Auctions for Online Display Advertising Exchanges: Approximations and Design

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

Graduate School of Business, Columbia University srb2155@columbia.edu, ob2105@columbia.edu, gyw2105@columbia.edu

September 20, 2012

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. Our first main contribution is to introduce the novel notion of a Fluid Mean Field Equilibrium (FMFE) that is behaviorally appealing, computationally tractable, and in some important cases yields a closed-form characterization. Moreover, we show that a FMFE approximates well the rational behavior of advertisers in large markets. Our second main contribution is to use this framework to provide sharp prescriptions for key auction design decisions that publishers face in these markets, such as the reserve price, the allocation of impressions to the exchange versus an alternative channel, and the disclosure of viewers' information. Notably, 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.

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 (The Economist, 2011). 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 posts the ad slot in the exchange together with potentially some user information known to her; e.g., the user's geographical location and her *cookies*. Advertisers (or bidders) interested in advertising on the site post bids as a function of their targeting criteria and the user information provided by the publisher. Then, an auction is run to determine the winning advertiser 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. The publisher repeatedly offers slots to display advertisements on her web-site as users arrive; typically, a given publisher runs millions of these auctions per day. On its part, advertisers participate in the exchange with the objective of fulfilling marketing campaigns. In practice, such campaigns are commonly based on a given 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 presence of this new channel presents a host of new strategic and tactical questions for publishers. How should the allocation of impressions be divided between bilateral contracts and exchange(s)? How should Ad Exchange auctions be designed to maximize profits given that advertisers are budget constrained? What is the role of user information and to what extent should it be disclosed? In this paper, we propose a framework that captures some key characteristics of an exchange, through which one may start quantifying the central trade-offs in these markets and answering such questions. Our first main contribution is to introduce a novel equilibrium concept for advertisers' competition in an exchange which is behaviorally appealing, computationally tractable, and in some important cases yields closed-form characterizations. Our second main contribution is to use this framework to provide prescriptions and insights for key auction design decisions that publishers face. We explain these next.

1.1 Main Contributions

Our model incorporates advertisers' budget constraints that are prevalent in these markets. These constraints link the different auctions over time, and therefore advertisers 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. Profit is measured as the difference between the sum of the values generated by the impressions won during the length of their campaign minus the expenditures. In many cases, advertisers 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. Then, given the auction format, advertisers compete in a *dynamic* game. In order to quantify the impact of auction design parameters, we focus first on 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, a notion of equilibrium that one may consider is Perfect Bayesian Equilibrium (PBE) in which advertisers maintain priors on the states of all other bidders, and update them accordingly using Bayes' rule. Even if priors could be succinctly updated, bidders are left with the problem of computing a best response, which is a highdimensional dynamic program. 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 first contribution of this paper is a novel notion of equilibrium that is tractable and provides a good approximation to the strategic interactions among budgetconstrained bidders in an Ad Exchange. Our new notion of equilibrium combines 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 mean field approximation is that, when the number of competitors 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. This type of approximations have appeared in other auction and industrial organization applications (see, e.g., Iyer et al. (2011); Weintraub et al. (2008)). Moreover, in Ad Exchange markets the number of participants is typically large. 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, it also fits well in the context of Ad Exchanges.

Using these two approximations, we define the notion of a *Fluid Mean Field Equilibrium* (FMFE). A FMFE is an appealing equilibrium concept for Ad Exchanges from a behavioral perspective. Moreover, we justify the use of the FMFE as an equilibrium concept in this setting by proving that the FMFE provides a good approximation to the rational behavior of agents in large markets.

Notably, when a second-price auction is conducted, the resulting FMFE strategy has a simple, yet

¹In 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.

appealing, form: an advertiser needs to shade her values by a *constant factor*. Intuitively, when budgets are tight, advertisers shade their bids, because there is an option value for future good opportunities. We leverage the latter characterization to analyze properties of FMFE. In particular, we show that an FMFE always exists and provide a broad 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 the strategies under the FMFE, and of the resulting competitive landscape in the case of homogeneous bidders, i.e., when all advertisers have the same budget and campaign length. Such closed forms for equilibria of dynamic games are remarkably rare and one may significantly leverage such a result when studying the publisher's problem.

Auction Design. Given our characterization of the outcome of the interactions among advertisers, we study the auction design problem for a publisher that maximizes expected profits. In particular, we analyze the impact of three different decision variables on the publisher's profits when running second-price auctions: the reserve price, the allocation of impressions, and the disclosure of information. 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.

We start by analyzing the model with homogeneous bidders, for which we can prove analytical results. First, we provide a complete characterization of the optimal reserve price. Second, we derive the optimal rate of impressions to allocate to the exchange (vis-à-vis collecting the opportunity cost upfront). We show that when the reserve price is fixed, profits initially increase with the allocation of impressions, but it is not necessarily optimal to send all impressions to the exchange. This result stems from the fact that beyond a certain level, the publisher may not extract further revenues because of the budget constraints, and allocating more impressions increases the opportunity cost. When jointly optimizing over the rate of impressions and the reserve price, however, we establish that the publisher is always better off increasing the allocation of impressions as much as possible. In this case, because the reserve price optimization considers the alternative channel, the exchange becomes a "free option" that is always worth testing.

When the publisher posts an impression in the exchange she can decide which 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. We apply our framework to a stylized model for information disclosure, and show that if the publisher reacts to thinner markets by setting an appropriate reserve price, then disclosing more information will always increase the publisher's profits; an appropriately set reserve price allows to extract surplus even in thin markets.

The results for homogeneous advertisers are of independent interest. Moreover, they provide insights that may be valid for the more general case with heterogeneous bidders, i.e., when advertisers have different budgets and campaign lengths. More specifically, the structure of the optimal reserve price suggests how the publisher should balance extracting revenues from budget-constrained bidders with minimizing the opportunity cost. In addition, the last two results highlight the importance of performing reserve price optimization when adjusting the other two auction design levers, namely, the allocation of impressions and the level of users' information disclosure. We also performed numerical experiments to check for the robustness of the previous results in the model with heterogeneous bidders.

Overall, our results provide sharp insights on the design of Ad Exchange markets and on the publisher's profit maximization problem. Notably, the structure of the optimal auction design decisions are simple and intuitive, and may have implications on the design of such auctions in practice. At the same time, 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. In fact, we expect that FMFE may have additional applications beyond the one presented in this paper. This work also contributes to the revenue management literature; the publisher's optimization of impression allocation and selling mechanism are core revenue management problems and so is, in some way, the advertisers's scarce resource allocation problem.

1.2 Related Work

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 not considered. Vulcano et al. (2002) considers a related problem in the context of a single-leg revenue management problem. From the advertiser's perspective, Ghosh et al. (2009) study the design of a bidding agent that implements a campaign in the presence of an exogenous market.

Regarding the disclosure of information, Levin and Milgrom (2010) discuss how targeting can increase efficiency by improving the match between users and advertisers, but at the same time reduce publisher's revenues by creating thinner market. With this motivation, Celis et al. (2011) introduce a new randomized auction mechanism that experimentally performs better than an optimized secondprice auction in markets that become thin due to targeting. This work, however, considers a one shot auction and does not take into account the dynamics introduced by budget constraints.

