

Auctions for Online Display Advertising Exchanges: Approximation Result Technical Report

Santiago Balseiro, Omar Besbes, Gabriel Weintraub

Graduate School of Business, Columbia University

srb2155@columbia.edu, ob2105@columbia.edu, gyw2105@columbia.edu

First version: September, 2012

This version: November 9, 2012

1 Approximation

In this technical report we show 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. Before stating the result, we proceed by formalizing the scaling under consideration.

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) in the main paper). Thus, irrespectively of the scaling, the FMFE strategy is given by $\beta^F = \{\beta_\theta^F\}_{\theta \in \Theta}$ 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 $J_{\theta}^{\kappa}(\beta, \beta^{\text{F}})$, 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, $J_{\theta}^{\kappa}(\beta_{\theta}^{\text{F}}, \beta^{\text{F}})$ measures the *actual* expected payoff of the FMFE strategy for the advertisers in the exchange.¹ We have the following result.

Theorem 1. *Suppose that $r \in (0, \bar{V})$ 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 non-anticipating and adaptive strategy $\beta^{\kappa} \in \mathbb{B}$. The relative expected payoff of this deviation with respect to the FMFE strategy $\beta_{\theta}^{\text{F}}$ satisfies*

$$\limsup_{\kappa \rightarrow \infty} \frac{J_{\theta}^{\kappa}(\beta^{\kappa}, \beta^{\text{F}})}{J_{\theta}^{\kappa}(\beta_{\theta}^{\text{F}}, \beta^{\text{F}})} \leq 1,$$

*when the initial states of the advertisers in the market are drawn from an appropriately pre-specified distribution.*²

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 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

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

²We discuss the nature of this distribution in Section 3 of this report 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 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.

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 1 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 1) holds. This generalization is likely to have other applications in mean-field models beyond the one presented in this paper.

Preliminaries. In the rest of this report, we drop the dependence on the scaling κ when clear from the context. Throughout this report we assume that the reserve price is positive, that is, $r > 0$. The latter excludes the possibility of an advertiser winning an impression for free.

As a preamble to proving the steps, we argue that the FMFE is invariant to the scaling. Define the *budget-per-auction* as the ratio of budget to expected number of matching auctions during the campaign length, given by $g_\theta = b_\theta / (\alpha_\theta \eta s_\theta)$. Clearly, the budget-per-auction is invariant to the scaling.

From optimization problem (1) of the main paper, it is not hard to see that strategies depend solely on the budget-per-auction. Moreover, for all θ , both the expected number of matching bidders $\alpha_\theta \lambda s_\theta$ and the probability that a matching bidder is of that type $\mathbb{P}_\Theta\{\theta\}$ are invariant to the scaling. Hence, the scaling does not impact the equilibrium distribution of the maximum bid. These two facts imply that the FMFE is invariant to the scaling.

Outline. Theorem 1 is proven in two main steps, as outlined following the statement of the result. We first analyze the Budget-constrained Mean Field Model and the performance of the FMFE strategies in the latter. We then justify the mean field assumption through a propagation of chaos and fluid limit arguments. The required definitions and intermediary results are first presented in § 2 and § 3. The proof of the main result, Theorem 1 is provided in § A.1 and the proofs of the intermediary results are then presented in § A.2 and § A.3.

2 Budget-constrained Mean Field Model

In this section, we study a budget-constrained mean-field model (BMFM) in which advertisers are only allowed to bid when they have positive budgets. The main distinction between the real and the BMFM system is that, in the latter, all interactions are assumed to be independent.

BMFM Model. We study the performance of a fixed θ -type advertiser in the following mean-field system. We assume all advertisers (including the one in consideration) employ the FMFE strategy profile β^F . We refer to the advertiser in consideration as the *zeroth* advertiser. We assume that the zeroth bidder will participate in a random number of independent auctions over the course of his campaign, and the states of the competing matching bidders are independent across bidders, across auctions, and of the evolution of the zeroth advertiser's process. Let $X_\theta^{\text{MF}}(t) = (b_\theta(t), s_\theta(t)) \in \mathbb{R}_+^2$ denote the state of the zeroth advertiser at time t as given by the remaining budget $b_\theta(t)$ and the remaining time in system $s_\theta(t) = s_\theta - t$. The mean-field assumption implies that one need not keep track of the evolution of the market, and thus the process $X_\theta^{\text{MF}} = \{X_\theta^{\text{MF}}(t)\}_{t \in [0, s_\theta]}$ is Markov.

We next describe the evolution of the continuous time Markov process X_θ^{MF} . Initially, we have that $X_\theta^{\text{MF}}(0) = (b_\theta, s_\theta)$. The arrival of matching impressions corresponds to the jumps of a Poisson process $\{N_\theta(t)\}_{t \geq 0}$ with intensity $\alpha_\theta \eta$. We denote the sequence of jump times by $\{t_{\theta, n}\}_{n \geq 1}$. The number of competing matching bidders at the n -th auction, denoted by M_n , is drawn independently from a Poisson random variable with mean $\lambda \mathbb{E}[\alpha_\Theta s_\Theta]$. We denote by $e_{n, k} = (b_{n, k}, s_{n, k}, \theta_{n, k}) \in \mathbb{R}_+^3 \times \Theta$ the *extended* state of the k -th competing bidder in the n -th auction, which includes the relevant information to determine the agent's bid. The first component $b_{n, k}$ denotes the remaining budget, the second component $s_{n, k}$ denotes the remaining campaign length, and the last component $\theta_{n, k}$ denotes the type. The extended states of all competing bidders are drawn independently from a given distribution \mathbb{P}_e . Once the states

are revealed the realization of the values for the impression, are determined by $v_{n,k} = F_{\theta_{n,k}}^{-1}(u_{n,k})$, where $u_{n,k}$ are independent draws from a Uniform distribution with support $[0, 1]$ and $F_{\theta_{n,k}}$ is the valuation distribution of type $\theta_{n,k}$. Then, competing bids are determined by $w_{n,k} = \beta_{\theta_{n,k}}^F(v_{n,k})\mathbf{1}\{b_{n,k} > 0\}$, that is, bidders are allowed to bid only when they have a positive budget.³ Using these bids together with the bid of the zeroth advertiser $w_{n,0} = \beta_{\theta}^F(v_{n,0})\mathbf{1}\{b(t_{\theta,n}^-) > 0\}$, the exchange runs a second-price auction with reserve price r , and determines the allocation vector $x_{n,k}$ and payments $d_{n,k}$. The zeroth advertiser's budget is updated as $b_{\theta}(t_{\theta,n}) = b_{\theta}(t_{\theta,n}^-) - x_{n,0}d_{n,0}$.

In order to determine the evolution of the process X_{θ}^{MF} one needs to specify the distribution of the extended states \mathbb{P}_e . To make the dependence explicit we write the Markov process when extended states are drawn from \mathbb{P}_e as $X_{\theta}^{\text{MF}}(\mathbb{P}_e) = \{X_{\theta}^{\text{MF}}(t; \mathbb{P}_e)\}_{t \in [0, s_{\theta}]}$. Recall that in our model, the dynamics of the advertisers campaigns are governed by an $M/G/\infty$ queue. Then, the probability that a matching advertiser is of type θ is proportional to the arrival rate $p_{\theta}\lambda$, matching probability α_{θ} and campaign length s_{θ} . The latter implies that the steady-state probability that a competing advertiser is of type θ is $\mathbb{P}\{\hat{\Theta} = \theta\} = (p_{\theta}\alpha_{\theta}s_{\theta}) / \sum_{\theta'} p_{\theta'}\alpha_{\theta'}s_{\theta'}$. Additionally, given that the randomly sampled competing advertiser is of type θ , the advertiser can be at any point of her campaign with uniform probability, because arrivals are governed by a Poisson process. Thus motivated, we impose the following consistency requirement in the BMFM model: the distribution of a uniform sampling in time of the resulting mean-field process $X_{\theta}^{\text{MF}}(\mathbb{P}_e)$ of an advertiser of type θ competing against bidders sampled according to \mathbb{P}_e should be consistent with the distribution initially postulated \mathbb{P}_e . More formally, we define the notion of a consistent BMFM.

Definition 1 (Consistent BMFM). *A BMFM is said to be consistent if for any Borel-measurable set of states $\mathcal{X} \subset \mathbb{R}_+^2$, and type θ , the extended state measure \mathbb{P}_e satisfies*

$$\mathbb{P}_e\{\mathcal{X}, \theta\} = \mathbb{P}\{X_{\theta}^{\text{MF}}(U[0, s_{\theta}]; \mathbb{P}_e) \in \mathcal{X}; \hat{\Theta} = \theta\} \quad (1)$$

with $U[0, s_{\theta}]$ an independent uniform random variable with support $[0, s_{\theta}]$, and $\hat{\Theta}$ denoting the steady-state distribution of types in the system.

2.1 Existence of a consistent BMFM

The consistency equation (1) can be simplified by recognizing that the fluid-based strategies are independent of the state, and solely dependent on the realization of the values and the type. Thus, it suffices to know whether the competing bidders have a positive budget to determine their bids. Denote

³In this model a bidder's total expenditure may exceed her budget if at some point the payment exceeds the remaining budget. This assumption has a small impact on the performance of the system, but simplifies the analysis. The actual bid would be given by $w_{n,k} = \min\{b_{n,k}, \beta_{\theta_{n,k}}^F(v_{n,k})\}$.

by $a_{n,k} = \mathbf{1}\{b_{n,k} > 0\}$ the *active indicator*, which is one when the k -th advertiser of the n -th auction has a positive budget and zero otherwise. For our purpose here, we can reduce the extended state to $\{a_{n,k}, v_{n,k}, \theta_{n,k}\}$. In this formulation the distribution of the active indicator given a type θ is Bernoulli with success probability q_θ . Intuitively, the *active probability* q_θ denotes the likelihood that a θ -type bidder has positive budget at a uniformly random time of her campaign. Let $\mathbf{q} = \{q_\theta\}_{\theta \in \Theta}$ be a vector of active probabilities. Since values and types are independent, equation (1) implies the consistency of active probabilities, i.e., $q_\theta = \mathbb{P}\{b_\theta(U[0, s_\theta]; \mathbf{q}) > 0\}$. Moreover, using the fact that $U[0, s_\theta]$ is uniform and independent of the process one may write q_θ as the expected fraction of time that the advertiser has positive budget. Indeed,

$$q_\theta = \mathbb{P}\{b_\theta(U[0, s_\theta]; \mathbf{q}) > 0\} = \frac{1}{s_\theta} \int_0^{s_\theta} \mathbb{P}\{b_\theta(t; \mathbf{q}) > 0\} dt = \mathbb{E} \left[\frac{1}{s_\theta} \int_0^{s_\theta} \mathbf{1}\{b_\theta(t; \mathbf{q}) > 0\} dt \right]. \quad (2)$$

The next result establishes that the κ^{th} mean-field model is well defined in the sense that there always exists a consistent vector of probabilities \mathbf{q}^κ satisfying the fixed-point equation (2).

Proposition 1. *For every scaling κ , there exists a vector of active probabilities \mathbf{q}^κ satisfying the consistency equation (2). Moreover, the consistent probability distribution of extended states is given by $\mathbb{P}_e^\kappa\{\mathcal{X}, \theta\} = \mathbb{P}\{X_\theta^{\text{MF}(\kappa)}(U[0, s_\theta^\kappa]; \mathbf{q}^\kappa) \in \mathcal{X}; \hat{\Theta} = \theta\}$.*

To prove this proposition we first show that the right-hand side of equation 2 is continuous in \mathbf{q} , by using a coupling argument; and then conclude by invoking Brouwer's Fixed-Point Theorem. The previous result, however, does not exclude the existence of multiple distinct active probability vectors consistent with the BMFM.

2.2 Active Bidders

As the number of opportunities in the horizon increases, one would expect that advertisers deplete their budgets closer to the end of their campaign, and that the fraction of time bidders are active gets close to one. The next result shows that this conjecture is asymptotically correct, that is, as the scaling increases the vector of active probabilities converges to one. Additionally, we show that the distribution of the maximum competing bid in the κ^{th} consistent BMFM, denoted by D^κ , converges in distribution to the steady-state maximum bid D of the FMFE.⁴

Proposition 2. *Suppose that $r > 0$ and that there are at most two bidders' types. Every sequence of consistent active probability vectors $\{\mathbf{q}^\kappa\}_\kappa$ satisfies $\lim_{\kappa \rightarrow \infty} \|\mathbf{1} - \mathbf{q}^\kappa\|_\infty = 0$. Additionally, the dis-*

⁴We note that this result and the results in Subsection 2.3 only require that the number of auctions advertiser participates on and the budgets grow to infinity at the same rate; they do not require the arrival rate increases to infinity. The latter part of the scaling is required for the results in Section 3.

tribution of the maximum competing bid in the κ^{th} consistent BMFM, denoted by D^κ , converges in distribution to the steady-state maximum bid D of the FMFE.

The latter result justifies the underlying assumption of the FMFE that all bidders were active throughout their stay. We prove the result under the assumption that there are at most two types. This assumption is needed to show that $\mathbf{q} = 1$ is the unique consistent active probability in the limiting case. The argument in the proof of the result shows, however, how this assumption can be weakened under further technical conditions; in particular, by imposing that an appropriately defined Jacobian matrix is a P-matrix.

2.3 Payoff Evaluation in the BMFM

In this section we study the expected payoff of the zeroth advertiser in the budget-constrained mean-field model when all competing bidders follow the FMFE strategy. Let $J_\theta^{\text{MF}(\kappa)}(\beta; \beta^{\text{F}})$ be the expected payoff of a θ -type advertiser when the market evolves according to the BMFM when she implements strategy β , and the competing bidders implement the FMFE strategies β^{F} .

First, we provide an asymptotic lower bound for the normalized expected payoff of the FMFE strategy in any consistent BMFM. To do so, we define the normalized objective value of the fluid problem (1) of the main paper as $\bar{J}_\theta^{\text{F}}(F_d) \triangleq J_\theta^{\text{F}}(F_d)/(\alpha_\theta \eta s_\theta)$. We also define F_d as the distribution of the maximum bid in the FMFE.

Proposition 3. (Lower Bound). *Consider any consistent BMFM with scaling κ in which all competitor bidders follow the FMFE strategy β^{F} . Suppose that the zeroth advertiser of type θ implements the FMFE strategy β_θ^{F} . The expected payoff of the zeroth advertiser in the BMFM, denoted by $J_\theta^{\text{MF}(\kappa)}(\beta_\theta^{\text{F}}; \beta^{\text{F}})$, is lower bounded by*

$$\liminf_{\kappa \rightarrow \infty} \frac{1}{\alpha_\theta^\kappa \eta^\kappa s_\theta} J_\theta^{\text{MF}(\kappa)}(\beta_\theta^{\text{F}}; \beta^{\text{F}}) \geq \bar{J}_\theta^{\text{F}}(F_d).$$

The intuition underlying the proof of this result relies heavily on Proposition 2. By the latter, in any consistent BMFM, advertisers will be active for most of their campaign as the scale of the system increases. Given the latter, the proof revolves around lower bounding the zeroth advertiser's performance by its performance in an alternative system where it may bid when it runs out of budget, but pays a penalty of \bar{V} for any such bid. It is possible to show that the first result of Proposition 2 implies that as the scale κ increases, the penalties paid will be relatively "small", and hence the advertiser's performance, when normalized, is close to $\bar{J}_\theta^{\text{F}}(F_d^\kappa)$, which itself is close to $\bar{J}_\theta^{\text{F}}(F_d)$ (by the second part of Proposition 2).

Next, we upper bound the normalized expected payoff of *any* strategy in a consistent BMFM.

