scaling $\alpha\sigma(\alpha) = 1$ we need to show that

$$\Pi_0(r_c^*(\alpha), \alpha) = \Pi_y(\alpha\sigma(\alpha) r^*_{c/\sigma(\alpha)}, \alpha) = \Pi_y(r^*_{\alpha c}, \alpha)$$

is non-increasing in α. Since $r^*_{c\alpha}$ is the optimal reserve price for Π_y and the budget constraint is not binding, we may invoke the Envelope Theorem to get that

$$\begin{aligned}\frac{\mathrm{d}\Pi_y(r^*_{\alpha c}, \alpha)}{\mathrm{d}\alpha} &= \frac{\partial \Pi_y}{\partial \alpha}(r^*_{\alpha c}, \alpha) + \frac{\partial \Pi_y}{\partial y}(r^*_{\alpha c}, \alpha)\frac{\mathrm{d}r^*_{\alpha c}}{\mathrm{d}\alpha} = \frac{\partial \Pi_y}{\partial \alpha}(r^*_{\alpha c}, \alpha) \\ &= -(\lambda s)^2 \eta \int_{r^*_{\alpha c}}^{\infty} \bar{F}_v^2(x)\Big(\xi_v(x) - 1\Big) e^{-\alpha\lambda s \bar{F}_v(x)}\,\mathrm{d}x - c\lambda s\eta \bar{F}_v(r^*_{\alpha c}) e^{-\alpha\lambda s \bar{F}_v(r^*_{\alpha c})},\end{aligned}$$

where the third equation follows from differentiating under the integral sign, which is valid because the derivative of the integrand is continuous on its domain. The IGFR assumption and the fact that the optimal reserve price is increasing with the opportunity cost imply that for all $x \geq r^*_{\alpha c}$, $\xi_v(x) \geq \xi_v(r^*_{\alpha c}) \geq \xi_v(r^*_0) = 1$ and hence the integrand above is positive. We conclude that the derivative is negative.

Case 2. Suppose that $\alpha\eta s G_0(r_c^*(\alpha), \alpha) > b$, i.e., the expenditure at $r_c^*(\alpha)$ exceeds the budget. Then $\bar{r}(\alpha) = \sup\{r \geq 0 : \alpha\eta s G_0(r, \alpha) = b\}$ is the optimal reserve price. Using the scaling and integrating by parts, we obtain that the optimal reserve price $\bar{r}(\alpha)$ satisfies the equation

$$b = \alpha\eta s G_0(\bar{r}(\alpha), \alpha) = \eta s \int_{\alpha\bar{r}(\alpha)}^{\infty} \bar{F}_v(x)\Big(\xi_v(x) - 1\Big) e^{-\alpha\lambda s \bar{F}_v(x)}\,\mathrm{d}x. \tag{25}$$

Now advertisers deplete their budgets in expectation and the publisher's profit is given by

$$\Pi_0(r, \alpha) = \lambda b - c\eta\left(1 - e^{-\alpha\lambda s \bar{F}_v(\alpha r)}\right).$$

Applying the change of variables $y = \alpha r$, and defining $\bar{y}(\alpha)$ as the scaled optimal reserve price; we obtain that the optimal profit is given by $\Pi_0(\bar{r}(\alpha), \alpha) = \Pi_y(\bar{y}(\alpha), \alpha)$. Taking derivatives w.r.t. the matching probability we obtain

$$\frac{\mathrm{d}\Pi_y(\bar{y}(\alpha), \alpha)}{\mathrm{d}\alpha} = \frac{\partial \Pi_y}{\partial \alpha}(\bar{y}(\alpha), \alpha) + \frac{\partial \Pi_y}{\partial y}(\bar{y}(\alpha), \alpha)\frac{\mathrm{d}\bar{y}(\alpha)}{\mathrm{d}\alpha}.$$

To conclude that the profit is non-increasing we shall show that both terms are non-positive. Indeed, the partial derivative w.r.t. the matching probability is $\partial\Pi_y/\partial\alpha = -c\lambda s\eta \bar{F}_v(\bar{y}(\alpha)) e^{-\alpha\lambda s \bar{F}_v(\bar{y}(\alpha))} \leq 0$. Similarly, the partial derivative w.r.t. the scaled reserve price is $\partial\Pi_y/\partial y = c\eta\alpha\lambda s f_v(\bar{y}(\alpha)) e^{-\alpha\lambda s \bar{F}_v(\bar{y}(\alpha))} \geq 0$. Finally, invoking the Implicit Function Theorem we get from equation (25) that the total derivative of the scaled optimal reserve price is

$$\frac{\mathrm{d}\bar{y}(\alpha)}{\mathrm{d}\alpha} = -\frac{\lambda s \int_{\bar{y}(\alpha)}^{\infty} \bar{F}_v^2(x)\Big(\xi_v(x) - 1\Big) e^{-\alpha\lambda s \bar{F}_v(x)}\,\mathrm{d}x.}{\bar{F}_v(y(\alpha))\Big(\xi_v(y(\alpha)) - 1\Big) e^{-\alpha\lambda s \bar{F}_v(y(\alpha))}} \leq 0.$$

For the last inequality recall that, by assumption, $\bar{r}(\alpha) > r_c^*(\alpha)$, which implies that $\bar{y}(\alpha) > r_{\alpha c}^* \geq r_0^*$. Using the IGFR assumption we obtain that $\xi_v(y(\alpha)) > \xi_v(r_0^*) \geq 1$, and then both the numerator and the denominator are non-negative. Hence, the optimal reserve price is non-increasing with the matching probability.

Putting it all together. The optimal profit is given by

$$\Pi(\alpha) = \Pi_0(\max\{r_c^*(\alpha), \bar{r}(\alpha)\}, \alpha) = \Pi_y(\max\{r_{\alpha c}^*, \bar{y}(\alpha)\}, \alpha),$$

where $\Pi_y(y, \alpha)$ is jointly continuous in y and α. From case 1 and 2, we know that $r_{\alpha c}^*$ is continuous and increasing in α, while $\bar{y}(\alpha)$ is continuous and non-increasing in α. Thus, $\Pi(\alpha)$ is continuous in α; $r_{\alpha c}^* = \bar{y}(\alpha)$ in at most one point; and the profit is non-decreasing in α. This concludes the proof.