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 contributes to recent work in mean-field approximations to dynamic games. Weintraub et al. (2008) and Adlakha et al. (2011) use the related notions of oblivious equilibrium and mean field equilibrium, respectively, to approximate Markov perfect equilibrium in dynamic oligopoly models that are commonly studied in industrial organization. More related to our work is Iyer et al. (2011) that use a mean-field notion of equilibrium to study dynamic repeated auctions in which bidders learn about their own private values over time. Our mean field approximation build on theirs. However, in our setting dynamics are driven by budget constraints as opposed to learning. Moreover, our fluid approximation to the bidders' control problem enables us to prove sharper results regarding the equilibrium characterization and auction design. Closest to our paper is the very recent study of Gummadi et al. (2012) that, in simultaneous and independent work, also study budget-constrained bidders in repeated auctions and define a related mean field equilibrium concept. However, they do not provide approximation nor auction design results, which are a key part of our contribution.

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 literature. 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.

Our work also relates from both methodological and approach standpoints 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 notion of fluid, or open-loop, equilibrium. Other papers studying on dynamic games in revenue management (all focusing on price competition) include Farias et al. (2011), de Albéniz and Talluri (2011), and Dudey (1992).

2 Model Description

We study a continuous-time infinite horizon setting in which users arrive to an online publisher's webpage 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. 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 value an advertiser perceives for the impression, and the amount she is willing to bid. The last two components of the type, α_k and γ_k , determine the targeting criteria and the valuation distributions 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 $F_v(\cdot; \gamma_k)$, parametrized by $\gamma_k \in \mathbb{R}$. The distributions have compact support $[\underline{V}, \overline{V}] \subset \mathbb{R}_+$ and continuously differentiable density.² Later, we will explain how the publisher can affect the value of α_k and the distribution parameter γ_k by changing her user information disclosure policy.

At the moment of arrival, an advertiser's type is drawn independently from a common knowledge distribution with support Θ , 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

 $^{^{2}}$ 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.

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 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. We analyze three levers that the publisher may use to do so: (i) the reserve price r to set for the auctions, (ii) the volume of ads to send to the exchange given the positive external opportunity cost c, and (iii) the amount of information she discloses to the advertisers that we will represent with a real number ι .³

Notation. Given a random variable X, we denote a realization x with lower case, its sample space 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,

³To simplify notation, we will not make this dependence explicit until needed.

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 be 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 might be 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 bidders 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).

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, and bidding strategies ignore the individual state and are only dependent on the actual realization of the bidder's value. This can be shown to provide advertisers with provably good policies 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 $\beta_{\theta}^{\text{F}}(v; F_d)$ that places a bid depending solely on her value v for the impression. Notice that 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

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

s.t.
$$\alpha \eta s \mathbb{E} \Big[\mathbf{1} \{ D \le w(V) \} D \Big] \le b,$$
 (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 1. Suppose that $\mathbb{E}[D] < \infty$. An optimal bidding strategy that solves (1) for type θ is given by

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

where μ_{θ}^* is an optimal solution of the dual problem $\inf_{\mu \ge 0} \Psi_{\theta}(\mu; F_d)$ with $\Psi_{\theta}(\mu; F_d) = \alpha \eta s \mathbb{E} \Big[V - (1 + \mu) D \Big]^+ + \mu b.$

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 the multiplier μ_{θ}^* is the shadow price 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. The proof of the result relies on an analysis of the dual of problem (1). Note that (1) is not a convex program and hence standard strong duality arguments do not apply. As a result, 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 Q(t) the set of indices of the advertisers in the exchange at time t, and by Q(t) = |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, $\Omega(t) = \{Q(t), \{b_k(t), s_k(t), \theta_k\}_{k \in Q(t)}\}$, 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.

Suppose that all bidders conjecture a common distribution of the maximum bid they compete against, and implement the optimal fluid strategy described in the preceding section. These strategies induce, through the market dynamics, a new distribution of bids. In equilibrium we require a consistency check: the resulting steady-state distribution of the maximum bid coincides with the one that was originally postulated.

A difficulty with the 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). We rigorously justify that this additional layer of approximation is asymptotically correct in large markets (see Appendix B and Balseiro et al. (2012)).⁴

Since arrival and departures of advertisers are governed by an $M/G/\infty$ queue and campaign lengths

 $^{^{4}}$ 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. We show that this layer of approximation is asymptotically correct. Furthermore, it plays a key role to obtain tractability in our analysis.

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 $(\hat{M}, \{\hat{\Theta}_k\}_{k=1}^{\hat{M}})$ be a random vector that describes the number of matching bidders, 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 $\theta \in \Theta$, and independent of \hat{M} .⁵

For a fluid-based strategy profile $\boldsymbol{\beta} = \{\beta_{\theta}(\cdot) : \theta \in \boldsymbol{\Theta}\}$ with $\beta_{\theta} : [\underline{V}, \overline{V}] \to \mathbb{R}$, 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 1 (Fluid Mean Field Equilibrium). A fluid-based strategy profile β constitutes a FMFE if for every advertiser's type $\theta \in \Theta$, the bidding function β_{θ} is optimal for problem (1) given that the distribution of the maximum bid of other advertisers is given by $F_d(\beta)$ (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.⁶ Finally, for a given set of auction design parameters (r, η, ι) , we denote by $C_{\text{FMFE}}(r, \eta, \iota)$ the set of corresponding FMFEs.

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 in Appendix B, we show that when all advertisers implement the FMFE strategy, the relative

⁵For a type $\theta \in \Theta$ 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.

⁶Note that by the PASTA property of a Poisson arrival process this assumption is in fact correct.

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.

3.3 Publisher's Problem

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 the followers. In particular, the publisher first selects the reserve price in the second-price auction r, the rate of impressions η , and the extent of information disclosed ι . Then, after observing these design parameters, 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 optimization problem may be written as

$$\max_{r,\eta,\iota} \quad \eta \mathbb{E} \Big[\mathbf{1} \{ \beta_{\hat{\Theta}}(V_{\hat{\Theta}})_{1:\hat{M}} \ge r \} \Big(\max \{ \beta_{\hat{\Theta}}(V_{\hat{\Theta}})_{2:\hat{M}}, r \} - c \Big) \Big]$$
(3a)

s.t.
$$\boldsymbol{\beta} \in \mathcal{C}_{\text{FMFE}}(r, \eta, \iota)$$
 (3b)

where $\beta_{\hat{\Theta}}(V_{\hat{\Theta}})_{1:\hat{M}}$ is the highest bid among the bids $\left\{\beta_{\hat{\Theta}_k}(V_{\hat{\Theta}_k})\right\}_{k=1}^{\hat{M}}$ submitted in the auction and $\beta_{\hat{\Theta}}(V_{\hat{\Theta}})_{2:\hat{M}}$ the second-highest bid. The publisher's problem amounts to maximizing her long run average profit rate (3a) considering the opportunity cost of the alternative channel, subject to the constraint that β is an FMFE. For given (η, r, ι) , a priori, a FMFE may not exist or even if one exists it may not be unique. Hence, the optimization problem above is not well posed in general. However, we will prove in the coming sections that a FMFE exists and is unique for an important class of models.