Proposition 4. (Upper Bound). *Consider any consistent BMFM with scaling κ in which all competitor bidders follow the FMFE strategy β^F . Suppose that the zeroth advertiser of type θ implements an alternative strategy β^κ . The expected payoff of the zeroth advertiser, denoted by $J_\theta^{\text{MF}(\kappa)}(\beta^\kappa; \beta^F)$, is upper bounded by*

$$\limsup_{\kappa \rightarrow \infty} \frac{1}{\alpha_\theta^\kappa \eta^\kappa s_\theta} J_\theta^{\text{MF}(\kappa)}(\beta^\kappa; \beta^F) \leq \bar{J}_\theta^F(F_d).$$

To prove the result, we first upper bound the performance of an arbitrary strategy by that of a strategy with the benefit of hindsight (which has complete knowledge of the future realizations of bids and values). This is akin to what is typically done in revenue management settings (see, e.g., Talluri and van Ryzin (1998)), with the exception that here, the competitive environment (which is the counterpart of the demand environment in RM settings) is endogenous and determined through the BMFM consistency requirement. In turn, we upper bound (asymptotically) the normalized performance of the hindsight strategy by the objective value of the normalized value of the fluid problem when D has the FMFE distribution. Here, the second part of Proposition 2 is once again key to ensure that the distribution of the maximum bid in a consistent BMFM converges to that postulated in the FMFE. The conjunction of the two propositions above imply that the FMFE strategy is near-optimal when an advertiser faces the competition induced by the BMFM.

3 Propagation of Chaos in the BMFM

The critical assumptions of the BMFM are that the actions of an advertiser do not affect the market, that the states of competitors are independent, and that the number of matching bidders in successive auctions is independent. However, in the actual system there are two effects that undermine the independence assumption. The first is an *interaction effect*. Because advertisers may interact between themselves more than once directly, their states and bids in the same and in successive auctions may be correlated. Even when two advertisers meet for the first time, their states may be correlated if both were influenced by a third advertiser in the past. The second is a *queueing effect*. Because of the queueing dynamics of the advertisers' arrival and departure process, the total number of advertisers in the exchange exhibits temporal correlation. As a consequence, the number of matching bidders in successive auctions may also be correlated.

The next result compares the expected performance of a strategy β^κ in the real system when all other advertisers implement the FMFE strategy, denoted by $J_\theta^\kappa(\beta^\kappa; \beta^F)$, to the performance of the same strategy in the BMFM, denoted by $J_\theta^{\text{MF}(\kappa)}(\beta^\kappa; \beta^F)$. The comparison is conducted under the assumption that, in the actual system, the initial states of the advertisers in the market are drawn independently from a consistent BMFM probability distribution. This initial conditions differ from the steady-state of the actual system, though one would expect them to be close as the scaling increases.

The initial conditions are as follows: (i) the number of bidders $Q(0)$ is Poisson with mean $\lambda^\kappa \mathbb{E}[s_\Theta]$; and (ii) the state of each advertisers is drawn independently from the measure \mathbb{P}_e^κ of a consistent BMFM.

Proposition 5. *Consider a κ -scaled market in which competitor bidders follow the FMFE strategy β^F . Initially, the number of advertisers is drawn from a Poisson distribution with mean $\lambda^\kappa \mathbb{E}[s_\Theta]$, and the state of each one of them is drawn independently from a consistent BMFM probability distribution \mathbb{P}_e^κ . Suppose that the zeroth advertiser of type θ arrives to the market at time zero, and implements a non-anticipating and adaptive strategy $\beta^\kappa \in \mathbb{B}$. The difference between the expected payoffs verifies*

$$\lim_{\kappa \rightarrow \infty} \frac{1}{\alpha_\theta^\kappa \eta^\kappa s_\theta} \left| J_\theta^\kappa(\beta^\kappa; \beta^F) - J_\theta^{\text{MF}(\kappa)}(\beta^\kappa; \beta^F) \right| = 0.$$

The result revolves around establishing that *i.*) with high probability an advertiser interacts with distinct advertisers during her campaign, and that the same applies recursively with those advertisers she competes with; and *ii.*) the queueing dynamics and their temporal correlation have little impact on the number of matching bidders, which intuitively follows from the fact that advertisers match at random with an impression. Thus, as the scaling increases the impact of the interaction effect and the queueing effect become negligible, the real system behavior is “close” to that in the BMFM, and the predictions in the BMFM carry over, in an appropriate sense, to the real system.

A difficulty in establishing the previous result is that in the AdX market the number of agents in the system is not fixed. Instead, advertisers arrive and depart from the market according to the dynamics of a $M/G/\infty$ queue; resulting in an *open* system. In order to analyze this system during a fixed time horizon $[0, T]$, we consider an alternate *closed* system in which all advertisers are present at time 0, but they are allowed to bid only during their campaigns, which start at a uniformly random time in the horizon. In this system, the number of advertisers originally present is random and equates to the number of arrivals during the horizon plus the number of advertisers currently running a campaign at time zero. For this purpose, in Section B we study a general mean-field model for closed systems with a random number of agents; an analysis that may be of independent interest. This construction allow us to appropriately extend previous propagation of chaos arguments by Graham and Méléard (1994) and Iyer et al. (2011) for closed systems with a *fixed* number of agents. In Section C we formally show that our AdX market can be modeled as such system.

References

- Aven, T. (1985), ‘Upper (lower) bounds on the mean of the maximum (minimum) of a number of random variables’, *Journal of Applied Probability* **22**(3), 723–728.
- Chauvin, B. and Giroux, G. (1990), ‘A geometric rate of convergence to the equilibrium for boltzmann processes with multiple particle interactions’, **27**(3), 510–520.

- Facchinei, F. and Pang, J.-S. (2003), *Finite-Dimensional Variational Inequalities and Complementarity Problems, Volume 1*, Springer Series in Operations Research and Financial Engineering, Springer.
- Graham, C. (1992), ‘Nonlinear diffusion with jumps’, *Annales de l’I.H.P. Probabilités et statistiques* **28**(3), 393–402.
- Graham, C. and Méléard, S. (1994), ‘Chaos hypothesis for a system interacting through shared resources’, *Probability Theory and Related Fields* **100**(2), 157–174.
- Iyer, K., Johari, R. and Sundararajan, M. (2011), Mean Field Equilibria of Dynamic Auctions with Learning. SSRN eLibrary.
- 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.

A Proofs

A.1 Proof of Theorem 1

Write the relative expected payoff of the deviation as

$$\begin{aligned} \frac{J_\theta^\kappa(\beta^\kappa, \beta^F)}{J_\theta^\kappa(\beta_\theta^F, \beta^F)} &= \frac{J_\theta^\kappa(\beta^\kappa, \beta^F) - J_\theta^{\text{MF}(\kappa)}(\beta^\kappa, \beta^F) + J_\theta^{\text{MF}(\kappa)}(\beta^\kappa, \beta^F)}{J_\theta^\kappa(\beta_\theta^F, \beta^F) - J_\theta^{\text{MF}(\kappa)}(\beta_\theta^F, \beta^F) + J_\theta^{\text{MF}(\kappa)}(\beta_\theta^F, \beta^F)} \\ &= \frac{(\alpha_\theta^\kappa \eta^\kappa s_\theta)^{-1} \left(J_\theta^\kappa(\beta^\kappa, \beta^F) - J_\theta^{\text{MF}(\kappa)}(\beta^\kappa, \beta^F) \right) + (\alpha_\theta^\kappa \eta^\kappa s_\theta)^{-1} J_\theta^{\text{MF}(\kappa)}(\beta^\kappa, \beta^F)}{(\alpha_\theta^\kappa \eta^\kappa s_\theta)^{-1} \left(J_\theta^\kappa(\beta_\theta^F, \beta^F) - J_\theta^{\text{MF}(\kappa)}(\beta_\theta^F, \beta^F) \right) + (\alpha_\theta^\kappa \eta^\kappa s_\theta)^{-1} J_\theta^{\text{MF}(\kappa)}(\beta_\theta^F, \beta^F)} \end{aligned}$$

Proposition 5 gives the convergence of the expected payoff under the real system to the expected payoff under the BMFM, which implies that the first term of the numerator and denominator converge to zero. Proposition 4 implies that the limsup of the second term of the numerator is bounded from above by $\bar{J}_\theta^F(F_d)$ and Proposition 3 implies that the liminf of the second term of the denominator is bounded from below by $\bar{J}_\theta^F(F_d)$. The result follows.

A.2 Proofs of Propositions 1 - 4

A.2.1 Proof of Proposition 1

The second part of the statement is direct. We prove the first one. Let $\mathbf{f} : [0, 1]^{|\Theta|} \rightarrow [0, 1]^{|\Theta|}$ be a mapping such that $f_\theta(\mathbf{q})$ determines the fraction of time that a zeroth θ -type bidder is active when the competing advertisers are active a fraction \mathbf{q} of their campaign. Like in equation (2) we have that

$$f_\theta(\mathbf{q}) = \mathbb{E} \left[\frac{1}{s_\theta} \int_0^{s_\theta} \mathbf{1}\{b_\theta(t; \mathbf{q}) > 0\} dt \right].$$

Notice that the domain of \mathbf{f} is compact, and coincides with its codomain. To show that the consistency equation $\mathbf{f}(\mathbf{q}) = \mathbf{q}$ admits a solution, it suffices to show that the functions $f_\theta(\mathbf{q})$ are continuous in \mathbf{q} and invoke Brouwer's Fixed-Point Theorem.

Next, we show that for each type θ the function $f_\theta(\mathbf{q})$ is Lipschitz continuous using a coupling argument. Fix $\theta \in \Theta$, and let \mathbf{q} and \mathbf{q}' be two distinct vectors of active probabilities. Let X_θ^{MF} and $X_\theta'^{\text{MF}}$ be the state processes in the BMFM when competing bidders are drawn according to \mathbf{q} and \mathbf{q}' , respectively. Consider a coupling Y_θ and Y_θ' of the processes in a common probability space such that both processes coincide in (i) the number of impressions, (ii) the realization of values of the zeroth advertisers, (iii) the number of matching bidders in each auction, and (iv) the types and values of the competing matching bidders. The processes only differ in the realization of the active indicators, which are distinct Bernoulli random variables coupled through a common uniform distribution. That is, for the k -th bidder of the n -th auction the active indicator in Y_θ is $a_{n,k} = \mathbf{1}\{U_{n,k} \leq q_{\theta_{n,k}}\}$, while for Y_θ' is $a'_{n,k} = \mathbf{1}\{U_{n,k} \leq q'_{\theta_{n,k}}\}$, with $U_{n,k}$ uniform in $[0, 1]$.

Notice that, by construction, the laws of the coupled processes coincide with the original ones, i.e., $\mathcal{L}(X_\theta^{\text{MF}}) = \mathcal{L}(Y_\theta)$ and $\mathcal{L}(X_\theta'^{\text{MF}}) = \mathcal{L}(Y_\theta')$. Let $A = \left\{ \bigcup_{n=1}^{N_\theta(s_\theta)} \bigcup_{k=1}^{M_n} a_{n,k} \neq a'_{n,k} \right\}$ be the event that some pair of active indicator differs, where $N_\theta(s_\theta)$ denoted the number of matching auctions for the zeroth advertiser during her campaign, and M_n denoted the number of matching bidders in the n -th auction. We have that the coupled processes coincide in the complement event \bar{A} , and thus the difference in total variation of the two processes satisfies

$$\|\mathcal{L}(X_\theta^{\text{MF}}) - \mathcal{L}(X_\theta'^{\text{MF}})\|_{TV} \leq \mathbb{P}\{Y_\theta \neq Y_\theta'\} \leq \mathbb{P}\{A\}.$$

To bound the probability of the event A , note that for the (n, k) -bidder with type $\theta_{n,k}$ the active indicator differs only if the uniform distribution lies within the interval $\left(\min\{q_{\theta_{n,k}}, q'_{\theta_{n,k}}\}, \max\{q_{\theta_{n,k}}, q'_{\theta_{n,k}}\} \right]$, an event occurring with probability $|q_{\theta_{n,k}} - q'_{\theta_{n,k}}|$. Taking expectations over the types, we have that $\mathbb{P}\{a_{n,k} \neq a'_{n,k}\} = \sum_{\theta \in \Theta} \mathbb{P}\{\hat{\Theta}_{n,k} = \theta\} |q_\theta - q'_\theta| \leq \|\mathbf{q} - \mathbf{q}'\|_\infty$. Using a union bound together with the independence assumption on the primitives, one may bound the probability of the event A by

$$\begin{aligned} \mathbb{P}\{A\} &\leq \mathbb{E}_{N_\theta(s_\theta), M} \left[\sum_{n=1}^{N_\theta(s_\theta)} \sum_{k=1}^{M_n} \mathbb{P}\{a_{n,k} \neq a'_{n,k}\} \right] \leq \|\mathbf{q} - \mathbf{q}'\|_\infty \mathbb{E} \left[\sum_{n=1}^{N_\theta(s_\theta)} M_n \right] \\ &= \|\mathbf{q} - \mathbf{q}'\|_\infty (\alpha_\theta \eta s_\theta) (\lambda \mathbb{E}[\alpha_\Theta s_\Theta]). \end{aligned}$$

We conclude by noting that $f_\theta(\cdot) \in [0, 1]$ to get that $|f_\theta(\mathbf{q}) - f_\theta(\mathbf{q}')| \leq |1-0| \cdot \|\mathcal{L}(X_\theta^{\text{MF}}) - \mathcal{L}(X_\theta'^{\text{MF}})\|_{TV} \leq C \|\mathbf{q} - \mathbf{q}'\|_\infty$, with $C = (\alpha_\theta \eta s_\theta) (\lambda \mathbb{E}[\alpha_\Theta s_\Theta]) < \infty$.

A.2.2 Proof of Proposition 2

We prove the results in four steps. First, we show that the sequence of functions $\{\mathbf{f}^\kappa\}_\kappa$ converges point-wise to a continuous function \mathbf{f}^∞ . Second, we show that the unique fixed point of the function \mathbf{f}^∞ is 1, that is, in the limit all advertisers are active. Third, we prove that all fixed-points of the functions $\{\mathbf{f}^\kappa\}_\kappa$ converge to the unique fixed point of \mathbf{f}^∞ . Fourth, we prove the convergence in distribution of the maximum bid.

Step 1 (The point-wise convergence). Fix a type θ , and the active probability vector \mathbf{q} . Consider a coupled process $Y_\theta^\kappa(\mathbf{q}) = \{Y_\theta^\kappa(t; \mathbf{q})\}_{t \in [0, \infty)}$ in which the advertiser is allowed to bid beyond the length of her campaign so that the laws of $X_\theta^{\text{MF}(\kappa)}(\mathbf{q})$, and $Y_\theta^\kappa(\mathbf{q})$ coincide for $t \in [0, s_\theta]$. Notice that $f_\theta^\kappa(\mathbf{q}) = \mathbb{E} \left[\min\{\tilde{S}_\theta^\kappa(\mathbf{q})/s_\theta, 1\} \right]$, where $\tilde{S}_\theta^\kappa(\mathbf{q}) = \inf\{s \geq 0 : b_\theta^\kappa(s) \leq 0\}$ is the first time that the budget is non-positive (defined with respect to the process $Y_\theta^\kappa(\mathbf{q})$).

In order to study the hitting time $\tilde{S}^\kappa(\mathbf{q})$ we consider the sequence $\{Z_{\theta,n}(\mathbf{q})\}_{n \geq 1}$ of expenditures of the zeroth advertiser in each auction when the active probability vector is \mathbf{q} . In view of our mean-field assumption the sequence of expenditures is i.i.d. and independent of the impressions' inter-arrival times. Before proceeding we characterize the maximum bid and the expenditure in the BMFM as a function of the active probability vector. The maximum competing bid at the n -th auction is given by $D_{\theta,n}(\mathbf{q}) = \max \left(\left\{ \left\{ \beta_\theta(V_{n,k}) \right\}_{k=1}^{M_{\theta,n}^a(\mathbf{q})} \right\}_{\theta \in \Theta}, r \right)$, where $M_{\theta,n}^a(\mathbf{q})$ denotes the number of matching bidders of type θ with positive budget, which is distributed as a Poisson random variable with mean $p_\theta q_\theta \alpha_\theta \lambda s_\theta$ (where $p_\theta = \mathbb{P}_\Theta\{\theta\}$) since each advertiser is active independently with probability q_θ , and type θ advertisers arrive to the exchange with rate $p_\theta \lambda$. In this notation we have that the expenditure of the zeroth bidder in the n -th auction is $Z_{\theta,n}(\mathbf{q}) = \mathbf{1}\{D_{\theta,n}(\mathbf{q}) \leq \beta_\theta(V_{\theta,n})\} D_{\theta,n}(\mathbf{q})$, where $V_{\theta,n}$ is a drawn of the zeroth advertiser value. Notice that, for a fixed active probability vector, both the distribution of the maximum bid and of the zeroth advertiser's expenditure are *invariant* to the scaling.