Note that both r and η directly affect the publisher's objective. In addition, it is clear from equations (1), that they could affect bidders' FMFE strategies. Furthermore, recall that ι affects the matching probabilities $\alpha_k(\iota)$ and the valuation parameters $\gamma_k(\iota)$. Hence, changing ι also directly affects the publisher's objective and the FMFE being played. In the next sections, we make these dependencies more explicit and analyze the optimal selection of each of those design decisions.

4 Fluid Mean Field Equilibrium Characterization

In this section we prove the existence, provide conditions for uniqueness, and characterize the FMFE. Proposition 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

Consider fixed values of the auction parameters (r, η, ι) . We first prove the existence of a FMFE. Recall from Proposition 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 $\boldsymbol{\mu} = {\{\mu_{\theta}\}}_{\theta \in \Theta}$ a vector of multipliers in $\mathbb{R}_{+}^{|\Theta|}$ for the different advertiser's 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 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

$$\mu_{\theta}^* \in \arg\min_{\mu \ge 0} \Psi_{\theta}(\mu; \mu^*), \quad \text{for all types } \theta \in \Theta.$$
(4)

One may establish that the system of equations (4) always admits a solution to obtain the following.

Theorem 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 $\mathbf{G} : \mathbb{R}_{+}^{|\Theta|} \times \mathbb{R}_{+} \to \mathbb{R}_{+}^{|\Theta|}$ 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 bidder 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 $D = \max\left(\left(V_{\hat{\Theta}}/\left(1+\mu_{\hat{\Theta}}\right)\right)_{1:\hat{M}}, r\right)^{-7}$.

Assumption 1 (P-matrix). The Jacobian of $-\mathbf{G}$ with respect to $\boldsymbol{\mu}$ is a P-matrix for all $\boldsymbol{\mu}$ in $\mathbb{R}_{+}^{|\Theta|}$

A matrix is P-matrix if all its principal minors are positive (Horn and Johnson, 1991, p.120). Assumption 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 Θ is a singleton. In Appendix C, we provide examples of settings with heterogeneous advertisers in which it also holds. The next theorem shows that the equilibrium is unique under the P-matrix assumption.

 $^{^{7}}$ Note that consistent with the FMFE assumption and the PASTA property, the bidder competes against the *market* steady-state maximum bid.

Theorem 2. Suppose Assumption 1 holds. Then, there is a unique FMFE of the form $\beta_{\theta}(v) = v/(1 + \mu_{\theta}), \theta$ in Θ .

We prove the result by formulating the FMFE conditions as a Non-linear Complementarity Problem (NCP) as presented in Corollary 1 below, and employing a Univalence Theorem to show that the expected expenditure mapping is injective (Facchinei and Pang, 2003a). Moreover, it is possible to establish under further mild regularity conditions 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 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.

4.2 Equilibrium Characterization

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

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

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

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

Note that the expected expenditure for a bidder of type θ over its campaign when bidders use a vector of multipliers μ is given by $\alpha_{\theta}\eta_{s_{\theta}}G_{\theta}(\mu, r)$, because on average she faces $\eta_{s_{\theta}}$ 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 (2003b) 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. In this case, we shall see that Assumption 1 always holds and there exists a unique FMFE. Let

$$G_0(r) = G_\theta(0, r) \tag{5}$$

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 $\alpha\eta sG_0(r)$. This quantity plays a key role in the FMFE characterization.

Proposition 2. Suppose Θ 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 $v \in [\underline{V}, \overline{V}]$, 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., $\alpha\eta sG_0(r) < b$), then in equilibrium advertisers will bid truthfully. If however, budgets are tight (i.e., $\alpha\eta sG_0(r) \ge b$), 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 $\mu^* = (\alpha\eta sG_0(0)/b - 1)^+$.

4.3 Reformulation of the Publisher's Problem

When Assumption 1 holds, given the existence and uniqueness of a FMFE, one may now properly and explicitly formulate the publisher's problem. Let $I(\mu, r) = 1 - F_d(r; \mu)$ denote the probability that the impression is won by some advertiser in the exchange when advertisers shade according to the profile μ and the publisher sets a reserve price r. Using the characterization of an FMFE in Corollary 1, one may alternatively write the publisher's problem in terms of multipliers rather than bidding functions, and obtain the following Mathematical Program with Equilibrium Constraints (MPEC)

$$\max_{r,\eta,\iota} \quad \lambda \sum_{\theta \in \Theta} p_{\theta} \alpha_{\theta} \eta s_{\theta} G_{\theta}(\boldsymbol{\mu}, r) - \eta c I(\boldsymbol{\mu}, r)$$
s.t.
$$\mu_{\theta} \ge 0 \quad \perp \quad \alpha_{\theta} \eta s_{\theta} G_{\theta}(\boldsymbol{\mu}, r) \le b_{\theta}, \quad \forall \theta \in \Theta,$$
(6)

where $p_{\theta} \triangleq \mathbb{P}_{\Theta} \{\theta\}$ is the probability that an arriving advertiser is of type θ . We denote by $\Pi(\boldsymbol{\mu}, r, \eta, \iota)$ the objective function of the MPEC. The first term in the objective is the publisher's revenue rate obtained from all bidders' types, 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.

When studying the publisher's maximization problem, it will be useful to separate the impact of changing each of the design decisions on the objective. First, there is a *direct effect*: assuming advertiser's strategies remain unchanged, modifying a decision directly impacts the objective. Second, there is an *indirect effect*; modifying a decision may change the induced FMFE strategy and this, in turn, impacts the objective.

5 Auction Design and Allocation Decisions

In this section we focus on the publisher's profit maximization problem. The framework developed in the previous sections can be applied to general inputs with regard to advertiser heterogeneity and market parameters as one could solve the resulting MPEC profit maximization problem. However, to gain further insights on the publisher's problem, and the trade-offs at hand, we first focus on the case of homogeneous bidders. In this case, quasi-closed form solutions may be obtained for the publisher's optimal decisions. In §5.2, we illustrate numerically how some of these insights may generalize to the case of heterogeneous bidders, and help better understand the latter.

5.1 The Case of Homogeneous Advertisers

We first consider the case in which 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 . That is, bidders are homogeneous and the space of types Θ is a singleton. By Proposition (2) 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)).⁸</sup>

5.1.1 Optimal Reserve Price

In the absence of budget constraints, 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.

⁸For instance the uniform, exponential, triangular, truncated normal, gamma, Weibull, and log-normal distribution have IGFR.

Theorem 3. (Optimal reserve price). Suppose that η and ι are fixed. 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^* , $\alpha\eta sG_0(r_c^*)$, 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^* ($\alpha\eta sG_0(r_c^*) \leq b$), 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 (6) 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) the *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 (2) the *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 $\bar{r} = \sup \mathcal{R}^*$, that minimizes the opportunity cost in \mathcal{R}^* , is optimal.

It is worthwhile to compare the result above with the work in 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^* . 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 $\alpha\eta s G_0(r_c^*) \geq b$. 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.

5.1.2 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 4. (Optimal allocation of impressions). Suppose that r and ι are fixed. If $cI_0(r) \ge \alpha \lambda sG_0(r)$, then the publisher is better off not participating in the exchange, and the optimal rate of impressions is $\eta^* = 0$. If $cI_0(r) < \alpha \lambda sG_0(r)$, 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 $\eta^0(r) = b/(\alpha sG_0(r))$.

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 2). 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 4, 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 2. (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 3. 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 $\mathcal{R}^*(\eta) = \{r : \alpha\eta s G_0(r) \ge b\}$.

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 3 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 $\eta \ge \eta^0(r_c^*)$, 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 $\eta \ge \eta^0(r_c^*)$, 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.

5.1.3 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.⁹

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

$$F_{v(\alpha)}(x) = F_v(x/\sigma(\alpha)),$$

$$\alpha \sigma(\alpha) = 1, \text{ for all } \alpha \in (0, 1].$$

$$(7)$$

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 $\Pi(\alpha) = \max_{r\geq 0} \Pi(\mu(r, \alpha), r, \alpha)$,

⁹Previous 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.

where $\mu(r, \alpha)$ denotes the equilibrium multiplier for the given auction parameters.

Theorem 5. (Joint information disclosure and reserve price optimization). Suppose that η is fixed and that advertisers' valuations follow (7). 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 3, 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 $\bar{r}(\alpha) = \sup\{r \ge 0 : \alpha\eta s G_0(r, \alpha) \ge b\}$, and advertisers deplete their budgets 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.¹⁰ Nonetheless, it is possible to show that disclosing more information and further decreasing α may cause profits to decrease.

5.2 Numerical Results for the Case of Heterogeneous Advertisers

While it was possible to obtain essentially closed-form solutions for the publisher's decisions in the case of homogeneous advertisers, it is not, in general, possible to derive such results for the case of

¹⁰This 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.

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 1 for different auction parameters. In this section, we conduct a series of numerical experiments to explore the impact of the allocation of impressions and the reserve price on the publisher's profit to conduct some robustness check on some of the conclusions of §5.1.

The setup 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. Finally, we assume that values are exponentially distributed with each type's mean independently and 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]. In total, we consider 150 different scenarios. Additionally, for each scenario, we experiment with increasing levels of impressions allocated to the exchange, as given by η .¹¹

For each combination of parameters and value of η , we optimize problem (6) over the reserve price; and compute the optimal reserve price r^* and the optimal profit for the publisher II^{*}. Then, we compute the latter quantities for a grid of values of the rate of impressions η . An important conclusion in §5.1 was that it is optimal for the publisher to send all the impressions to the exchange as long as the reserve price was properly adjusted. For all sampled parameters, we indeed find that, when the publisher accounts for the optimal reserve price, her profit is increasing with the allocation of impressions. As before, the publisher is better off sending all impressions to the exchange, even in the presence of an alternative channel.

To explore an illustrative example in further detail, Figure 1 depicts, for a given set of parameters with two types, the optimal publisher's profit (a), the equilibrium advertisers' expenditures (b) and multipliers (c) at the optimal solution, as well as the optimal reserve price (d) (all as a function of the allocation of impressions to the exchange η). Notice that when the publisher prices optimally, the high-budget type always bids truthfully. However, this is not necessarily true for the low-budget type. As opposed to the homogeneous case, now the publisher cannot perfectly discriminate between types and for some levels of supply, low-type advertisers will shade their bids under the optimal reserve price.

Focusing on the optimal reserve price, as expected, 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

¹¹For fixed auction parameters, solving for the FMFE takes a few seconds on a laptop computer.



Figure 1: Optimal profit, expenditures, multipliers and reserve price as a function of the rate of impressions for an instance with $\alpha = 0.1$, $\lambda = 1$, s = 40, Unif[0,2] 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) the solid line denotes the optimal profit. (b) the solid lines correspond to the campaign expenditures for each type, while the dashed lines denote the budgets. (c) equilibrium multipliers as a function of the allocation of impressions. (d) the solid line corresponds to optimal reserve price, while the dashed lines denote the optimal advertisers share the same budget (either b_1 or b_2).

optimal reserve price. As the publisher shares more impressions with the exchange, 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 surplue 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 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 this point the optimal reserve price starts increasing.

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. As previously mentioned, in these experiments, we also observed that the publisher's profit rate is increasing in the rate of impressions. As an illustration, in Figure 2 we provide similar plots than the ones above for a market with five advertiser types.

The numerical experiments above illustrate the behavior and tradeoff one observes with heterogeneous advertisers, focusing on the allocation of impressions. In the same manner, given an arbitrary mapping from information to distribution of values, it would be possible to explore the optimal amount of information disclosure. The latter is an important question in these markets and our framework could allow to study the impact of different information disclosure policies on publisher's profits when advertisers are heterogeneous.

6 Conclusions

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 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



Figure 2: Optimal profit, expenditures, multipliers and reserve price as a function of the rate of impressions for an instance with $\alpha = 0.1$, $\lambda = 4$, s = 10, Exponential(1) distribution of values, c = 1, $b_{\theta} = 1, 2, 3, 4, 5$, and types probabilities are equal to 1/5 each.

data in exchanges, a direction we are currently exploring.

Acknowledgments

The authors thank Vahab Mirrokni for helpful insights as well as the participants at the 2012 MSOM Conference and ISMP 2012 for their feedback. They also thank the NET Institute (www.NETinst.org) and the Deming Center (www.gsb.columbia.edu/deming) for financial support.

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.
- Asmussen, S. (2003), *Applied Probability and Queues*, Stochastic Modelling and Applied Probability, Vol. 51, 2 edn, Springer.
- 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.
- 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.
- 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. (2003a), Finite-Dimensional Variational Inequalities and Complementarity Problems, Volume 1, Springer Series in Operations Research and Financial Engineering, Springer.
- Facchinei, F. and Pang, J.-S. (2003b), 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.

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.
- 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.
- 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.
- 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.
- The Economist (2011), 'Innovation in online advertising: Mad men are watching you'.

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.
- 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

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

Lemma 1. *i.*) The distribution of the maximum competing bid when $x \ge r$ can be characterized 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

$$G_{\theta}(\boldsymbol{\mu}, r) = r\bar{F}_{v_{\theta}}((1+\mu_{\theta})r)F_{d}(r; \boldsymbol{\mu}) + \int_{r}^{V} x\bar{F}_{v_{\theta}}((1+\mu_{\theta})x) dF_{d}(x; \boldsymbol{\mu}).$$

Proof of (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.

Proof of (ii). The expenditure function can be written as

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

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 2. *i.*) For any μ , the maximum bid $D \sim F_d(\mu)$ 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 $r \in [\underline{V}, \overline{V}]$, $\partial G_{\theta}(\boldsymbol{\mu}, r) / \partial \mu_{\theta} < 0$.