In the following we drop the dependence on \mathbf{q} . Let $C_{\theta,n} = \sum_{j=1}^n Z_{\theta,j}$ denote the cumulative expenditure incurred after the n -th auction, and let $\tilde{N}_\theta^\kappa = \inf\{n \geq 1 : C_{\theta,n} \geq b_\theta^\kappa\}$ be the number of auctions until the cumulative expenditure exceeds the budget b_θ^κ , which is a stopping time for the sequence. Since expenditures are bounded, $Z_{\theta,j} \leq \bar{V} < \infty$ a.s., the cumulative expenditure at the stopping time can be bounded from below and above by

$$b_\theta^\kappa \leq C_{\theta, \tilde{N}_\theta^\kappa} \leq b_\theta^\kappa + \bar{V}.$$

Dividing by the expected number of impressions on the campaign $\alpha_\theta^\kappa \eta^\kappa s_\theta$ we obtain that

$$g_\theta \leq \frac{\tilde{N}_\theta^\kappa}{\alpha_\theta^\kappa \eta^\kappa s_\theta} \frac{1}{\tilde{N}_\theta^\kappa} \sum_{j=1}^{\tilde{N}_\theta^\kappa} Z_{\theta,j} \leq g_\theta + \frac{\bar{V}}{\alpha_\theta^\kappa \eta^\kappa s_\theta}.$$

Note that $\lim_{\kappa \rightarrow \infty} \tilde{N}_\theta^\kappa = \infty$ almost surely since $b_\theta^\kappa \leq C_{\theta, N_\theta^\kappa} \leq \tilde{N}_\theta^\kappa \bar{V}$, and $\lim_{\kappa \rightarrow \infty} b_\theta^\kappa = \infty$. Hence, taking the limit as $\kappa \rightarrow \infty$ and using the Strong Law of Large Numbers (SLLN), one obtains that $(1/\tilde{N}_\theta^\kappa) \sum_{j=1}^{\tilde{N}_\theta^\kappa} Z_{\theta, j}$ converges to $\mathbb{E}Z_\theta$ a.s. In turn, we obtain that $\tilde{N}_\theta^\kappa / (\alpha_\theta^\kappa \eta^\kappa s_\theta) \rightarrow g_\theta / \mathbb{E}Z_\theta$ a.s.

Next, notice that \tilde{S}_θ^κ is a sum of a random number \tilde{N}_θ^κ of exponential random variables. More formally, $\tilde{S}_\theta^\kappa = t_{\theta, \tilde{N}_\theta^\kappa}^\kappa = \sum_{n=1}^{\tilde{N}_\theta^\kappa} (t_{\theta, n}^\kappa - t_{\theta, n-1}^\kappa)$, where $t_{\theta, n}^\kappa - t_{\theta, n-1}^\kappa$ is the inter-arrival time of the n -th matching impression for the zeroth advertiser. Since inter-arrival times are independent of \tilde{N}_θ^κ and exponentially distributed with rate $\alpha_\theta^\kappa \eta^\kappa$, we may invoke the SLLN again to obtain that

$$\frac{\tilde{S}_\theta^\kappa}{s_\theta} = \frac{\tilde{N}_\theta^\kappa}{\alpha_\theta^\kappa \eta^\kappa s_\theta} \frac{1}{\tilde{N}_\theta^\kappa} \sum_{n=1}^{\tilde{N}_\theta^\kappa} \alpha_\theta^\kappa \eta^\kappa (t_{\theta, n}^\kappa - t_{\theta, n-1}^\kappa) \rightarrow \frac{g_\theta}{\mathbb{E}Z_\theta}, \quad \text{a.s. as } \kappa \rightarrow \infty.$$

The Dominated Convergence Theorem enables to conclude that

$$f_\theta^\infty = \lim_{\kappa \rightarrow \infty} f_\theta^\kappa = \lim_{\kappa \rightarrow \infty} \mathbb{E} \left[\min\{\tilde{S}_\theta^\kappa / s_\theta, 1\} \right] = \mathbb{E} \left[\min\left\{ \lim_{\kappa \rightarrow \infty} \tilde{S}_\theta^\kappa / s_\theta, 1 \right\} \right] = \min \left\{ \frac{g_\theta}{\mathbb{E}Z_\theta}, 1 \right\},$$

point-wise in all active probability vectors \mathbf{q} .

Step 2 (Fixed-points of \mathbf{f}^∞). In this section we study the fixed-points of the limit function \mathbf{f}^∞ , and show that 1 is the unique fixed-point of the mapping. We proceed by considering the related functions $H_\theta(\mathbf{q}) = q_\theta \mathbb{E}[Z_\theta(\mathbf{q})]$, and using the fact that the set of fixed-points of the function $f_\theta^\infty(\mathbf{q}) = \min\{g_\theta / \mathbb{E}[Z_\theta(\mathbf{q})], 1\}$ coincide with the solutions of the NCP

$$H_\theta(\mathbf{q}) \leq g_\theta \quad \perp \quad 0 \leq q_\theta \leq 1, \quad \forall \theta \in \Theta, \quad (3)$$

where the complementary condition is with the inequality $q_\theta \leq 1$. Note further that $G_\theta(\boldsymbol{\mu}) = H_\theta(1)$, and thus from the equilibrium condition of the FMFE we get that 1 is a solution of (3). Additionally, it is not hard to show that the functions

$$H_\theta(\mathbf{q}) = q_\theta \mathbb{E}[Z_\theta(\mathbf{q})] = q_\theta r \bar{F}_v((1 + \mu_\theta)r) F_d(r; \mathbf{q}) + q_\theta \int_r^{\bar{V}} x \bar{F}_v((1 + \mu_\theta)x) dF_d(x; \mathbf{q}),$$

are differentiable. Also, by Lemma 1 (stated and proved in Appendix D), the Jacobian of \mathbf{H} is a P-matrix. Then, by Facchinei and Pang (2003, Proposition 3.5.10) we conclude that 1 is the unique vector of active probabilities that solves (3), and thus the unique fixed-point of \mathbf{f}^∞ .

Step 3 (Convergence of the fixed-points). Let $\{\mathbf{q}^\kappa\}_\kappa$ be a sequence of fixed-points of the sequence of functions $\{\mathbf{f}^\kappa\}_\kappa$, i.e., $\mathbf{f}^\kappa(\mathbf{q}^\kappa) = \mathbf{q}^\kappa$ for every scaling κ . If the convergence of the sequence of functions to \mathbf{f}^∞ is *uniform*, together with the continuity of the mapping, one would be able to invoke Lemma 2 (stated and proved in Appendix D) to conclude that these fixed-points converge to the unique

fixed-point of \mathbf{f}^∞ , i.e., $\lim_{\kappa \rightarrow \infty} \|1 - \mathbf{q}^\kappa\| = 0$. Next, we prove that the mappings converge uniformly by showing that the sequence of functions is uniformly Cauchy.

Fix a vector of active probabilities \mathbf{q} , and let κ, κ' be two different scalings. We may bound the difference between two different scalings as follows

$$\begin{aligned}
|f_\theta^\kappa(\mathbf{q}) - f_\theta^{\kappa'}(\mathbf{q})| &= \left| \mathbb{E} \left[\min\{\tilde{S}_\theta^\kappa(\mathbf{q})/s_\theta, 1\} \right] - \mathbb{E} \left[\min\{\tilde{S}_\theta^{\kappa'}(\mathbf{q})/s_\theta, 1\} \right] \right| \\
&\leq \mathbb{E} \left| \min\{\tilde{S}_\theta^\kappa(\mathbf{q})/s_\theta, 1\} - \min\{\tilde{S}_\theta^{\kappa'}(\mathbf{q})/s_\theta, 1\} \right| \\
&\leq \mathbb{E} \left| \tilde{S}_\theta^\kappa(\mathbf{q})/s_\theta - \tilde{S}_\theta^{\kappa'}(\mathbf{q})/s_\theta \right| \\
&\leq \frac{1}{s_\theta} \mathbb{E} \left| \tilde{S}_\theta^\kappa(\mathbf{q}) - \mathbb{E} \left[\tilde{S}_\theta^\kappa(\mathbf{q}) \right] \right| + \frac{1}{s_\theta} \mathbb{E} \left| \tilde{S}_\theta^{\kappa'}(\mathbf{q}) - \mathbb{E} \left[\tilde{S}_\theta^{\kappa'}(\mathbf{q}) \right] \right| + \frac{1}{s_\theta} \left| \mathbb{E} \left[\tilde{S}_\theta^\kappa(\mathbf{q}) \right] - \mathbb{E} \left[\tilde{S}_\theta^{\kappa'}(\mathbf{q}) \right] \right| \\
&\leq \frac{1}{s_\theta} \sqrt{\text{Var}[\tilde{S}_\theta^\kappa(\mathbf{q})]} + \frac{1}{s_\theta} \sqrt{\text{Var}[\tilde{S}_\theta^{\kappa'}(\mathbf{q})]} + \frac{1}{s_\theta} \left| \mathbb{E} \left[\tilde{S}_\theta^\kappa(\mathbf{q}) \right] - \mathbb{E} \left[\tilde{S}_\theta^{\kappa'}(\mathbf{q}) \right] \right|, \tag{4}
\end{aligned}$$

where the first inequality follows from the convexity of the absolute value and Jensen's inequality, the second from the fact that $\min\{x, 1\}$ is Lipschitz continuous with constant 1, the third from the triangular inequality; and the fourth from Lyapunov's inequality. We next turn to the problem of bounding the mean and variance of the hitting time \tilde{S}_θ^κ .

Note that with positive probability the advertiser spends at least $r > 0$ and thus $\mathbb{E}\tilde{N}_\theta < \infty$. Hence, we may employ Wald's identities to bound the mean and variance of the stopping time \tilde{N}_θ . First, Wald's first identity implies that $\mathbb{E}[C_{\theta, \tilde{N}_\theta}] = \mathbb{E}\tilde{N}_\theta \mathbb{E}Z_\theta$. Using the fact that $C_{\theta, \tilde{N}_\theta} \geq b_\theta$, one obtains that the mean is bounded from below by $\mathbb{E}[\tilde{N}_\theta] \geq b_\theta/\mathbb{E}[Z_\theta]$. Using the fact that $C_{\theta, \tilde{N}_\theta} \leq b_\theta + \bar{V}$, one may also bound the mean from above by $\mathbb{E}[\tilde{N}_\theta] \leq (b_\theta + \bar{V})/\mathbb{E}[Z_\theta]$. Second, the variance is bounded from above by $\text{Var}(\tilde{N}_\theta) \leq (b_\theta + \bar{V})\text{Var}(Z_\theta)/\mathbb{E}[Z_\theta]^3 + \bar{V}^2/\mathbb{E}[Z_\theta]^2$ (use Wald's second identity to get $\mathbb{E}[C_{\theta, \tilde{N}_\theta} - \tilde{N}_\theta \mathbb{E}Z_\theta]^2 = \text{Var}(Z_\theta)\mathbb{E}\tilde{N}_\theta$). Next, recall that \tilde{S}_θ is a sum of a random number \tilde{N}_θ of independent exponential random variables. Thus, we obtain by taking conditional expectations that $\mathbb{E}[\tilde{S}_\theta] = (\alpha_\theta \eta)^{-1} \mathbb{E}[\tilde{N}_\theta]$, and $\text{Var}[\tilde{S}_\theta] = (\alpha_\theta \eta)^{-2} (\mathbb{E}[\tilde{N}_\theta] + \text{Var}[\tilde{N}_\theta])$ (see, e.g., Ross (1996, pp.22)).

Thus, we have that the mean of the hitting time is bounded as follows

$$\frac{g_\theta}{\mathbb{E}[Z_\theta(\mathbf{q})]} \leq \frac{\mathbb{E}[\tilde{S}_\theta^\kappa(\mathbf{q})]}{s_\theta} \leq \frac{g_\theta}{\mathbb{E}[Z_\theta(\mathbf{q})]} \left(1 + \frac{\bar{V}}{b_\theta^\kappa} \right)$$

where $g_\theta = b_\theta^\kappa / (\alpha_\theta^\kappa \eta^\kappa s_\theta)$ is the budget-per-auction for type θ , which is *invariant* to the scaling. Simi-

larly, the variance can be upper bounded by

$$\begin{aligned} \frac{\text{Var}[\tilde{S}_\theta^\kappa(\mathbf{q})]}{s_\theta^2} &= \frac{1}{(\alpha_\theta^\kappa \eta^\kappa s_\theta)^2} \left(\mathbb{E}[\tilde{N}_\theta^\kappa(\mathbf{q})] + \text{Var}[\tilde{N}_\theta^\kappa(\mathbf{q})] \right) \\ &\leq \frac{b_\theta^\kappa + \bar{V}}{(\alpha_\theta^\kappa \eta^\kappa s_\theta)^2} \frac{\text{Var}[Z_\theta(\mathbf{q})] + \mathbb{E}[Z_\theta(\mathbf{q})]^2}{\mathbb{E}[Z_\theta(\mathbf{q})]^3} + \frac{\bar{V}^2}{(\alpha_\theta^\kappa \eta^\kappa s_\theta)^2} \frac{1}{\mathbb{E}[Z_\theta(\mathbf{q})]^2} \leq \frac{K}{b_\theta^\kappa} \end{aligned}$$

for some $K > 0$ independent of the scaling, the vector of active probabilities, and the type. The last follows from the facts that (i) $\text{Var}[Z_\theta(\mathbf{q})] \leq \bar{V}^2/4$ and $\mathbb{E}[Z_\theta(\mathbf{q})] \leq \bar{V}$ because $0 \leq Z_\theta(\mathbf{q}) \leq \bar{V}$ almost surely; (ii) for sufficiently large scaling we have that $\bar{V} \leq b_\theta^\kappa$; and (iii) because the reserve price is strictly positive and there is a positive probability that the advertiser wins the auction the expected expenditure can never drop to zero, i.e., $\inf_{0 \leq \mathbf{q} \leq 1} \mathbb{E}[Z_\theta(\mathbf{q})] > 0$.

Combining the last bounds we obtain that the difference in (4) is bounded by

$$\begin{aligned} |f_\theta^\kappa(\mathbf{q}) - f_\theta^{\kappa'}(\mathbf{q})| &\leq \sqrt{\frac{K}{b_\theta^\kappa}} + \sqrt{\frac{K}{b_\theta^{\kappa'}}} + \frac{\bar{V} g_\theta}{\mathbb{E}[Z_\theta(\mathbf{q})]} \left| \frac{1}{b_\theta^\kappa} - \frac{1}{b_\theta^{\kappa'}} \right| \\ &\leq \sqrt{\frac{K}{b_\theta^\kappa}} + \sqrt{\frac{K}{b_\theta^{\kappa'}}} + \frac{K'}{b_\theta^\kappa} + \frac{K'}{b_\theta^{\kappa'}}, \end{aligned}$$

where the second bound follows from the triangle inequality and property (iii) from above. Since the Cauchy difference converges to zero as $\kappa, \kappa' \rightarrow \infty$ uniformly in \mathbf{q} , we get that the sequence of functions is uniformly convergent.

Step 4 (Convergence in distribution of the maximum bid.) Let \mathbf{q}^κ be a consistent probability vector of the κ -th mean field system. The cumulative distribution function of $D(\mathbf{q}^\kappa)$ for any $x \geq r$ is given by

$$F_d^\kappa(x) = \mathbb{P}\{D(\mathbf{q}^\kappa) \leq x\} = \exp\left(-\lambda^\kappa \sum_{\theta \in \Theta} p_\theta q_\theta^\kappa \alpha_\theta^\kappa s_\theta \bar{F}_{v_\theta}((1 + \mu_\theta)x)\right)$$

which converges to the FMFE distribution of the maximum bid for all continuity points, since $\|1 - \mathbf{q}^\kappa\| \rightarrow 0$, and $\lambda^\kappa \alpha_\theta^\kappa$ is invariant to the scaling.

A.2.3 Proof of Proposition 3

Fix a type θ and the scaling κ . Let \mathbf{q}^κ be a consistent vector of active probabilities for the κ -th scaling (which exists according to Proposition 1). As in the proof of Proposition 2 we consider the coupled process $Y_\theta^\kappa(\mathbf{q}^\kappa)$ in which the zeroth advertiser is allowed to bid beyond the length of her campaign.

Let $\{(Z_{\theta,n}(\mathbf{q}), U_{\theta,n}(\mathbf{q}))\}_{n \geq 1}$ be the sequence of realized expenditures and utilities of the zeroth

advertiser in each auction when the vector of active probabilities is \mathbf{q} , which in view of our mean-field assumption in the BMFM is i.i.d. The zeroth advertiser's utility in the n -th auction is $U_{\theta,n}(\mathbf{q}) = \mathbf{1}\{D_{\theta,n}(\mathbf{q}) \leq \beta_{\theta}(V_{\theta,n})\}(V_n - D_{\theta,n}(\mathbf{q}))$. Again, it is the case that, for a fixed active probability vector, the distribution of the utility is *invariant* to the scaling. Moreover, using the fact that $\beta_{\theta}^F(x) = x/(1 + \mu_{\theta})$ we get that,

$$U_{\theta,n}(\mathbf{q}) = (V_n - (1 + \mu_{\theta})D_{\theta,n}(\mathbf{q}))^+ + \mu_{\theta}Z_{\theta,n}(\mathbf{q}),$$

which after taking expectations implies that $\mathbb{E}[U_{\theta}(\mathbf{q})] = \bar{\Psi}_{\theta}(\mu_{\theta}; F_d(\mathbf{q})) + \mu_{\theta}(\mathbb{E}[Z_{\theta}(\mathbf{q})] - g_{\theta})$, where the normalized dual function is defined as $\bar{\Psi}_{\theta}(\mu; F_d) \triangleq \Psi_{\theta}(\mu; F_d)/(\alpha_{\theta}\eta s_{\theta}) = \mathbb{E}[V - (1 + \mu)D]^+ + \mu g_{\theta}$, and $F_d(\mathbf{q})$ is the distribution of the maximum of the competitors' bids for a given vector \mathbf{q} of active probabilities.

Next, we lower bound the expected payoff of the zeroth advertiser. Recalling that $N_{\theta}^{\kappa}(s_{\theta})$ is the number of auctions the zeroth advertiser participates during her campaign, and $\tilde{N}_{\theta}^{\kappa}(\mathbf{q}^{\kappa})$ is the number of auctions until the cumulative expenditure exceeds the budget, we have that

$$\begin{aligned} J_{\theta}^{\text{MF}(\kappa)}(\beta^F; \beta^F) &= \mathbb{E} \left[\sum_{n=1}^{\tilde{N}_{\theta}^{\kappa}(\mathbf{q}^{\kappa}) \wedge N_{\theta}^{\kappa}(s_{\theta})} U_{\theta,n}(\mathbf{q}^{\kappa}) \right] \geq \mathbb{E} \left[\sum_{n=1}^{N_{\theta}^{\kappa}(s_{\theta})} U_{\theta,n}(\mathbf{q}^{\kappa}) \right] - \bar{V} \mathbb{E}[N_{\theta}^{\kappa}(s_{\theta}) - \tilde{N}_{\theta}^{\kappa}(\mathbf{q}^{\kappa})]^+ \\ &\geq \mathbb{E} \left[\sum_{n=1}^{N_{\theta}^{\kappa}(s_{\theta})} U_{\theta,n}(\mathbf{q}^{\kappa}) \right] - \bar{V} \mathbb{E}[N_{\theta}^{\kappa}(s_{\theta}) - \alpha_{\theta}^{\kappa} \eta^{\kappa} s_{\theta}]^+ - \bar{V} \mathbb{E}[\alpha_{\theta}^{\kappa} \eta^{\kappa} s_{\theta} - \tilde{N}_{\theta}^{\kappa}(\mathbf{q}^{\kappa})]^+, \end{aligned}$$

where the first inequality follows from the fact that $0 \leq U_{\theta,n}(\mathbf{q}) \leq \bar{V}$; and the second from the fact that for every $a, b, c \in \mathbb{R}$ we have that $(a - c)^+ \leq (a - b)^+ + (b - c)^+$. In the remainder of the proof we will show that, the first term on the right-hand side, normalized by the expected number of auctions, converges to $\bar{J}_{\theta}^F(F_d)$, and the second and last terms to zero. We study one term at a time.

For the first term, notice that the number of matching impressions is independent of the utility, and thus

$$\frac{1}{\alpha_{\theta}^{\kappa} \eta^{\kappa} s_{\theta}} \mathbb{E} \left[\sum_{n=1}^{N_{\theta}^{\kappa}(s_{\theta})} U_{\theta,n}(\mathbf{q}^{\kappa}) \right] = \mathbb{E}[U_{\theta}(\mathbf{q}^{\kappa})] = \bar{\Psi}_{\theta}(\mu_{\theta}; F_d^{\kappa}) + \mu_{\theta}(\mathbb{E}[Z_{\theta}(\mathbf{q}^{\kappa})] - g_{\theta}).$$

Notice that $\bar{\Psi}_{\theta}(\mu_{\theta}; F_d^{\kappa}) \rightarrow \bar{\Psi}_{\theta}(\mu_{\theta}; F_d)$ as $\kappa \rightarrow \infty$, since $F_d^{\kappa} \Rightarrow F_d$ from Proposition 2, and $\bar{\Psi}_{\theta}$ is continuous w.r.t. the distribution of the maximum bid from the proof of Theorem 1 of the main paper. Furthermore, since μ_{θ} is an optimal dual variable we get that $\bar{\Psi}_{\theta}(\mu_{\theta}; F_d) = \bar{J}_{\theta}^F(F_d)$, in view of Proposition 1 of the main paper. Additionally, from Proposition 2 we have $\mathbb{E}[Z_{\theta}(\mathbf{q})]$ is continuous in \mathbf{q} , and thus $\mathbb{E}[Z_{\theta}(\mathbf{q}^{\kappa})] \rightarrow \mathbb{E}[Z_{\theta}(1)]$ as $\kappa \rightarrow \infty$. From the complementarity condition between the equilibrium multiplier μ_{θ} and the expected expenditure $\mathbb{E}[Z_{\theta}(1)]$ of the FMFE, we get that last term

goes to zero.

For the second term, note that for any random variable X and constant x , we have that $\mathbb{E}(X - x)^+ \leq (\mathbb{E}X - x)^+ + \sqrt{\text{Var}(X)/2}$, by the upper bound on the maximum of random variables given in Aven (1985). Since $N_\theta^\kappa(s_\theta)$ is Poisson with mean $\alpha_\theta^\kappa \eta^\kappa s_\theta$ we get that

$$\frac{1}{\alpha_\theta^\kappa \eta^\kappa s_\theta} \mathbb{E}[N_\theta^\kappa(s_\theta) - \alpha_\theta^\kappa \eta^\kappa s_\theta]^+ \leq (2\alpha_\theta^\kappa \eta^\kappa s_\theta)^{-1/2},$$

with the right-hand side converging to zero as the scaling increases.

For the third term, we use a similar bound on the expected value of the maximum together with the bounds on the mean and variance of the hitting time $\tilde{N}_\theta^\kappa(\mathbf{q})$ developed in Proposition 2 to get

$$\begin{aligned} \frac{1}{\alpha_\theta^\kappa \eta^\kappa s_\theta} \mathbb{E}[\alpha_\theta^\kappa \eta^\kappa s_\theta^\kappa - \tilde{N}_\theta^\kappa(\mathbf{q}^\kappa)]^+ &\leq \left(1 - \frac{\mathbb{E}[\tilde{N}_\theta^\kappa(\mathbf{q}^\kappa)]}{\alpha_\theta^\kappa \eta^\kappa s_\theta}\right)^+ + \sqrt{\frac{\text{Var}[\tilde{N}_\theta^\kappa(\mathbf{q}^\kappa)]}{2(\alpha_\theta^\kappa \eta^\kappa s_\theta)^2}} \\ &\leq \left(1 - \frac{g_\theta}{\mathbb{E}[Z_\theta(\mathbf{q}^\kappa)]}\right)^+ + \sqrt{g_\theta \frac{\text{Var}[Z_\theta(\mathbf{q}^\kappa)] + \bar{V} \mathbb{E}[Z_\theta(\mathbf{q}^\kappa)]}{\mathbb{E}[Z_\theta(\mathbf{q}^\kappa)]^3}} (\alpha_\theta^\kappa \eta^\kappa s_\theta)^{-1/2}. \end{aligned}$$

The first term of the right-hand side converges to $(1 - g_\theta/\mathbb{E}[Z_\theta(1)])^+ \leq 0$, since the expected expenditure in the FMFE never exceeds the budget, that is, $\mathbb{E}[Z_\theta(1)] \leq g_\theta$. The last term of the right-hand side follows by the previous bound on $\text{Var}[\tilde{N}_\theta^\kappa(\mathbf{q}^\kappa)]$ and the fact that $\bar{V} \leq b_\theta^\kappa$ for large enough κ , and it converges to zero.

A.2.4 Proof of Proposition 4

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

Fix a type θ and a scaling κ . Let $J_\theta^{\text{MF}(\kappa)}(\beta^H, \beta^F)$ denote the expected payoff under perfect hindsight, which is obtained by looking at the optimal expected payoff when the realization of the number of impressions, the competing bids and values for the whole horizon is revealed up-front. No strategy can perform better than the perfect hindsight strategy β^H and we have that

$$J_\theta^{\text{MF}(\kappa)}(\beta; \beta^F) \leq J_\theta^{\text{MF}(\kappa)}(\beta^H; \beta^F).$$

Given a sample path ω , which determines the number of matching impressions $N_\theta^\kappa(s_\theta)(\omega) = n_\theta$ and the realization of the competing bids and values $\{(D_{n,0}(\omega), V_{n,0}(\omega))\}_{n=1}^{N_\theta^\kappa(s_\theta)(\omega)} = \{(d_{n,0}, v_{n,0})\}_{n=1}^{n_\theta}$, the advertiser only needs to determine which auctions to win (since bidding an amount $\epsilon > 0$ larger than

the maximum bid guarantees her winning the auction). Let the decision variable $x_{n,0} \in \{0, 1\}$ indicate whether the zeroth advertiser decides to win the auction or not. In hindsight, the zeroth advertiser needs to solve, for each realization ω , the following knapsack problem

$$J_\theta^{\text{H}(\kappa)}(\omega) = \max_{x_{n,0} \in \{0,1\}} \sum_{n=1}^{n_\theta} x_{n,0}(v_{n,0} - d_{n,0}) \quad (5a)$$

$$\text{s.t. } \sum_{n=1}^{n_\theta} x_{n,0}d_{n,0} \leq b_\theta. \quad (5b)$$

The perfect hindsight bound is obtained by averaging over all possible realizations consistently with the strategy of the other bidders and the BMFM, or equivalently $J_\theta^{\text{MF}(\kappa)}(\beta^{\text{H}}, \beta^{\text{F}}) = \mathbb{E}_\omega [J_\theta^{\text{H}(\kappa)}(\omega)]$.

Consider the continuous relaxation of the hindsight program (5) in which we replace the integrality constraints by $0 \leq x_{n,0} \leq 1$. Let μ_θ be the equilibrium multiplier of the FMFE for type θ . Introducing dual variables $\mu \geq 0$ for the budget constraint and $z_n \geq 0$ for the constraints $x_{n,0} \leq 1$, we get by weak duality that

$$\begin{aligned} J_\theta^{\text{H}(\kappa)}(\omega) &\leq \min_{\mu \geq 0, z_n \geq 0} \left\{ \sum_{n=1}^{n_\theta} z_n + \mu b_\theta \text{ s.t. } z_n \geq v_{n,0} - (1 + \mu)d_{n,0}, \forall n = 1, \dots, n_\theta \right\} \\ &= \min_{\mu \geq 0} \left\{ \sum_{n=1}^{n_\theta} [v_{n,0} - (1 + \mu)d_{n,0}]^+ + \mu b_\theta \right\} \\ &\leq \sum_{n=1}^{n_\theta} [v_{n,0} - (1 + \mu_\theta)d_{n,0}]^+ + \mu_\theta b_\theta \end{aligned}$$

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

$$\frac{1}{\alpha_\theta^\kappa \eta^\kappa s_\theta} J_\theta^{\text{MF}(\kappa)}(\beta^{\text{H}}; \beta^{\text{F}}) \leq \bar{\Psi}_\theta(\mu_\theta; F_d^\kappa).$$

We conclude by noting that $\lim_{\kappa \rightarrow \infty} \bar{\Psi}_\theta(\mu_\theta; F_d^\kappa) = \bar{J}_\theta^{\text{F}}(F_d)$ as in the proof of Proposition 3.

A.3 Proof of Proposition 5

Section C shows that the AdX market may be modeled as a closed system with a random number of agents. Furthermore, Proposition 7 shows that when the initial conditions are set according to the BMFM, we obtain a consistent distribution for the mean-field model of the closed market, in which

the evolution of an advertiser during her campaign coincides with that given by the BMFM.

Next, we should compare an agent's evolution in the closed system to the evolution of the same agent in the mean-field model. That is, suppose that we "attach" a new agent to the real system with its own initial condition and its own strategy, referred as the zeroth agent, independently of everything else. When the number of agents is large, one would expect that presence of this extra agent and the arbitrary strategy that she implements would not affect considerably the evolution of the system. Corollary 1 shows that the law of the state of the zeroth agent in the closed system is close to the law of her state in the closed mean-field model, in a total variation sense. This result uses a propagation of chaos argument to show that the interaction effects in the real system become negligible as the scale increases. We conclude by noting that the law of the zeroth advertiser in the closed mean-field model is equal to the law of an advertiser in the BMFM.