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

A.1 Proof of Proposition 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

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

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

$$\inf_{\mu \ge 0} \sup_{w(\cdot)} \mathcal{L}_{\theta}(w, \mu) = \inf_{\mu \ge 0} \left\{ \alpha \eta s \sup_{w(\cdot)} \left\{ \mathbb{E} \left[\mathbf{1} \{ D \le w(V) \} \left(V - (1+\mu)D \right) \right] \right\} + \mu b \right\}$$
$$= \inf_{\mu \ge 0} \left\{ \alpha \eta s \mathbb{E} \left[\mathbf{1} \{ (1+\mu)D \le V \} \left(V - (1+\mu)D \right) \right] + \mu b \right\}$$
$$= \inf_{\mu \ge 0} \left\{ \alpha \eta s \mathbb{E} \left[V - (1+\mu)D \right]^{+} + \mu b \right\},$$

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 $\Psi_{\theta}(\mu) = \alpha \eta s \mathbb{E} \left[V - (1+\mu)D \right]^+ + \mu b$. 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 $\alpha\eta s\mathbb{E}\left[\mathbf{1}\left\{D \leq V\right\}D\right] \leq b$, 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 $\alpha\eta s\mathbb{E}\left[\mathbf{1}\left\{D \leq V\right\}D\right] > b$. 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 $\alpha\eta s\mathbb{E}\left[\mathbf{1}\left\{D \leq \frac{V}{1+\mu}\right\}D\right] = b$, 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 $\beta_{\theta}^{\mathrm{F}}(v) = v/(1 + \mu^*)$ satisfy $\mu^* \left(b - \alpha \eta s \mathbb{E} \left[\mathbf{1} \left\{ D \leq \beta_{\theta}^{\mathrm{F}}(V) \right\} D \right] \right) = 0$, and thus the complementary slackness conditions hold. Additionally from the first-order conditions of the dual, we get that the bid function $\beta_{\theta}^{\mathrm{F}}(\cdot)$ is primal feasible. We conclude by showing that the primal objective of the proposed bid function attains the dual objective. That is,

$$\alpha\eta s\mathbb{E}\left[\mathbf{1}\{D \leq \beta_{\theta}^{\mathrm{F}}(V)\}\left(V - D\right)\right] = \mathcal{L}_{\theta}(\beta_{\theta}^{\mathrm{F}}, \mu^{*}) + \mu^{*}\left(b - \alpha\eta s\mathbb{E}\left[\mathbf{1}\{D \leq \beta_{\theta}^{\mathrm{F}}(V)\}D\right]\right)$$
$$= \mathcal{L}_{\theta}(\beta_{\theta}^{\mathrm{F}}, \mu^{*}) = \Psi_{\theta}(\mu^{*}),$$

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 $\beta_{\theta}^{\mathrm{F}}$ is the optimal bid function.

A.2 Proof of Theorem 3

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 (6), 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 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 $r_c^* \ge c$, 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 $[\underline{V}, \overline{V}]$, 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

$$\frac{\partial \Pi}{\partial r}(0,r) = \alpha \lambda \eta s \left(G_0'(r) + c e^{-\alpha \lambda s \bar{F}_v(r)} f_v(r) \right) = \alpha \lambda \eta s \bar{F}_v(r) e^{-\alpha \lambda s \bar{F}_v(r)} \left(1 - \frac{r-c}{r} \xi(r) \right), \quad (8)$$

where the second equation follows by Lemma 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 $[\underline{V}, \overline{V}]$.

- (ii) Then the set \mathcal{R}^* is a closed bounded interval. The proof follows by noticing that setting c = 0in (8) implies that $G_0(r)$ is strictly quasi-concave in r (in the interval $[\underline{V}, \overline{V}]$). 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 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 2. 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 2, $\mu > 0$ 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 $\mu'r + 1 + \mu \ge 0$ 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,

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

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 $G(\mu, r) = g \ge 0$, and that $G(\mu, r)$ is decreasing in μ for fixed r by Lemma 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 $\Pi(\mu(r), r) \leq \Pi(0, \bar{r}_c)$. 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., $G_0(r_c^*) \ge g$). 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 $r_c^* \le \bar{r} \le r$, we have that $\Pi(0, \bar{r}) \ge \Pi(0, r)$. 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.

A.3 Proof of Corollary 2

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 3 we have that

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

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 $\alpha\eta sG_0(\bar{r}) = b$, one may invoke the Implicit Function Theorem to write $d\bar{r}/d\eta = -b/(\alpha\eta^2 sG'_0(\bar{r}))$. 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 1(ii), one obtains

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

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.

B Approximation

In this section, we present a result that establishes that the FMFE provides a good approximation to the rational behavior of agents when the markets are large and the number of bidding opportunities per advertiser are also large. More specifically, we show that when all advertisers implement the FMFE strategy, the relative increase in payoff of any unilateral deviation to a strategy that keeps track of all information available to the advertiser in the market becomes negligible as the market scale increases. Hence, FMFE strategies become asymptotically optimal.

We consider a sequence of markets indexed by a positive parameter κ , referred to as the scaling; such that the higher the scaling, the larger the market "size". On the demand side, a θ -type advertiser matching probability decreases as $\alpha_{\theta}^{\kappa} \propto \kappa^{-1}$, while the budget increases as $b_{\theta}^{\kappa} \propto \log \kappa$. Additionally, the arrival rate of advertisers increases as $\lambda_{\theta}^{\kappa} \propto \kappa$; and both the distribution of values and the length of the campaign are invariant to the scaling. On the supply side, the arrival rate of impressions increases as $\eta^{\kappa} \propto \kappa \log \kappa$. Hence, the mean number of auctions an advertiser participates in, $\alpha_{\theta}^{\kappa} \eta^{\kappa} s_{\theta} \propto \log \kappa$, grows at the same rate that the budget. The scaling is such that auctions occur more frequently, but the expected number of matching bidders in each auction, $\alpha_{\theta}^{\kappa} \lambda_{\theta}^{\kappa} s_{\theta}$, remains constant. Additionally, the FMFE is *invariant* to the scaling, because advertisers aim to satisfy the budget constraints in expectation and strategies are state-independent (see Eq. (1) and (2)). Thus, irrespectively of the scaling, the FMFE strategy is given by $\beta^{\rm F} = \{\beta_{\theta}^{\rm F}\}_{\theta \in \Theta}$ and is described by a vector of multipliers.

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 advertiser's arrival time to the system; her initial budget; length of stay in the exchange; the realizations of her values up to that time; her bids; and whether she won or not the auctions, and in the cases she did win, the payments made to the publisher. We define a pure strategy β for advertiser k as a mapping from histories to bids. A strategy specifies, given an history $h_k(t)$ and assuming the advertiser participates in an auction at time t, an amount to bid $\beta(h_k(t))$. We denote by \mathbb{B} be the space of strategies that are non-anticipating and adaptive to the history.