Next, we show that the bound on the total variation of the laws $g'(\eta, F_\kappa, \alpha, T)$, as defined in Corollary 1, converges to zero as κ goes to infinity. Let $Y^\kappa = \bar{\alpha}^\kappa K^\kappa$, and $T = \bar{s}$, where $\bar{\alpha}^\kappa = \max_\theta \alpha_\theta^\kappa$ and $\bar{s} = \max_\theta s_\theta$. Then the bound can be written as $\mathbb{E}_{Y^\kappa} [g^\kappa(Y^\kappa)]$,

$$g^\kappa(y) = \left(A^\kappa(y) + \frac{(\bar{\alpha}^\kappa)^2 \eta^\kappa \bar{s}}{2} C^\kappa(y) \right) \frac{e^{2\bar{\alpha}^\kappa \eta^\kappa \bar{s} B(y)} - 1}{2B(y)},$$

with $A^\kappa(y) = 2\bar{\alpha}^\kappa + \sqrt{\text{Var}Y^\kappa} + |y - \mathbb{E}[Y^\kappa]|$, $B(y) = y$, and $C^\kappa(y) = (y)(2 + y - \alpha^\kappa)$. Using Cauchy-Schwartz inequality, together with Minkowski's inequality, and denoting by $\|X\|_2 = \sqrt{\mathbb{E}[X^2]}$ the L_2 norm we obtain that

$$\mathbb{E}_{Y^\kappa} [g^\kappa(Y^\kappa)] \leq \left(\|A^\kappa(Y^\kappa)\|_2 + \frac{(\bar{\alpha}^\kappa)^2 \eta^\kappa \bar{s}}{2} \|C^\kappa(Y^\kappa)\|_2 \right) \left\| \frac{e^{2\bar{\alpha}^\kappa \eta^\kappa \bar{s} Y^\kappa} - 1}{2Y^\kappa} \right\|_2.$$

The first term in parenthesis can be bounded as

$$\begin{aligned} \|A^\kappa(Y^\kappa)\|_2 &\leq 2\bar{\alpha}^\kappa + \sqrt{\text{Var}Y^\kappa} + \|Y^\kappa - \mathbb{E}[Y^\kappa]\|_2 \\ &= 2\bar{\alpha}^\kappa + 2\sqrt{\text{Var}Y^\kappa} \leq 2\bar{\alpha}^\kappa + 2\sqrt{\bar{\alpha}^\kappa} \sqrt{2\bar{\alpha}^\kappa \lambda^\kappa \bar{s}} = O(\kappa^{-1/2}), \end{aligned}$$

where the first inequality follows from Minkowski's inequality, the equality from the variance formula, the second inequality from the fact that the variance of K^κ is at most $2\bar{\alpha}^\kappa \lambda^\kappa \bar{s}$, and the last from the fact that the number of matching bidders is invariant to the scaling, i.e., $\alpha^\kappa \lambda^\kappa \bar{s} = O(1)$. For the second term in parenthesis we obtain

$$\frac{(\bar{\alpha}^\kappa)^2 \eta^\kappa \bar{s}}{2} \|C^\kappa(Y^\kappa)\|_2 \leq \frac{(\bar{\alpha}^\kappa)^2 \eta^\kappa \bar{s}}{2} (\|(Y^\kappa)^2\|_2 + 2\|Y^\kappa\|_2) = O(\kappa^{-1} \log \kappa),$$

since $\|Y^\kappa\|_2$ and $\|(Y^\kappa)^2\|_2$ are $O(1)$. For the last factor we use the fact that $(e^{\xi y} - 1)/y \leq \xi e^{\xi y}$ for

$y, \xi \geq 0$ to obtain

$$\begin{aligned}
\left\| \frac{e^{2\bar{\alpha}^\kappa \eta^\kappa \bar{s} Y^\kappa} - 1}{2Y^\kappa} \right\|_2 &\leq \bar{\alpha}^\kappa \eta^\kappa \bar{s} \|e^{2\bar{\alpha}^\kappa \eta^\kappa \bar{s} Y^\kappa}\|_2 = \bar{\alpha}^\kappa \eta^\kappa \bar{s} \sqrt{\mathbb{E} \left[\exp(2\bar{\alpha}^\kappa \eta^\kappa \bar{s} Y^\kappa)^2 \right]} \\
&= \bar{\alpha}^\kappa \eta^\kappa \bar{s} \sqrt{\mathbb{E} \left[\exp(4(\bar{\alpha}^\kappa)^2 \eta^\kappa \bar{s} K^\kappa) \right]} = \bar{\alpha}^\kappa \eta^\kappa \bar{s} \exp \left(\lambda^\kappa \bar{s} \left(e^{4(\bar{\alpha}^\kappa)^2 \eta^\kappa \bar{s}} - 1 \right) \right) \\
&= \bar{\alpha}^\kappa \eta^\kappa \bar{s} \exp \left(4(\bar{\alpha}^\kappa \lambda^\kappa \bar{s})(\bar{\alpha}^\kappa \eta^\kappa \bar{s}) \frac{e^{4(\bar{\alpha}^\kappa)^2 \eta^\kappa \bar{s}} - 1}{4(\bar{\alpha}^\kappa)^2 \eta^\kappa \bar{s}} \right) = O(\kappa^\epsilon),
\end{aligned}$$

where the second equality follows from $Y^\kappa = \alpha^\kappa K^\kappa$, third equality from the moment generating function of the Poisson random variable; and the last from the fact that $(e^x - 1)/x = O(1)$ around zero and that $\bar{\alpha}^\kappa \eta^\kappa \bar{s} = O(\log \kappa)$. Note that the exponent $\epsilon > 0$ can be made arbitrarily small by choosing a suitable large base in the logarithmic growth of number of opportunities as given by η^κ and b^κ . Choosing the scaling so that $\epsilon < 1/2$ we obtain that the bound converges to zero.

Define the *extended state* of the zeroth advertiser as the budget remaining, campaign remaining and last realization of her value. Let $H_0(t)$ denote the entire history of the extended states for the zeroth agent until time t (note that is a proper subset of the history defined in Section B.5, which includes the histories of the competing agents). The zeroth advertiser's strategy β^κ maps a history $H_0(t)$ to a bid $\beta^\kappa(H_0(t))$. The total payoff-per-auction of the zeroth advertiser for a given sample path is defined by

$$\begin{aligned}
&\frac{1}{\alpha_\theta^\kappa \eta^\kappa s_\theta} \sum_{n=1}^{N(s_\theta)} \mathbf{1}\{\beta^\kappa(H_0(t_n)) > d_{n,0}, b_0(t_n^-) > 0\} (v_{n,0} - d_{n,0}) A_{n,0}, \\
&= \frac{1}{\alpha_\theta^\kappa \eta^\kappa s_\theta} \sum_{n=1}^{N(s_\theta)} \mathbf{1}\{b_0(t_n^-) - b_0(t_n) > 0\} \left(v_{n,0} - (b_0(t_n^-) - b_0(t_n)) \right) A_{n,0},
\end{aligned}$$

where $A_{n,0} = 1$ whenever the zeroth advertiser participates in the n^{th} auction, and the second equation follows from the fact that the zeroth advertiser's payment is $b_0(t_n^-) - b_0(t_n)$. Note that the payoff-per-auction function is measurable and bounded. Measurability follows from the fact that the strategies are non-anticipating and adaptive w.r.t. the history $H_0(t)$. Boundedness follows from the fact that the utility per auction is bounded by \bar{V} , an advertiser can win at most b_θ/r auctions, and thus that the ratio of total utility to number of auctions is bounded by $g_\theta \bar{V}/r$. Thus, the convergence of the payoff functions follows from the convergence in total variation of the processes' laws given by Corollary 1 for the extended states.

B Mean-field Model for Systems with a Random Number of Agents

In this section we consider a general mean-field model for a system in which the number of agents is random and determined up-front when the system is created. We present our model and results in full generality, since these may be of independent interest. We start by considering a model with homogeneous agents and then we move on to generalize it to heterogeneous agents.

B.1 Real System

Let $K \in \mathbb{Z}_+$ be the number of agents, which is drawn from some discrete distribution $F_k(\cdot)$. After the number of agents in the system is drawn, it remains fixed for the whole time horizon. We denote the state of agent k at time t by $X_k(t) \in \mathbb{X}$ where $\mathbb{X} \subset \mathbb{R}^d$.

The dynamics of the system are as follows. First, the number of agents in the system is drawn. Then, the initial states of the agents $\{X_k(0)\}_{k=1}^K$ are determined as i.i.d. draws from a random variable X_0 . The evolution of the states of the agents is governed by a deterministic drift, and a stochastic jump process that determines the agents' interactions. The deterministic drift depends exclusively on the agent's own state and is oblivious to the other agents' states. That is, the drift is given by a function $v : \mathbb{X} \times \mathbb{R} \rightarrow \mathbb{X}$, which determines the instantaneous change in an agent's state $v(x, t)$ at time t when the current state is x . The drift is assumed to be uniformly bounded, and Borel-measurable in its first argument.

Before defining the interactions we need some notation. Let $\mathbb{X}^{\mathbb{N}}$ be the space of finite length sequences on \mathbb{X} . For a sequence $\vec{x} = \langle x_1, x_2, \dots, x_{|\vec{x}|} \rangle$ we denote by $|\vec{x}|$ the length of the sequence. Given two sequences \vec{x} and \vec{y} we define the concatenation of these sequences as $\vec{x} \cdot \vec{y} = \langle x_1, \dots, x_{|\vec{x}|}, y_1, \dots, y_{|\vec{y}|} \rangle$. The concatenation operator is similarly defined for an element of the space \mathbb{X} and a sequence.

The interactions are governed by the jumps of a Poisson process $N(t)$ with intensity η , where we denote by $\{t_n\}_{n \geq 1}$ the sequence of jump times. Each agent participates in the interaction randomly and independently of other agents with probability α . We denote by $A_{k,n}$ a Bernoulli random variable with success probability α indicating whether the k^{th} agent participates in the n^{th} interaction or not.⁵ The indices of the participating agents is given by the set $\mathcal{M}_n = \{k = 1, \dots, K : A_{k,n} = 1\}$, and the total number of agents in the interaction by $M_n = |\mathcal{M}_n|$. We allow some random noise term $\xi_{k,n} \in \mathbb{E}$ to be associated to each agent participating in the interaction. These noise terms are drawn independently from some common distribution $F_\xi(\cdot)$.

Once the identities of the participating agents is determined the states are updated according to an *interaction function* $f : \mathbb{X}^{\mathbb{N}} \times \mathbb{E}^{\mathbb{N}} \times \mathbb{R} \rightarrow \mathbb{X}^{\mathbb{N}}$, such that $f(\vec{x}, \vec{\xi}, t) = \vec{y}$ gives an additive change in the

⁵As we will later see, in the context of our Ad Exchange, this Bernoulli random variable will be equal to one if both the bidder is "alive" in its campaign and if it matches the targeting criteria.

participating agents' states at time t when their states before the interaction are \vec{x} and the noise terms are $\vec{\xi}$. The interaction function is defined whenever $|\vec{x}| = |\vec{\xi}|$ and satisfies the following properties. First, the length of the input and output state sequences should be consistent, that is, $|\vec{y}| = |\vec{x}|$. Second, the interaction functions is symmetric in its arguments, that is, for every permutation π of the indices $\{1, \dots, |\vec{x}|\}$ we have that $f(\vec{x}_\pi, \vec{\xi}_\pi, t) = \vec{y}_\pi$. Third, the interaction function is uniformly bounded and Borel-measurable.

The dynamics of the agents in the system can be informally defined in terms of the following system of coupled stochastic differential equations (SDE)

$$dX_k(t) = v(X_k(t), t)dt + f_1 \left(X_k(t) \cdot \vec{X}_{-k}(t), \xi_k(t) \cdot \vec{\xi}_{-k}(t), t \right) A_k(t) dN(t), \quad (6)$$

for all agent $k = 1, \dots, K$. In the previous equation we denote by $A_k(t) \triangleq A_{k, N(t)}$ the indicator that agent k participates in the last event before time t , $\xi_k(t) \triangleq \xi_{k, N(t)}$ her noise terms, $\mathcal{M}(t) \triangleq \mathcal{M}_{N(t)}$ the set of indices of agents interacting at the event before time t , the sequence of states of the agents interacting with k by $\vec{X}_{-k}(t) = \langle X_i(t) \rangle_{i \in \mathcal{M}(t) \setminus k}$, and the sequence of noise terms associated to the agents interacting with k by $\vec{\xi}_{-k}(t) = \langle \xi_i(t) \rangle_{i \in \mathcal{M}(t) \setminus k}$. All terms in the right-hand side of the SDE are evaluated at time t^- to preserve predictability. The symmetry of the interaction function f allows one to write the system dynamics as if the agent in consideration was the first argument.

B.2 Mean-field Model

Next, we study the evolution of a fixed agent in a mean-field model associated with the previous system. In the mean-field model the agent in consideration (i) interacts with a random number of agents that is independent of the total number of agents in the system, and (ii) the states of the interacting agents are independent draws from a time-dependent distribution.

We refer to the agent in consideration as the zeroth agent. Let $\tilde{X}_0(t)$ be the state of the agent in consideration in the mean-field model. As in the real system, the initial state of the agent is drawn from the random variable X_0 . At the n^{th} interaction the number of agents in the system \tilde{K}_n is drawn independently from the distribution $F_k(\cdot)$, and the number of participating agents (excluding 0) is given by \tilde{M}_n , which is a Binomial random variable with success probability α and \tilde{K}_n trials. Note that, in the mean-field model, the number of agents in the system is re-drawn at each interaction. In order to determine the evolution of the process $\tilde{X}_0 = \{\tilde{X}_0(t)\}_{t \geq 0}$ one needs to specify the distribution of the interacting agents. In the following, we assume that the states of the agents interacting are drawn from some distribution $\mathbb{P}_c : \mathcal{B}(\mathbb{X}) \times \mathbb{R} \rightarrow [0, 1]$, where $\mathbb{P}_c(\mathcal{X}, t)$ gives the probability that, at time t , the state of an interacting agent lies in the Borel set \mathcal{X} .

The dynamics of the agent in the mean-field model are governed by the following stochastic differ-

ential equation

$$d\tilde{X}_0(t) = v(\tilde{X}_0(t), t)dt + f_1 \left(\tilde{X}_0(t) \cdot \vec{\tilde{X}}_{-0}(t), \xi_0(t) \cdot \vec{\xi}_{-0}(t), t \right) A_0(t) dN(t). \quad (7)$$

In the previous equation the sequence of states of the agents interacting with 0 in the n^{th} event at time t are given by $\vec{\tilde{X}}_{-0,n} = \langle \tilde{X}_{1,n}, \dots, \tilde{X}_{\tilde{M}_n,n} \rangle$, with $\tilde{X}_{k,n}$ drawn i.i.d. from the distribution $\mathbb{P}_c(\cdot; t)$. Similarly, the sequence of noise terms associated to the agents interacting with 0 are given by $\vec{\xi}_{-0,n} = \langle \xi_{1,n}, \dots, \xi_{\tilde{M}_n,n} \rangle$. These noise terms are drawn from the same distribution as in the real system, with the exception that now the noise vector has \tilde{M}_n components.

We emphasize that in order to determine the evolution of the process \tilde{X}_0 one needs to specify the distribution of the interacting agents' states \mathbb{P}_c . For the system to be consistent this distribution should be endogenously determined from the model itself. That is, suppose that one postulates a candidate distribution \mathbb{P}_c , and let the mean-field system evolve with interacting agents' states drawn from that distribution. It should be the case that the state at time t of the zeroth advertiser in the mean-field model is distributed as $\mathbb{P}_c(\cdot, t)$. We formalize this concept next.

Definition 2. *A distribution $\mathbb{P}_c : \mathcal{B}(\mathbb{X}) \times \mathbb{R} \rightarrow [0, 1]$ is T-consistent if for any Borel-measurable set of states \mathcal{X} and time t*

$$\mathbb{P}_c(\mathcal{X}, t) = \mathbb{P} \left\{ \tilde{X}_0(t) \in \mathcal{X} \mid \text{interacting agents states drawn from } \mathbb{P}_c \right\}.$$

Note that at both sides of the previous fixed point equation the distribution \mathbb{P}_c is time-dependent. Uniqueness of a T-consistent distribution for problems with a bounded jump rate can be proved using a contraction argument on probability measures (see, e.g., Graham (1992)).

B.3 Boltzmann Tree

Given a distribution for the initial conditions, a T-consistent distribution for the mean-field model can be constructed by considering the associated *Boltzmann tree*, which we detail next. We shall construct the tree in two steps. In the first step we move backwards from time T until time 0 to determine all the interactions in the horizon. Here we are not concerned about the states; instead, we focus on determining the time in which interactions occur and the agents involved in these interactions. In the second step, we move forwards in time to determine the evolution of the states. We start by specifying the initial conditions and the noise terms for the interactions, and then the state processes are computed deterministically using the system dynamics.

We partition the lifetime of an agent in the system during time $[0, T]$ into a countable sequence of *slices*, where one slice is the lifetime of the agent between two consecutive interactions. We shall label

each slice by a finite sequence of indices $\vec{k} \in \mathbb{Z}_+^N$, denoted as $\vec{k} = \langle k_1, k_2, \dots \rangle$.

Step 1. We construct a tree rooted in the zeroth agent by moving backwards in time. Let $\langle 0 \rangle$ be the first slice for the zeroth agent from the time of the last interaction until time T . We associate to each slice \vec{k} a Poisson process $N_{\vec{k}}(t)$. The slice ends at the time $T_{\vec{k}}^+$ and begins at a time $T_{\vec{k}}^-$ that corresponds to the last jump of the Poisson process $N_{\vec{k}}(t)$ before the time $T_{\vec{k}}^+$. At time $T_{\vec{k}}^-$ the agent interacts with an independent random number of agents, denoted by $\tilde{M}_{\vec{k}}$. To determine the number of agents competing we first draw the number of agents in system $\tilde{K}_{\vec{k}}$ from $F_k(\cdot)$, and then draw $\tilde{M}_{\vec{k}}$ as a Binomial random variable with success probability α and $\tilde{K}_{\vec{k}}$ trials (all quantities are drawn independently). This event corresponds to the *branching* of the tree. At this point $\tilde{M}_{\vec{k}} + 1$ new slices are constructed and attached to the tree. The new slices are labeled $i \cdot \vec{k}$ with i from 0 to $\tilde{M}_{\vec{k}}$. Here, the first slice $0 \cdot \vec{k}$ corresponds to the previous slice of the incumbent agent \vec{k} , which is referred as the *creator* of the interaction. The remaining slices correspond to the other agents interacting with \vec{k} at that point. These slices $i \cdot \vec{k}$ end at time $T_{i \cdot \vec{k}}^+ = T_{\vec{k}}^-$. These steps are repeated recursively until all slices reach time 0.

Each time a slice \vec{k} reaches time 0, we associate to it an agent whose lifetime would extend until she participates in an interaction in which she is not the creator. From that point on, we are not concerned about the state of the agent since it is not relevant to determine the evolution of the zeroth agent. That is, if $\vec{k} = \langle 0, \dots, 0, k_{n+1}, k_{n+2}, \dots, k_{|\vec{k}|} \rangle$ with $k_{n+1} \neq 0$, then the agent participated in n different events in which she was the creator, and was created by interacting with slice $\langle k_{n+2}, \dots, k_{|\vec{k}|} \rangle$ in the $n + 1^{\text{th}}$ event. Let \mathcal{K} be the set of slices that reach time 0.

Step 2. Once the tree is constructed, we assign a state process $\tilde{X}_{\vec{k}} \triangleq \{\tilde{X}_{\vec{k}}(t)\}_{t \in [T_{\vec{k}}^-, T_{\vec{k}}^+]}$ to each slice. The evolution of the state are determined by following the SDE forward in time. First, for the each slice $\vec{k} \in \mathcal{K}$, we set $\tilde{X}_{\vec{k}}(0)$ according to i.i.d. draws from the initial distribution X_0 . Then, the states evolve deterministically according to the drift v during the slice $[T_{\vec{k}}^-, T_{\vec{k}}^+]$. At the point of an interaction, noise terms $\xi_{\vec{k}}$ are drawn independently for each slice, and the state of the creator after the interaction is determined using the interaction function f . That is, if agents $i \cdot \vec{k}$ participate in the interaction, we have that

$$\tilde{X}_{\vec{k}}(T_{\vec{k}}^-) = f_1 \left(\left\langle \tilde{X}_{i \cdot \vec{k}}(T_{i \cdot \vec{k}}^+) \right\rangle_{i=0}^{\tilde{M}_{\vec{k}}}, \left\langle \xi_{i \cdot \vec{k}} \right\rangle_{i=0}^{\tilde{M}_{\vec{k}}}, T_{\vec{k}}^- \right).$$

We proceed in this manner until slice $\langle 0 \rangle$ is reached, which corresponds to the last slice of the zeroth agent.

Once we conclude with the forward evolution, the state process for the whole lifetime of the zeroth agent can be reconstructed by concatenating the slices $\langle 0, 0, \dots, 0 \rangle$, $\langle 0, \dots, 0 \rangle$, and so forth until slice $\langle 0 \rangle$. The state process is just the concatenation of the state processes of each slice $\{\tilde{X}_0(t)\}_{t \in [0, T]} =$

$\bigcup_{\vec{k}:k_i=0} X_{\vec{k}}$.

In the Boltzmann tree agents evolve without self-interactions since each agent interacts with agents whom themselves evolve independently within trees. Thus, each agent in the tree evolves as in the mean-field model, and the law of the process $\{\tilde{X}_0(t)\}_{t \in [0, T]}$ constructed above using the Boltzmann tree is T-consistent for the mean-field model (Definition 2). This is proved formally in, for example, Chauvin and Giroux (1990).

B.4 Propagation of Chaos

Let $\mathcal{L}(X_k | K \geq k)$ be the law for the process of the k^{th} agent's state in the real system conditioning on the number of agents being greater or equal than k (that is, under the condition that the k^{th} agent is in the system). The next result shows that the law of that agent is close, in total variation norm, to the law of an agent in the mean-field model.

Proposition 6. *Let \mathbb{P}_c be a T-consistent distribution for the mean-field model. Then*

$$\|\mathcal{L}(X_k | K \geq k) - \mathbb{P}_c\|_{[0, T]} \leq g(\eta, F_k, \alpha, T, k),$$

where $\|\cdot\|_{[0, T]}$ denotes the total variation norm over the time horizon $[0, T]$, and

$$g(\eta, F_k, \alpha, T, k) = \mathbb{E}_K \left[\left(A + \frac{\alpha^2 \eta T}{2} C \right) \frac{e^{2\alpha \eta T B} - 1}{2B} \Big| K \geq k \right],$$

with $A = \alpha + \alpha \sqrt{\text{Var}K} + \alpha |K - \mathbb{E}[K]|$, $B = \alpha(K - 1)$, and $C = (\alpha K - \alpha)(2 + \alpha K - 2\alpha)/2$.

Proof: The proof follows from the combination of a propagation of chaos argument for the interactions (such as that used in Graham and Méléard (1994) and Iyer et al. (2011)) and a fluid limit for the number in system. The result is proven in four steps. In the first step, we present a path-wise construction of the real system with the minimal information necessary to determine all interactions that occur in the system. In the analysis we are not concerned about the evolutions of the states, and thus we shall not describe the draws of the noise terms and outcomes of the interactions. Indeed, we shall restrict our attention to the interaction times and the identity of agents interacting in each event. In the second step, we present some sufficient conditions under which the evolution of the k -th agent in the real system is “close” to that of the same agent in a Boltzmann tree. Namely, that (i) an agent interacts with distinct agents who do not share any past common influence, and that the same applies recursively to those agents she interacts with; and (ii) the stochastic deviations in number of agents initially in the system do not affect significantly the number of agents interacting in the successive events. In the third step, we show that the complement of condition (i) occurs with low probability; in the last step we show that that the same holds for condition (ii).

Step 1: Path-wise construction of the real system. Here we present a path-wise construction of the real system with the minimal information necessary to determine all interactions that occur in the system during time $[0, T]$. The initial conditions are as follows. First, the initial number of agents is drawn from $K|K \geq k$. Second, the state of each agents is drawn independently from the initial distribution X_0 .

Independently, we have that events occur according to the jumps of the Poisson process $\{N(t)\}_{t \geq 0}$ with rate η . Recall that the jump times were denoted by t_n . At the time of the n^{th} event we need to determine which agents participate in the event. We do so by assigning independent Bernoulli coins to the agents currently in the system so that agents interact whenever the coin is one. Let $\mathbf{M} = \{m_{n,i}\}_{i \in \mathbb{N}, n \in \mathbb{N}}$ be an infinite matrix of independent Bernoulli random variables with success probability α , which act as the indicators of whether the agent interact in each event. More formally, in our construction the k^{th} agent participates in the n^{th} event if her associated coin $m_{n,k}$ is one. We refer to these at the *interacting coins*. Let $\mathcal{M}(t_n)$ and $M(t_n)$ denote the indices and number of agents interacting in the events. This information suffices to determine all the interactions of the real system.

Step 2: Interaction Graphs and Coupling. The evolution of the state of an agent is directly affected by other agents that interact with her in the events, and indirectly influenced by other agents who recursively affect the agents she interacts with. Graham and Méléard (1994) introduced the *interaction graph* construction to summarize the past history of an agent, including all the agents that have influenced her evolution of the state. If we prove that all agents that may have influenced the one in consideration share no common influence we will be able to prove that the real system evolved as the mean-field model.

We define the interaction graphs as follows. Let $\Gamma_k(t) \in \mathcal{P}(\mathbb{R}_+ \times \mathcal{P}(\mathbb{N}_0))$ be the interaction graph of agent k at time t , where $\mathcal{P}(X)$ denotes the power set of X . The interaction graph is a set of pairs (t', \mathcal{M}') indicating that at time t' agents with indices in the set \mathcal{M}' interacted in an event, and we shall see that it records all events that may affect directly or indirectly the state of this agent. The interaction graphs are built recursively as follows. First, at time zero the interaction graph is $\Gamma_k(0) = \{(0, \{k\})\}$. Afterwards, the interaction graph of the agent in consideration remains unchanged until she participates in an event. If her interacting coin for the event at time t is one, the interaction graph is extended to include the interaction graphs of all agents that participate in the event, that is,

$$\Gamma_k(t) = \bigcup_{k' \in \mathcal{M}(t)} \Gamma_{k'}(t^-) \cup \{(t, \mathcal{M}(t))\}$$

where $\Gamma_k(t^-)$ denotes the interaction set just before the event. As a consequence, after an event the histories of all participants are appended in the graph; the current state of the agent may have been influenced by them. Note that the interaction graphs are deterministically determined once we fix the

path-wise construction of the system.

If at time t we have that interaction graphs of two agents k and k' are disjoint, that is $\Gamma_k(t) \cap \Gamma_{k'}(t) = \emptyset$, then there is no common agent that had influenced them in the past, and these agents have evolved independently. We say that the interaction graph $\Gamma_k(t)$ is a *tree* if for all $(t', \mathcal{M}') \in \Gamma_k(t)$ we have that $\Gamma_{k'}(t'^-) \cap \Gamma_{k''}(t'^-) = \emptyset$ for all pairs of agents $k' \neq k'' \in \mathcal{M}'$. This implies that the agent k evolved until time t without self interactions, that is, the agent interacts throughout her campaign with distinct agents who do not share any past common influence, and that the same applies recursively to those agents she competes with.

The fact that for the k^{th} agent its graph $\Gamma_k(T)$ is a tree guarantees that the interaction effect is not present in the evolution of the process. However, it can still be the case that the branching of the tree is correlated inter-temporally due to the fact that the number of agents in the real system is fixed while in the mean-field model this quantities are independently drawn at each event. For example, if the initial number of agents is large, one would expect that the tree would have more branches. For the correlation effect to be absent one needs that the number of interacting agents in the successive events in the graphs are uncorrelated.

From the perspective of one agent, in the real system the number of agents in system is $K - 1$, and the number of interacting agents is Binomial with success probability α and $K - 1$ trials. Now, let $\{\tilde{K}_n\}_{n \in \mathbb{N}}$ be a sequence of independent random variables drawn from $F_k(\cdot)$. We compare, using a coupling argument, the real system with an alternate system in which the number of interacting agents in each event is determined by the independent sequence \tilde{K}_n instead of the fixed amount $K - 1$, but keeping the same interacting coins.

Let $M_{n,k}$ be the number of interacting agents competing in the n^{th} event against the k^{th} agent (excluding the agent in consideration). Assuming that the k^{th} agent participates in the event, this quantity can be written as $M_{n,k} = \sum_{i=1}^K m_{n,i} - 1$. For the real system to evolve as in the mean-field, one needs that: (i) the number of agents competing is independent across events, and (ii) the number of competing agents is Binomial with success probability α and a random number of trials drawn from $F_k(\cdot)$. Keeping the same interacting coins, for the latter conditions to hold one needs that the number of interacting agents in each event coincides with $\tilde{M}_{n,k} = \sum_{i=1}^{\tilde{K}_n+1} m_{n,i} - 1$, whenever the k^{th} agent participates in the event. The extra term in the summation guarantees that the number of interacting agents coincides with that of the mean-field model, which excludes the agent in consideration. Indeed, the random variable $\tilde{M}_{n,k}$ is distributed as a Binomial with success probability α and a random number of trials drawn from $F_k(\cdot)$, which coincides with number of competing interacting agents of the mean-field model.

Let $\Delta(\Gamma_k(t)) = \left| M_{\mathcal{N}_k(t),k} - \tilde{M}_{\mathcal{N}_k(t),k} \right|$ be the maximum difference between the actual and mean-field model number of interacting agents competing with the k^{th} agent in the last event before time t ; where

we denoted by $\mathcal{N}_k(t) = \sup\{n \leq N(t) : m_{n,k} = 1\}$ the index of the last event before time t in which the k^{th} agent participated. Note that if $\Delta(\Gamma_k(t)) = 0$, the number of interacting agents in the last event is identical to that of the mean-field model. We say that an interaction graph $\Gamma_k(t)$ is *uncorrelated* if $\Delta(\Gamma_k(t)) = 0$ and for all pairs $(t', \mathcal{M}') \in \Gamma_k(t)$ we have that $\Delta(\Gamma_{k'}(t'^-)) = 0$ for all $k' \in \mathcal{M}'$. The latter condition guarantees that all events that may have influenced the state at time t of agent k have an independent number of interacting agents which coincides with that of the mean-field model.

Now, we are ready to state the conditions under which the evolution of the real system coincides with that of the Boltzmann tree, and therefore, with \mathbb{P}_c by the argument at the end of Section B.3. Recall that in the Boltzmann tree agents evolve without self-interactions since each agent interacts with agents whom themselves evolve independently within trees. Therefore, the evolution of the state of an agent k until time t in the real system coincides with that of the Boltzmann tree in the event that her interaction graph $\Gamma_k(t)$ is a tree and uncorrelated. In particular, using a coupling argument, we can show that the difference of the laws of both processes is bounded in total variation by

$$\begin{aligned} \|\mathcal{L}(X_k | K \geq k) - \mathbb{P}_c\|_{TV, [0, T]} &\leq 1 - \mathbb{P}\{\Gamma_k(T) \text{ is a tree and uncorrelated}\} \\ &\leq \mathbb{P}\{\Gamma_k(T) \text{ not a tree}\} + \mathbb{P}\{\Gamma_k(T) \text{ not uncorrelated}\}, \end{aligned}$$

where the second inequality follows from a union bound. In the remainder of the proof, we bound each term on the right-hand side.

Step 3: Correlation effect. In this step we shall bound the probability that an interaction graph is not correlated by conditioning on the number of agents, and then taking expectations with respect to the number of agents in the system.

Let $U(t; K)$ be a bound on the probability that the interaction graph $\Gamma_k(t)$ of an agent picked at random at time t is not uncorrelated, given that the number of agents in the system is K . We can obtain such bound by conditioning on the time of the last interacting event before t , and exploiting the recursive nature of the interactions graphs. In the process we shall obtain a functional inequality of the renewal kind. Indeed, by conditioning on the time $x \leq t$ of the last event before time t we obtain

$$\begin{aligned} U(t; K) &\leq \int_0^t \underbrace{\mathbb{P}\{\Delta(\Gamma_k(x)) \neq 0 \mid K, k \in \mathcal{M}(x)\}}_{(I)} \\ &\quad + \underbrace{\mathbb{P}\left\{\bigcup_{k' \in \mathcal{M}(x)} \Gamma_{k'}(x^-) \text{ not uncorrelated} \mid K, k \in \mathcal{M}(x)\right\}}_{(II)} d\bar{F}_{\alpha\eta}(t-x) \end{aligned} \quad (8)$$

where $F_{\alpha\eta}(\cdot)$ is the cumulative distribution function of the events inter-arrival time, which is exponential

with rate $\alpha\eta$. The first term of the integrand can be bounded as follows

$$\begin{aligned}
(I) &= \left(\mathbb{P} \left\{ \mathcal{M}_{N_k(x),k} \neq \widetilde{M}_{N_k(x),k} \mid K, k \in \mathcal{M}(x) \right\} \right) \\
&\leq \alpha \mathbb{E} \left[|\tilde{K}_n + 1 - K| \mid K \right] \leq \alpha + \alpha \mathbb{E} \left| \tilde{K}_n - \mathbb{E}[\tilde{K}_n] \right| + \alpha |K - \mathbb{E}[K]| \\
&\leq \alpha + \alpha \sqrt{\text{Var}K} + \alpha |K - \mathbb{E}[K]| = A,
\end{aligned}$$

where the second inequality follows from observing that the interacting number of agents differ if at least one the interacting coins in $(K, \tilde{K}_n + 1]$ or $(\tilde{K}_n + 1, K]$ is one, the second follows from the triangle inequality, the third from Lyapunov's inequality and the variance formula. Thus, we obtain that the probability that the number of interacting agents in the real system differs from that of the mean-field model is bounded uniformly over time by A that is a function of the random variable K .

For the second term on the rhs of (8), use that the expected number of agents in the interaction is $\alpha(K - 1)$, and a union bound to estimate the probability that each of the sub-interaction graphs are not correlated to obtain

$$(II) \leq (1 + \alpha(K - 1))U(x; K) = (1 + B)U(x; K)$$

where the second inequality follows the triangle inequality. In the latter, B , also a function of the random variable K , is a bound uniform over time on the expected number of competing interacting agents.