For a fixed scaling κ , we study the expected payoff of a fixed advertiser, referred to as the zeroth advertiser, from the moment she arrives to the exchange until her departure, when she implements a strategy $\beta \in \mathbb{B}$ and all other advertisers follow the FMFE strategies β^{F} . This expected payoff is denoted by $J^{\kappa}_{\theta}(\beta, \beta^{\rm F})$, where the expectation is taken over the actual market process, with the initial market state drawn at random from an appropriate distribution. In this notation, $J^{\kappa}_{\theta}(\beta^{\rm F}_{\theta}, \beta^{\rm F})$ measures the *actual* expected payoff of the FMFE strategy for the advertisers in the exchange.¹² We have the following result.

Theorem 6. Suppose that r > 0 and that there are at most two bidders' types. Consider a market with scaling κ in which all bidders, except the zeroth bidder, follow the FMFE strategy β^{F} . Suppose that a θ -type advertiser (the zeroth bidder), upon arrival to the market, deviates and implements a nonanticipating and adaptive strategy $\beta^{\kappa} \in \mathbb{B}$. The relative expected payoff of this deviation with respect to the FMFE strategy $\beta^{\text{F}}_{\theta}$ satisfies

$$\limsup_{\kappa \to \infty} \frac{J_{\theta}^{\kappa}(\beta^{\kappa}, \boldsymbol{\beta}^{\mathrm{F}})}{J_{\theta}^{\kappa}(\beta^{\mathrm{F}}_{\theta}, \boldsymbol{\beta}^{\mathrm{F}})} \leq 1,$$

when the initial states of the advertisers in the market are drawn from an appropriately pre-specified distribution.¹³

The result establishes that the payoff increase of a deviation to a strategy that keeps track of all available information, relative to the payoff of the FMFE strategy, becomes negligible as the scale of the system increases. Therefore, FMFE approximates well the rational behavior of advertisers, in the sense that unilateral deviations to more complex strategies do not yield significant benefits. The proof of this result is provided in Balseiro et al. (2012) and is beyond the scope of the present paper but we discuss its main ideas below.

The key simplifications in the FMFE were that: i.) All advertisers present in the market were allowed to bid and the possibility of them running out of budget was only taken into account to compute an appropriate shading parameter in the fluid optimization problem, but not when sampling competitors' bids; and ii.) The mean field model assumes that the actions of an advertiser do not affect the competitors in the market, and that competitors' states and the number of matching bidders in successive auctions are independent.

The first step of the proof consists of addressing i.). To that end, we introduce a new mean field model, referred to as budget-constrained mean-field model (BMFM), that is similar to the original fluid model, but that accounts explicitly for the fact that advertisers may run out of budget, and not participate in some auctions. We establish that in the BMFM, when the scale increases, the expected fraction of time that any bidder has positive budget during her campaign converges to one. Using this result and techniques borrowed from revenue management (see, e.g., Talluri and van Ryzin (1998)), we show that the FMFE strategy is near-optimal when an advertiser faces the competition induced by the BMFM. This result justifies our initial assumption in the FMFE that advertisers present in the market do not run out of budgets.

¹²Note that this performance metric may differ from the FMFE value function, given by the objective value of the approximation problem $J_{\theta}^{\rm F}(F_d)$ given in (1).

 $^{^{13}}$ We discuss the nature of this distribution in Balseiro et al. (2012) and we show that this distribution gets close to the FMFE steady-state distribution as the market scale increases. In addition, we show that the assumption on the maximum number of types can be relaxed under further technical conditions.

The second step of the proof consists of addressing ii.) above. Given our scaling, we show that with high probability an advertiser interacts throughout her campaign with distinct advertisers who do not share any past common influence, and that the same applies recursively to those advertisers she competes with. This implies that, in this regime, the states of the competitors faced by the zeroth advertiser are essentially independent, and that her actions have negligible impact on future competitors. Additionally, we show that the impact of the queueing dynamics on the number of matching bidders may be appropriately bounded, and that the number of matching bidders in successive auctions are asymptotically uncorrelated. These steps combine a propagation of chaos argument for the interactions (similar to that used in Graham and Méléard (1994) and Iyer et al. (2011)) and a fluid limit for the advertisers' queue. Thus, as the scaling increases the real market behaves like the BMFM.

We note that Theorem 6 is proved for a given family of scalings. We conjecture, however, that the family of scalings under which our approximation result is valid is broader. In fact, the first step above generalizes to other scalings. On the other hand, the second step relies quite heavily on the nature of the scaling. For this step, our scaling and techniques are similar to those present in the papers using a propagation of chaos argument mentioned in the previous paragraph. An interesting technical avenue for future research is the generalization of these techniques and the family of scalings under which the second step above (and ultimately Theorem 6) holds. This generalization is likely to have other applications in mean-field models beyond the one presented in this paper.

Supplementary Appendix

C Sufficient Conditions for P-matrix Assumption to Hold

We establish here sufficient conditions for Assumption 1, that was required for uniqueness of a FMFE, to hold.

Proposition 3. The P-matrix condition (Assumption 1) holds in either of the following cases.

- i.) Θ is a singleton.
- ii.) Θ 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 $n \ge 0$ 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 1 holds when values are uniformly distributed with support $[0, \overline{V}]$.

D Additional Proofs

D.1 Proof of Theorem 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 $\bar{\mu} \triangleq \bar{\alpha}\eta\bar{s}V/\underline{b}$ is a dominated strategy. To see this notice that for every $\mu > \bar{\mu}$ we have that

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

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

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

Step 2. Next, we show that for each type $\theta \in \Theta$ the objective function $\Psi_{\theta}(\mu; \mu)$ is jointly continuous in μ and μ . Consider a sequence $(\mu^n, \mu^n) \in U \times U^{|\Theta|}$ converging as $n \to \infty$ to some (μ, μ) 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 1(i)), we get that the same holds for the maximum bid. Denoting by D^n the maximum bid random variable associated to $\boldsymbol{\mu}$; the previous argument implies that D^n converges in distribution to D. Additionally, by Slutsky's Theorem we have get that $(1 + \mu^n)D^n \Rightarrow (1 + \mu)D$.

Consider the function $\ell(x) = \mathbb{E}[V - x]^+ = \int_x^\infty \bar{F}_v(y) \, dy$. 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 $\Psi_{\theta}(\mu; \mu) = \alpha \eta s \mathbb{E} \left[\ell \left((1 + \mu)D \right) \right] + \mu b$. By portmanteau theorem we have that $\mathbb{E} \left[\ell \left((1 + \mu^n)D^n \right) \right] \to \mathbb{E} \left[\ell \left((1 + \mu)D \right) \right]$, and thus Ψ is jointly continuous in (μ, μ) .

Step 3. Because the domain is compact, Ψ is jointly continuous in (μ, μ) , and convex in μ for fixed μ (cf. proof of Proposition 1), an FMFE is guaranteed to exist by Proposition 8.D.3 in Mas-Colell et al. (1995).

D.2 Proof of Theorem 2

Exploiting the fact that the dual objective is convex and differentiable, one may write the equilibrium condition (4) 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

$$\mu_{\theta}^{*} = 0, \frac{\partial \Psi_{\theta}}{\partial \mu}(\mu_{\theta}^{*}, \boldsymbol{\mu}^{*}) \ge 0,$$

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

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

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

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 2 we have that the mapping **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.