Using the two previous bounds in conjunction with equation (8), one obtains the functional equation $U(t; K) \leq AF_{\alpha\eta}(t) + (1 + B)(U(\cdot; K) * F_{\alpha\eta})(t)$; where we denoted the Stieltjes convolution by $(F * G)(x) = \int F(x-u) dG(u)$. Iterating the functional equation we obtain the following exponential bound on the probability that the tree is not uncorrelated

$$\begin{aligned}
U(t; K) &\leq A \sum_{i=0}^{\infty} (1 + B)^i F_{\alpha\eta}^{(i+1)}(t) = A \sum_{i=0}^{\infty} (1 + B)^i \int_0^t \frac{(\alpha\eta)^{i+1} x^i}{i!} e^{-\alpha\eta x} dx \\
&= A \int_0^t \alpha\eta e^{-\alpha\eta x} \sum_{i=0}^{\infty} \frac{(\alpha\eta x (1 + B))^i}{i!} dx = A \int_0^t \alpha\eta e^{\alpha\eta x B} dx \\
&= A \frac{e^{\alpha\eta t B} - 1}{B} \leq A \frac{e^{2\alpha\eta t B} - 1}{2B}
\end{aligned} \tag{9}$$

where $F_{\alpha\eta}^{(i)}(t) = \int_0^t \frac{(\alpha\eta)^i x^{i-1}}{(i-1)!} e^{-\alpha\eta x} dx$ denotes the i^{th} convolution of the distribution of inter-arrival times, which is Erlang with shape i and rate $\alpha\eta$; the second equation follows from non-negativity and Tonelli's Theorem; the third from the power series definition of the exponential function; and the last inequality from the fact that $(e^x - 1)/x \leq (e^{2x} - 1)/(2x)$ for $x \geq 0$.

Step 4: Interaction effect. We bound the probability that an interaction graph is not a tree by following closely the developments in Graham and Méléard (1994). Let $Q(t; K)$ be a bound on the probability that the interaction graphs of two distinct agents i and j drawn at random from the system at time t are not disjoint when K bidders are in the system, that is, $\Gamma_i(t) \cap \Gamma_j(t) \neq \emptyset$. When these interactions graphs are not disjoint there is at least one interaction (t', \mathcal{M}') that belongs to both graphs. It may be the case that neither of these agents participate in that interaction, but instead some other agents participated who later influenced indirectly the agents in considerations.

Given an interaction graph Γ and two agents i and j we define the interaction distance of these two agents, denoted by $\text{dist}(i, j; \Gamma)$, as the minimum number of agents in the chain of influence between i and j . The distance is zero whenever there is some event in which both i and j directly interacted, that is, $\text{dist}(i, j; \Gamma) = 0$ if there is some event $(t, \mathcal{M}) \in \Gamma$ such that $i, j \in \mathcal{M}$. If there is no direct interaction, it is defined recursively as one plus the minimum distance between j and all k that interacted with i in some event. That is, $\text{dist}(i, j; \Gamma) = 1 + \min \{ \text{dist}(k, j; \Gamma) : i, k \in \mathcal{M} \text{ and } (t, \mathcal{M}) \in \Gamma \}$. The distance is ∞ if there is no chain of influence between i and j in the graph.

We have that two interaction graphs are not disjoint, $\Gamma_i(t) \cap \Gamma_j(t) \neq \emptyset$, whenever there is some chain of influence between agents i and j in the union of their interaction graphs, that is, $\text{dist}(i, j; \Gamma_i(t) \cup \Gamma_j(t)) < \infty$. We provide the estimate on the probability that the interaction graphs of two agents are not disjoint by considering the probability that two agents drawn at random from the system at time t are at interaction distance of d , which we denote by $Q_d(t; K)$. Then, the total bound can be obtained as $Q(t; K) = \sum_{d \geq 0} Q_d(t; K)$.

When the distance is zero there is a direct interaction between i and j between time 0 and T , an event occurring with rate $\alpha^2 \eta$. Thus,

$$Q_0(t; K) = 1 - e^{-\alpha^2 \eta t} \leq \alpha^2 \eta t.$$

Next, we proceed by induction on d . Suppose that we have a bound for $Q_{d-1}(t; K)$ for all time $t \in [0, T]$. For a chain reaction of distance d to happen between some i and j , we first need that either of them interacts with a third agent k such that the distance from k and the interacting agent is $d - 1$. The first interaction occurs at rate $2\alpha\eta$ (the minimum of two exponentials with rate $\alpha\eta$), and the expected number of agents she interacts is bounded by $B = \alpha(K - 1)$. Thus, we obtain that

$$\begin{aligned} Q_d(t; K) &\leq \int_0^t B Q_{d-1}(x; K) d\bar{F}_{2\alpha\eta}(t-x) \\ &\leq B \left(Q_{d-1}(\cdot; K) * F_{2\alpha\eta} \right) (t) \leq B^d \left(Q_0(\cdot; K) * F_{2\alpha\eta}^{(d)} \right) (t), \end{aligned}$$

where $F_{2\alpha\eta}(\cdot)$ is the cumulative distribution function of an exponential with rate $2\alpha\eta$, and the third

inequality follows from iterating the function equation as done previously. Summing over all non-negative distances d we obtain the following estimate on the probability that two interaction graphs are connected

$$\begin{aligned}
Q(t; K) &= \sum_{d=0}^{\infty} Q_d(t; K) \leq \sum_{d=0}^{\infty} B^d \left(Q_0(\cdot; K) * F_{2\alpha\eta}^{(d)} \right) (t) \\
&\leq Q_0(t; K) + B \sum_{d=0}^{\infty} (1+B)^d \left(Q_0(\cdot; K) * F_{2\alpha\eta}^{(d+1)} \right) (t) \\
&= Q_0(t; K) + B \int_0^t Q_0(t-x; K) 2\alpha\eta e^{2\alpha\eta x B} dx = \frac{\alpha}{2B} (e^{2\alpha\eta t B} - 1), \tag{10}
\end{aligned}$$

where the second inequality follows from partitioning the sum, the non-negativity of the terms and using that $B \leq B + 1$; the second equation from the third equation from Tonelli's Theorem and the power series definition of the exponential function; and the last equality from integrating.

Now, let $L(t; K)$ be a bound on the probability that the interaction graph of an agent drawn at random at time t is not a tree. For this to hold we need that the agent interacts with other agents, whose interaction graphs are themselves trees, and that these interactions graphs are disjoint. The expected number of agents she interacts with is $B = \alpha(K - 1)$, and the expected number of pairs of agents is $C = (\alpha K - \alpha)(2 + \alpha K - 2\alpha)/2$, both functions of the random variable K . Thus,

$$\begin{aligned}
L(t; K) &\leq \int_0^t (CQ(x; K) + (1+B)L(x; K)) d\bar{F}_{\alpha\eta}(t-x) \\
&= C \left(Q(\cdot; K) * F_{\alpha\eta} \right) (t) + (1+B) \left(L(\cdot; K) * F_{\alpha\eta} \right) (t) \\
&\leq C \sum_{i=0}^{\infty} (1+B)^i \left(Q(\cdot; K) * F_{\alpha\eta}^{(i+1)} \right) (t) \\
&= C \int_0^t Q(t-x; K) \alpha\eta e^{\alpha\eta x B} dx \\
&\leq \frac{\alpha C}{2} \frac{e^{2\alpha\eta t B} + 1 - 2e^{\alpha\eta t B}}{B^2} \leq \frac{\alpha^2 \eta t C}{2} \frac{e^{2\alpha\eta t B} - 1}{2B}, \tag{11}
\end{aligned}$$

where the second inequality follows from iterating the functional equation; the second equality from Tonelli's Theorem and the power series definition of the exponential function; and the third inequality from the bound (10), integrating and discarding negative terms; and the last inequality from the fact that $(e^{2x} - e^x - 1)/x^2 \leq (e^{2x} - 1)/(2x)$ for $x \geq 0$.

Step 5: Putting it all together. We conclude by taking expectations with respect to the initial number of agents in the system K to obtain bounds $U(T) = \mathbb{E}_K[U(T; K) | K \geq k]$, and $L(T) =$

$\mathbb{E}_K[L(T; K)|K \geq k]$. Thus, we have that

$$U(T) + L(T) \leq \mathbb{E}_K \left[\left(A + \frac{\alpha^2 \eta T}{2} C \right) \frac{e^{2\alpha \eta T B} - 1}{2B} \middle| K \geq k \right].$$

◇

B.5 Evaluating Deviations

The previous result allows one to compare an agent’s evolution in the real-system to the evolution in the mean-field model. Now, consider a real system in which we “attach” a new agent with its own initial condition, referred as the zeroth agent, independently of everything else. Let $X'_0(t) \in \mathbb{X}'$ denote the state of the zeroth bidder at time t in the new real system, where the state space $\mathbb{X}' \subset \mathbb{R}^d$ may be different to that of the other agents. Whenever this agent interacts the dynamics are governed by a new interaction function f' which is not symmetric w.r.t. the zeroth agent. Moreover, this function is allowed to depend on the entire history of the agent’s states and noise terms for all interactions until time t , which we denote by $H'_0(t) = \left\{ X'_0(t_n^-), \vec{X}'_{-0}(t_n^-), \xi_0(t_n), \vec{\xi}_{-0}(t_n) \right\}_{t_n \leq t}$. In the case that the zeroth agent does not participate in the interaction, the dynamics remain unchanged.

When the number of agents is large, one would expect that the arbitrary interaction function f' and the presence of an extra agent would not affect considerably the evolution of the system. As such, in order to study the performance of the zeroth agent in this new system, one can consider an alternative mean-field model for the zeroth agent in which interactions are governed by f' and the states of the interacting agents drawn from the T-consistent distribution \mathbb{P}_c of the original system’s mean-field model. Let $\tilde{X}'_0(t)$ be the state of the zeroth agent in the alternate mean-field model. This would satisfy the SDE

$$d\tilde{X}'_0(t) = v(\tilde{X}'_0(t), t)dt + f'_1 \left(\tilde{H}'_0(t) \right) A_0(t)dN(t), \quad (12)$$

with the interacting agents’ states drawn from \mathbb{P}_c , and $\tilde{H}'_0(t)$ the history of the zeroth agent in the mean-field model as defined before. Let \mathbb{P}'_c denote the law of the zeroth agent in the alternative mean-field model. That is, for any Borel-measurable set of states \mathcal{X} and time t :

$$\mathbb{P}'_c(\mathcal{X}, t) = \mathbb{P} \left\{ \tilde{X}'_0(t) \in \mathcal{X} \mid \text{interacting agents states drawn from } \mathbb{P}_c \text{ and } f' \text{ is used} \right\}.$$

Using a similar argument that in the previous result we can show that the law of the zeroth agent in the alternative system is close to the law of in the mean-field model, in a total variation sense.

Corollary 1. *We have that*

$$\|\mathcal{L}(X'_0) - \mathbb{P}'_c\|_{[0,T]} \leq g'(\eta, F_k, \alpha, T)$$

where

$$g'(\eta, F_k, \alpha, T) = \mathbb{E}_K \left[\left(A + \frac{\alpha^2 \eta T}{2} C \right) \frac{e^{2\alpha \eta T B} - 1}{2B} \right],$$

with $A = 2\alpha + \alpha\sqrt{\text{Var}K} + \alpha|K - \mathbb{E}[K]|$, $B = \alpha K$, and $C = (\alpha K)(2 + \alpha K - \alpha)$.

Proof: The proof follows as in Proposition 6, but considering instead a Boltzmann tree in which states are updated using the interaction function f' whenever the zeroth agent is involved. The evolution of the state of the zeroth agent in the Boltzmann tree and in the real system coincide whenever her interaction graphs is a tree and uncorrelated. The probability of these events, given the number of agents in the system, may be bounded as in the proof of Proposition 6. We conclude by taking expectation w.r.t. the number of agents in the system, which is now equal to $K + 1$. \diamond

B.6 Heterogeneous interaction probabilities

Our model can be extended to accommodate heterogeneous interaction probabilities which are dependent on the agent's state and time. Consider an interaction probability function $\alpha : \mathbb{X} \times \mathbb{R} \rightarrow [0, 1]$, such that $\alpha(x, t)$ gives the probability that an agent interacts in an event at time t when her state is x . In the following we assume that this function is uniformly bounded from above by $\bar{\alpha}$.

In this context, the real system is defined as before, with the only exception that the indicator that the k^{th} agent interacts in the n^{th} event as given by $A_{k,n}$ is now Bernoulli with success probability $\alpha(X_k(t_n), t_n)$. In the mean-field model, the number of interacting agents at the n^{th} event at time t_n is set to be

$$\tilde{M}_n = \sum_{k=1}^{\tilde{K}_n} \tilde{A}_{k,n}, \tag{13}$$

where $\tilde{A}_{k,n}$ is Bernoulli with success probability $\tilde{\alpha}(t_n)$, and $\tilde{\alpha}(t) = \int_{\mathbb{X}} \alpha(x, t) d\mathbb{P}_c(x, t)$ is the *expected* probability that an agent interacts at time t . Thus, \tilde{M}_n is Binomial with success probability $\tilde{\alpha}(t_n)$ and a random number of trials \tilde{K}_n . Some states for an agent might be more likely, conditional on her interacting in the event. Indeed, an agent's state conditional on interacting is distributed as

$$\tilde{\mathbb{P}}_c(\mathcal{X}, t) \triangleq \frac{1}{\tilde{\alpha}(t)} \int_{\mathcal{X}} \alpha(x, t) d\mathbb{P}_c(x, t). \tag{14}$$

Note that in this case Definition 2 of a T-consistent distribution can be extended to

$$\mathbb{P}_c(\mathcal{X}, t) = \mathbb{P} \left\{ \tilde{X}_0(t) \in \mathcal{X} \mid \text{interacting agents states drawn from } \tilde{\mathbb{P}}_c \right\},$$

where $\tilde{\mathbb{P}}_c$ is given by equation (14).

We conclude this section by describing how to reduce the heterogeneous interaction probability model to the homogenous one. To perform the reduction we consider an homogenous model in which agents decide to interact in two rounds. In the first round, an agent interacts with a common probability $\bar{\alpha}$ independently of her state and the time. In the second round, an agent with state x at time t interacts with probability $\alpha(x, t)/\bar{\alpha}$. The first round is performed as in the homogenous model, and the second round is performed within a new interaction function \bar{f} . We formalize this next.

For the first round, let the interaction indicators $\bar{A}_{k,n}$ be Bernoulli with success probability $\bar{\alpha}$. For the second round, we extend the noise terms by $\bar{\xi}_{k,n} = (\xi_{k,n}, u_{k,n})$ with $u_{k,n}$ distributed as a Uniform random variable with support $[0, 1]$. Letting \bar{x} the states and \bar{u} the uniform noise terms of the agents that pass the first round, we denote the set of agents that pass the second round by $\bar{\mathcal{M}}(\bar{x}, \bar{u}, t) = \{i = 1, \dots, |\bar{x}| : u_i \leq \alpha(x_i, t)/\bar{\alpha}\}$. The previous construction guarantees that agents interact with their correct state and time-dependent probability. Then, the new interaction function is defined as $\bar{f}_i(\bar{x}, (\bar{\xi}, \bar{u}), t) = 0$ for $i \notin \bar{\mathcal{M}}(\bar{x}, \bar{u}, t)$ and

$$\bar{f}(\bar{x}, (\bar{\xi}, \bar{u}), t) \Big|_{\bar{\mathcal{M}}(\bar{x}, \bar{u}, t)} = f(\bar{x}|_{\bar{\mathcal{M}}(\bar{x}, \bar{u}, t)}, \bar{\xi}|_{\bar{\mathcal{M}}(\bar{x}, \bar{u}, t)}, t),$$

where $\bar{x}|_{\mathcal{I}} = \langle x_i \rangle_{i \in \mathcal{I}}$ is a slice of \bar{x} restricted to the set of indices \mathcal{I} .

The previous results extend to the heterogenous interaction model by considering the homogenous model with interacting probability $\bar{\alpha}$, noise terms $\bar{\xi}$ and interaction function \bar{f} .