D.3 Proof of Proposition 2

Fix $r \ge 0$. The existence of the equilibrium follows from Theorem 1. The uniqueness follows from the fact that Assumption 1 is automatically satisfied in the present case from item *iii*.) of Lemma 2. We next derive the characterization of the FMFE.

Suppose first that $\alpha \eta s G_0(r) < b$. By Lemma 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 $\alpha \eta s G_0(r) \ge b$, 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 1). Noting that $(1 + \mu^*)G(\mu^*, r) = G_0((1 + \mu^*)r)$ concludes the proof.

D.4 Proof of Theorem 4

We use the following lemma to prove the theorem.

Lemma 3. Let Y be a non-negative continuous random variable with increasing generalized failure rate. Then for all y > 0

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

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,

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

Thus, we have that the random variable P_y dominates Y_y in the failure rate order, which in turns implies that P_y first-order stochastically dominates Y_y (see, e.g., Ross (1996)). Thus,

$$\mathbb{E}[Y|Y \ge y] = \mathbb{E}[Y_y] \le \mathbb{E}[P_y] = y \frac{\xi_Y(y)}{\xi_Y(y) - 1}. \quad \Box$$

Proof of Theorem 4. Fix $r \ge 0$ 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 $\max_{0\le\eta\le\bar{\eta}}\Pi(\mu(\eta),\eta)$. We use Proposition 2 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 $\eta \ge \eta^0$ advertisers shade their bids so as to deplete their budgets in expectation and

the multiplier is the unique solution of the equation $\alpha \eta s G_0((1+\mu)r) = (1+\mu)b$. We deduce that

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

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., $cI_0(r) \ge \alpha\lambda sG_0(r)$), 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., $cI_0(r) < \alpha\lambda sG_0(r)$), 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 $\eta \ge \eta^0$, 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

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

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

$$\frac{\mathrm{d}\mu}{\mathrm{d}\eta} = -\frac{G(\mu, r)}{\eta G'_{\mu}(\mu, r)} = \frac{(1+\mu)b}{\eta (b-\alpha\eta sr G'_0((1+\mu)r))},$$

where the second equation follows from writing $G(\mu, r) = G_{(1 + \mu)r)/(1 + \mu)}$, and using the fact that $\alpha\eta sG(\mu, r) = b$. Note that from Lemma 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

$$\begin{split} \frac{\mathrm{d}\Pi}{\mathrm{d}\eta} &= -cI_0 + c\alpha\lambda s(1+\mu)rf_v(1-I_0)\frac{b}{b-\alpha\eta srG_0'}\\ &= \frac{c}{\lambda b - \alpha\lambda\eta srG_0'} \Bigg(\Big(\underbrace{\lambda b - \eta rI_0}_{(A)}\Big)\alpha\lambda s(1+\mu)rf_v(1-I_0) - \Big(\underbrace{\lambda b - \alpha\lambda\eta sr\bar{F}_v(1-I_0)}_{(B)}\Big)I_0 \Bigg). \end{split}$$

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

$$\begin{aligned} \lambda b - \eta r I_0 &= \eta \mathbb{E} \left[\mathbf{1} \{ \hat{W}_{1:\hat{M}} \ge r \} \left(\max\{ \hat{W}_{2:\hat{M}}, r \} \right) \right] - \eta r \mathbb{P} \{ \hat{W}_{1:\hat{M}} \ge r \} \\ &= \eta \mathbb{E} \left[\mathbf{1} \{ \hat{W}_{1:\hat{M}} \ge r \} \left(\hat{W}_{2:\hat{M}} - r \right)^+ \right] = \eta \mathbb{E} \left[\mathbf{1} \{ \hat{W}_{2:\hat{M}} \ge r \} \left(\hat{W}_{2:\hat{M}} - r \right) \right] \\ &= \eta \mathbb{E} \left[\mathbf{1} \{ \hat{W}_{2:\hat{M}} \ge r \} \hat{W}_{2:\hat{M}} \right] - \eta r \mathbb{P} \{ \hat{W}_{2:\hat{M}} \ge r \}, \end{aligned}$$
(10)

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 secondhighest 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 $\xi_{w_{2:M}}(x) = xf_{w_{2:M}}(x)/\bar{F}_{w_{2:M}}(x)$ we have that $\xi_{w_{2:M}}(x) = \xi_w(x)\psi(\alpha\lambda s\bar{F}_w(x))$, 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 3, one may bound from above term (A) above

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

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

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

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

$$\begin{aligned} \frac{\mathrm{d}\Pi}{\mathrm{d}\eta} &\leq \frac{c\eta \mathbb{E} \left[\mathbf{1} \{ \hat{W}_{2:\hat{M}} \geq r \} \hat{W}_{2:\hat{M}} \right]}{\lambda b - \alpha \lambda \eta sr G'_0} \left(\frac{1}{\xi_{w_{2:M}}(r)} \alpha \lambda s (1+\mu) r f_v (1-I_0) - I_0 \right) \\ &= \frac{c\eta \mathbb{E} \left[\mathbf{1} \{ \hat{W}_{2:\hat{M}} \geq r \} \hat{W}_{2:\hat{M}} \right]}{\lambda b - \alpha \lambda \eta sr G'_0} \left(\frac{\alpha \lambda s \bar{F}_v}{\psi (\alpha \lambda s \bar{F}_v)} e^{-\alpha \lambda s \bar{F}_v} - (1-e^{-\alpha \lambda s \bar{F}_v}) \right) \\ &= \frac{c\eta \mathbb{E} \left[\mathbf{1} \{ \hat{W}_{2:\hat{M}} \geq r \} \hat{W}_{2:\hat{M}} \right]}{\lambda b - \alpha \lambda \eta sr G'_0} \left(\phi \left(\alpha \lambda s \bar{F}_v \right) - 1 \right) \leq 0 \end{aligned}$$

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

D.5 Proof of Theorem 5

Fix α in (0, 1]. In view of Theorem 3, advertisers bid truthfully at the optimal reserve price. Note that the generalized failure rate of the value distribution (7) 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 3, 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 $\alpha\eta sG_0(r_c^*(\alpha), \alpha) < b$, 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 $r_{\alpha\alpha}^*$ is the optimal reserve price for Π_y and the budget constraint is not binding, we may invoke the Envelope Theorem to get that

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

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 $x \ge r_{\alpha c}^*$, $\xi_v(x) \ge \xi_v(r_{\alpha c}^*) \ge \xi_v(r_0^*) = 1$ 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 $\bar{r}(\alpha) = \sup\{r \ge 0 : \alpha \eta s G_0(r, \alpha) = b\}$ 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

$$b = \alpha \eta s G_0(\bar{r}(\alpha), \alpha) = \eta s \int_{\alpha \bar{r}(\alpha)}^{\infty} \bar{F}_v(x) \Big(\xi_v(x) - 1\Big) e^{-\alpha \lambda s \bar{F}_v(x)} \,\mathrm{d}x.$$
(11)

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 (11) that the total derivative of the scaled optimal reserve price is

$$\frac{\mathrm{d}\bar{y}(\alpha)}{\mathrm{d}\alpha} = -\frac{\lambda s \int_{\bar{y}(\alpha)}^{\infty} \bar{F}_{v}^{2}(x) \left(\xi_{v}(x) - 1\right) e^{-\alpha\lambda s \bar{F}_{v}(x)} \mathrm{d}x.}{\bar{F}_{v}(y(\alpha)) \left(\xi_{v}(y(\alpha)) - 1\right) e^{-\alpha\lambda s \bar{F}_{v}(y(\alpha))}} \le 0.$$

For the last inequality recall that, by assumption, $\bar{r}(\alpha) > r_c^*(\alpha)$, which implies that $\bar{y}(\alpha) > r_{\alpha c}^* \ge r_0^*$. Using the IGFR assumption we obtain that $\xi_v(y(\alpha)) > \xi_v(r_0^*) \ge 1$, 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 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.

D.6 Proof of Lemma 2

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,

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

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 1(i), the distribution of the maximum competing bid when $x \ge r$ 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 $\partial F_d/\partial\mu_{\theta} = F_d(x; \boldsymbol{\mu})\mathbb{E}[\alpha_{\Theta}\lambda s_{\Theta}]\hat{p}_{\theta}xf_{v_{\theta}}((1+\mu_{\theta})x)$, 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 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 $[\underline{V}, \overline{V}] \times U$ 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

$$(I) = \frac{\partial}{\partial \mu_{\theta}} \left(r \bar{F}_{v_{\theta}} ((1+\mu_{\theta})r) F_d(r; \boldsymbol{\mu}) \right)$$

= $-r^2 f_{v_{\theta}} ((1+\mu_{\theta})r) F_d(r; \boldsymbol{\mu}) + r \bar{F}_{v_{\theta}} ((1+\mu_{\theta})r) \frac{\partial F_d}{\partial \mu_{\theta}} (r; \boldsymbol{\mu})$

and

$$(II) = \frac{\partial}{\partial \mu_{\theta}} \int_{r}^{\bar{V}} x \bar{F}_{v_{\theta}} ((1+\mu_{\theta})x) \frac{\partial F_{d}}{\partial x} dx$$

$$= -\int_{r}^{\bar{V}} x^{2} f_{v_{\theta}} ((1+\mu_{\theta})x) \frac{\partial F_{d}}{\partial x} dx + \int_{r}^{\bar{V}} x \bar{F}_{v_{\theta}} ((1+\mu_{\theta})x) \frac{\partial^{2} F_{d}}{\partial \mu_{\theta} \partial x} dx$$

$$= -\int_{r}^{\bar{V}} x^{2} f_{v_{\theta}} ((1+\mu_{\theta})x) \frac{\partial F_{d}}{\partial x} dx - r \bar{F}_{v_{\theta}} ((1+\mu_{\theta})r) \frac{\partial F_{d}}{\partial \mu_{\theta}} (r; \boldsymbol{\mu})$$

$$-\int_{r}^{\bar{V}} \frac{\partial}{\partial x} \left(x \bar{F}_{v_{\theta}} ((1+\mu_{\theta})x) \right) \frac{\partial F_{d}}{\partial \mu_{\theta}} dx,$$

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 $w(x) \triangleq \mathbb{E}[\alpha_{\Theta}\lambda s_{\Theta}]F_{d}(x; \mu)$. Using this new notation and canceling terms we can write the partial derivative as

$$\frac{\partial G_{\theta}}{\partial \mu_{\theta}} = -\sum_{\theta' \neq \theta} (1 + \mu_{\theta'}) \hat{p}_{\theta'} \langle f_{\theta}, f_{\theta'} \rangle - \hat{p}_{\theta} \langle f_{\theta}, \bar{F}_{\theta} \rangle - r f_{\theta}(r) F_d(r; \boldsymbol{\mu}), \tag{12}$$

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

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

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 $D \geq V_{\Theta'}/(1+\mu_{\Theta'})$ 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 $\bar{F}_{v_{\theta}}(\cdot) \leq 1$, and $\lim_{x\to\infty} \bar{F}_{v_{\theta}}(x) = 0$.

D.7 Proof of Proposition 3

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 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 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

$$\frac{\partial G_{\theta}}{\partial \mu_{\theta'}} = r \bar{F}_{v} ((1+\mu_{\theta})r) \frac{\partial F_{d}}{\partial \mu_{\theta'}} + \frac{\partial}{\partial \mu_{\theta'}} \int_{r}^{V} x \bar{F}_{v} ((1+\mu_{\theta})x) \frac{\partial F_{d}}{\partial x} dx$$

$$= r \bar{F}_{v} ((1+\mu_{\theta})r) \frac{\partial F_{d}}{\partial \mu_{\theta'}} + \int_{r}^{\bar{V}} x \bar{F}_{v} ((1+\mu_{\theta})x) \frac{\partial^{2} F_{d}}{\partial \mu_{\theta'} \partial x} dx$$

$$= -\int_{r}^{\bar{V}} \frac{\partial}{\partial x} \left(x \bar{F}_{v} ((1+\mu_{\theta})x) \right) \frac{\partial F_{d}}{\partial \mu_{\theta'}} dx$$

$$= (1+\mu_{\theta}) \hat{p}_{\theta'} \langle f_{\theta} f_{\theta'} \rangle - \hat{p}_{\theta'} \langle f_{\theta'} \bar{F}_{\theta} \rangle, \qquad (13)$$

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 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

$$det(J) = \frac{\partial G_1}{\partial \mu_1} \frac{\partial G_2}{\partial \mu_2} - \frac{\partial G_1}{\partial \mu_2} \frac{\partial G_2}{\partial \mu_1}$$

= $(1 + \mu_1) \hat{p}_1^2 \langle f_1 f_2 \rangle \langle f_1 \bar{F}_1 \rangle + (1 + \mu_1) \hat{p}_1 \hat{p}_2 \langle f_1 f_2 \rangle \langle f_1 \bar{F}_2 \rangle$
+ $(1 + \mu_2) \hat{p}_1 \hat{p}_2 \langle f_1 f_2 \rangle \langle f_2 \bar{F}_1 \rangle + (1 + \mu_2) \hat{p}_2^2 \langle f_1 f_2 \rangle \langle f_2 \bar{F}_2 \rangle$
+ $\hat{p}_1 \hat{p}_2 \langle f_1 \bar{F}_1 \rangle \langle f_2 \bar{F}_2 \rangle - \hat{p}_1 \hat{p}_2 \langle f_1 \bar{F}_2 \rangle \langle f_2 \bar{F}_1 \rangle,$

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 $f_i(x) = x f_v((1 + \mu_i)x) = x h_v((1 + \mu_i)x) \overline{F_v}((1 + \mu_i)x) = (1 + \mu_i)^n x h_v(x) \overline{F_i}(x)$. Defining a new weight function $\tilde{w}(x) = x h_v(x) w(x)$ and using Cauchy-Schwartz inequality one gets that

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

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

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