C Ad Exchange Market as a Closed System

In this section we model our Ad Exchange market as a system with a random number of agents and heterogeneous interaction probabilities. In this context, advertisers are the agents, auctions correspond to the events, and matchings to interactions. A key characteristic of the exchange is that the market is *open*, that is, advertisers arrive at random points in time, run their campaigns for a fixed amount of time, and then depart. We can model arrival and departures by considering a *closed* market in which advertisers are present for the whole time horizon but are “alive” only during their campaign. In this system, the number of advertisers originally present is random and corresponds to the number of arrivals during the time horizon. We refer to this as the *closed system*.

We consider a time horizon $[0, T]$. The state of advertiser k at time t in the closed system is given by $X_k^C(t) \in \mathbb{R}^3 \times \Theta$; where given a state $x_k^C(t) = (b_k(t), s_k(t), \tau_k, \theta_k)$, we denote by $b_k(t)$ the budget remaining, by $s_k(t)$ the remaining campaign length, τ_k the campaign start time, and by θ_k the advertiser’s type (we adhere to the convention that capital letters denote random variable and lower case letters denote realizations). The last two quantities are time-invariant. Advertisers are alive only during their campaign, that is, they are only allow to match in an auction if $\tau_k \leq t \leq \tau_k + s_{\theta_k}$. The model is as follows:

- The initial number of advertiser K is Poisson with mean $\sum_{\theta} \lambda_{\theta}^T$, where $\lambda_{\theta}^T = \lambda_{\theta}(T + s_{\theta})$ is the number of advertisers originally in the market plus the arrivals during the horizon $[0, T]$.
- The interaction probability function is $\alpha((b, s, \tau, \theta), t) = \alpha_{\theta} \mathbf{1}\{\tau \leq t \leq \tau + s\}$, that is, advertisers match with their type-dependent probability only during their campaign.
- The deterministic drift is given by $v((b, s, \tau, \theta), t) = -e_s \mathbf{1}\{\tau \leq t \leq \tau + s\}$, where e_s is a unit vector that is one for the remaining campaign length coordinate and zero elsewhere. That is, the advertisers remaining campaign length decreases uniformly during their campaign.
- The noise terms $\xi_{k,n}$ are Uniform with support $[0, 1]$ and determine the realization of values through the mapping $F_{\theta_k}^{-1}(\cdot)$.
- The interaction function $f(\vec{x}, \vec{\xi}, t) = \vec{y}$ gives the expenditure \vec{y} when advertisers with states \vec{x} and noise terms $\vec{\xi}$ participate in a second-price auction with reserve price t . Let $x_k = (b_k, s_k, \tau_k, \theta_k)$ be the state of the k^{th} matching bidder. Her value is given $v_k = F_{\theta_k}^{-1}(\xi_k)$, and her bid is $w_k = \beta_{\theta_k}(v_k) \mathbf{1}\{b_k > 0\}$. The competing bid observed by the advertiser is $d_k = \max(r, \max_{i \neq k} w_i)$, while her payment is $p_k = d_k \mathbf{1}\{w_k > d_k\}$. Finally, the output additive change is such that the budget is decreased by the payment, i.e., $y_k = (-p_k, 0, 0, 0)$.

A few remarks are in order. First, note that the interaction function is symmetric and uniformly bounded by \bar{V} , while the probability interaction function is uniformly bounded by $\bar{\alpha} = \max_{\theta} \alpha_{\theta}$. Also, the dynamics guarantee that the budgets and campaign length remaining remain unchanged before and after the campaign. Finally, interaction function is independent of the matching bidders' time in system and campaign start time; since matching bidders are, by definition, alive at the time of the auction.

Until now we have specified the dynamics in the exchange, which together with the initial conditions would give the complete evolution of the advertisers in the exchange. Before specifying the initial conditions we define a distribution for the states of the agents interacting in the closed market based on a consistent distribution for the BMFM. In the following, let $\mathbb{P}_e^{\text{BMFM}}(\mathcal{B}, \mathcal{S} | \theta, t) \triangleq \mathbb{P}_e^{\text{BMFM}}(B \in \mathcal{B}, S \in \mathcal{S} | \Theta = \theta, t)$ be the time-dependent consistent distribution for the BMFM that gives the probability that budget and campaign remaining at time t of a θ -type advertiser in the BMFM lie in the Borel-sets \mathcal{B} and \mathcal{S} , respectively. Note that $\mathbb{P}_e^{\text{BMFM}}$ is constructed from the distribution specified in Definition 1 before considering the uniform sampling in the campaign length. That is, denoting by \mathbb{P}_e the time-invariant consistent distrusting of the BMFM, we have that the time-dependent distribution satisfies the equation

$$\hat{p}_{\theta} \int_0^{s_{\theta}} \mathbb{P}_e^{\text{BMFM}}(\mathcal{B}, s_{\theta} - u | \theta, u) \mathbf{1}_{\{s_{\theta} - u \in \mathcal{S}\}} \frac{1}{s_{\theta}} du = \mathbb{P}_e(\mathcal{B}, \mathcal{S}, \theta). \quad (15)$$

Note that in the BMFM, the time is relative to the start of the advertiser's campaign, as opposed to the time in the closed system, which is relative to the system creation.

Definition 3. Let $\mathbb{P}_e^{\text{BMFM}}(\mathcal{B}, \mathcal{S} | \theta, t)$ be a time-dependent consistent distribution for the BMFM. The induced distribution for the closed market, denoted by $\mathbb{P}_c^{\text{BMFM}}(B \in \mathcal{B}, S \in \mathcal{S}, T = \tau, \Theta = \theta, t)$, is given as follows. First, the probability that an advertiser is of type $\theta \in \Theta$ is

$$\mathbb{P}_c^{\text{BMFM}}(\Theta = \theta, t) = \frac{\lambda_{\theta}^T}{\sum_{\theta'} \lambda_{\theta'}^T} = p_{\theta} \frac{s_{\theta} + T}{\mathbb{E}s_{\Theta} + T}.$$

Second, conditional on an advertiser being of type θ , the advertiser's campaign start time is Uniform with support $[-s_{\theta}, T]$, that is,

$$\mathbb{P}_c^{\text{BMFM}}(T \in \mathcal{T} | \theta = \theta, t) = \frac{|[-s_{\theta}, T] \cap \mathcal{T}|}{T + s_{\theta}}.$$

Finally, the budgets and campaign remaining conditional on an arrival time and type are distributed as

$$\mathbb{P}_c^{\text{BMFM}}\{B \in \mathcal{B}, S \in \mathcal{S} | \Theta = \theta, T = \tau, t\} = \mathbb{P}_e^{\text{BMFM}}(\mathcal{B}, \mathcal{S} | \theta, \text{Proj}_{[0, s_{\theta}]}(t - \tau)),$$

where $\text{Proj}_{[a,b]}(x) = \min(\max(a, x), b)$ is the projection of x to the interval $[a, b]$.

When we specify the initial conditions X_0 as drawn from $\mathbb{P}_c^{\text{BMFM}}$ at time $t = 0$ we get that (i) the initial number of advertisers and their remaining campaign lengths are drawn as from the steady-state, and (ii) departures and arrivals during the horizon $[0, T]$ follow the queue dynamics. Additionally, in the case that the advertiser arrives after the system is created (the campaign starts after time zero) the budgets and campaign remaining are set to the initial values as given by the type. In the case that the advertiser arrived before the system is created, the initial states are drawn from the consistent distribution of the BMFM. We have the following result.

Proposition 7. *Let $\mathbb{P}_e^{\text{BMFM}}$ be a time dependent consistent distribution for the BMFM, and $\mathbb{P}_c^{\text{BMFM}}$ the induced distribution for the closed market, as given by Definition 3. Then, $\mathbb{P}_c^{\text{BMFM}}$ is T -consistent for the mean-field model of the closed market. Moreover, the law of the state of an advertiser in the closed mean-field model during her campaign coincides with that of an advertiser in the BMFM.*

Proof: In order to prove the result we look at the mean-field model associated to the closed system when the states of the competing advertisers are drawn from the distribution $\mathbb{P}_c^{\text{BMFM}}$, which is determined from a consistent distribution of the BMFM through Definition 3. First, we show that the number of matching bidders in the closed mean-field model is time-invariant and distributed as in the BMFM. Second, we show that the distribution of the states of the matching advertisers in the closed mean-field model is time-invariant and coincides with that of the BMFM. Finally, we show the latter two points, together with the consistency of the BMFM, imply the consistency for the closed mean-field model.

First, note that the probability that an advertiser matches in the closed mean-field model is given

$$\begin{aligned} \tilde{\alpha}(t) &= \int_{\mathbb{X}} \alpha(x, t) \, d\mathbb{P}_c^{\text{BMFM}}(x, t) = \int_{\mathbb{X}} \alpha_{\theta_x} \mathbf{1}\{\tau_x \in [t - s_{\theta_x}, t]\} \, d\mathbb{P}_c^{\text{BMFM}}(x, t) \\ &= \int_{\mathbb{X}} \alpha_{\theta_x} \frac{s_{\theta_x}}{T + s_{\theta_x}} \, d\mathbb{P}_c^{\text{BMFM}}(x, t) = \sum_{\theta} p_{\theta} \alpha_{\theta} \frac{s_{\theta}}{\mathbb{E}[s_{\Theta}] + T} = \frac{\mathbb{E}[\alpha_{\Theta} s_{\Theta}]}{\mathbb{E}[s_{\Theta}] + T}, \end{aligned} \quad (16)$$

where the second equality follows from the definition of the matching probability function, the third from conditioning on the type and taking expectation w.r.t. campaign arrival time, and the fourth from taking expectations w.r.t. the types. The resulting matching probability is time-invariant. Additionally, since \tilde{K}_n is Poisson with mean $\sum_{\theta} \lambda_{\theta}^T$ we get from (13) that the number of advertisers matching in the closed mean-field model \tilde{M}_n is Poisson with mean $\tilde{\alpha}(t)\mathbb{E}[K] = \lambda\mathbb{E}[\alpha_{\Theta} s_{\Theta}]$. The latter coincides with the number of competing advertisers in the BMFM.

Second, in the AdX market the remaining campaign length is determined by the campaign starting time, as given by $s = \text{Proj}_{[0, s_{\theta}]}(\tau + s_{\theta} - t)$, and conditioning on the campaign being active at time t it

should be the case that $s = \tau + s_\theta - t$. Therefore, using (14) we obtain that an agent's state conditional on interacting at time t is distributed as:

$$\begin{aligned}
& \tilde{\mathbb{P}}_c^{\text{BMFM}}(\mathcal{B}, \mathcal{S}, [-s_\theta, T], \theta, t) \\
&= \frac{1}{\tilde{\alpha}(t)} \int_{\mathbb{X}} \alpha(x, t) \mathbf{1}\{b_x \in \mathcal{B}, s_x \in \mathcal{S}, \tau_x \in [-s_\theta, T], \theta_x = \theta\} d\mathbb{P}_c^{\text{BMFM}}(x, t) \\
&= \frac{1}{\tilde{\alpha}(t)} \int_{\mathbb{X}} \alpha_{\theta_x} \mathbf{1}\{b_x \in \mathcal{B}, s_x \in \mathcal{S}, \tau_x \in [t - s_\theta, t], \theta_x = \theta\} d\mathbb{P}_c^{\text{BMFM}}(x, t) \\
&= \frac{1}{\tilde{\alpha}(t)} \int_{\mathbb{X}} \alpha_{\theta_x} \mathbb{P}_e^{\text{BMFM}}(\mathcal{B}, \tau_x + s_\theta - t \mid \theta_x, t - \tau_x) \dots \\
&\quad \mathbf{1}\{\tau_x + s_\theta - t \in \mathcal{S}, \tau_x \in [t - s_\theta, t], \theta_x = \theta\} d\mathbb{P}_c^{\text{BMFM}}(x, t) \\
&= \frac{s_\theta + T}{\mathbb{E}[\alpha_{\Theta} s_{\Theta}]} p_\theta \alpha_\theta \int_{t-s_\theta}^t \mathbb{P}_e^{\text{BMFM}}(\mathcal{B}, \tau + s_\theta - t \mid \theta, t - \tau) \mathbf{1}\{\tau + s_\theta - t \in \mathcal{S}\} \frac{1}{s_\theta + T} d\tau \\
&= \frac{p_\theta \alpha_\theta s_\theta}{\mathbb{E}[\alpha_{\Theta} s_{\Theta}]} \int_0^{s_\theta} \mathbb{P}_e^{\text{BMFM}}(\mathcal{B}, s_\theta - u \mid \theta, u) \mathbf{1}\{s_\theta - u \in \mathcal{S}\} \frac{1}{s_\theta} du \\
&= \hat{p}_\theta \int_0^{s_\theta} \mathbb{P}_e^{\text{BMFM}}(\mathcal{B}, s_\theta - u \mid \theta, u) \mathbf{1}\{s_\theta - u \in \mathcal{S}\} \frac{1}{s_\theta} du = \mathbb{P}_e(\mathcal{B}, \mathcal{S}, \theta)
\end{aligned}$$

where the second equality follows from the definition of the matching probability function; the third from conditioning on the type and campaign start time, taking expectations w.r.t. the budgets and campaign length remaining, using Definition 3, and using that at point $t - \tau$ of the campaign the campaign length remaining is $\tau + s_\theta - t$; the fourth from taking expectations w.r.t. the type and campaign starting time that is Uniform with support $[-s_\theta, T]$ and equation (16); the fifth from performing the change of variables $u = t - \tau$; the sixth from our formula for the probability that a matching advertiser is steady-state is of type θ ; and the last from equation (15). Note that the distribution of the states of the matching advertisers in the closed mean-field model is time-invariant, and coincides with that of the BMFM.

The previous results show that the number of competing bidders and the distribution of their states in both mean-field model coincide. Therefore, the dynamics of the closed mean-field model and the BMFM are the same during an advertiser's campaign. Moreover, one can observe that the initial conditions in both models coincide. In the case that the campaign starts after time zero, $\tau \geq 0$, the projection operator guarantees that the campaign starts with the initial budgets and campaign length of the BMFM. In the case that the campaign is already active at time zero, $\tau < 0$, the starting budget is drawn from the time-dependent BMFM. Thus, the law of the state of an advertiser in the closed mean-field model during her campaign coincides with that of an advertiser in the BMFM. Finally, T-consistency for the closed mean-field model follows from the consistency of the BMFM. \diamond

D Auxiliary Results

Lemma 1. Fix a vector of multipliers $\boldsymbol{\mu}$ and consider the differentiable vector function $\mathbf{H} : [0, 1]^{|\Theta|} \rightarrow \mathbb{R}_+^{|\Theta|}$ given by

$$H_\theta(\mathbf{q}) = q_\theta \mathbb{E}[Z_\theta(\mathbf{q})] = q_\theta r \bar{F}_v((1 + \mu_\theta)r) F_d(r; \mathbf{q}) + q_\theta \int_r^{\bar{V}} x \bar{F}_v((1 + \mu_\theta)x) dF_d(x; \mathbf{q}).$$

Suppose that there are at most two types. Then, the Jacobian of \mathbf{H} is a P-matrix.

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

Step 1. Since the cumulative distribution of values are differentiable, the distribution of the maximum bid is differentiable w.r.t. x and \mathbf{q} . Indeed, its partial derivatives are given by $\partial F_d / \partial q_\theta = -F_d(x; \mathbf{q}) \mathbb{E}[\alpha_\Theta \lambda s_\Theta] \hat{p}_\theta \bar{F}_{v_\theta}((1 + \mu_\theta)x)$, and $\partial F_d / \partial x = F_d(x; \mathbf{q}) \mathbb{E}[\alpha_\Theta \lambda s_\Theta] \sum_\theta \hat{p}_\theta q_\theta (1 + \mu_\theta) f_{v_\theta}((1 + \mu_\theta)x)$. Moreover, the second derivatives of the distribution of the maximum bid are continuous because densities $f_{v_\theta}(\cdot)$ are continuously differentiable.

The partial derivative of one type's expenditure w.r.t. her active probability is

$$\frac{\partial H_\theta}{\partial q_\theta} = \mathbb{E}[Z_\theta(\mathbf{q})] + q_\theta \frac{\partial}{\partial q_\theta} \mathbb{E}[Z_\theta(\mathbf{q})]$$

where

$$\begin{aligned} \frac{\partial}{\partial q_\theta} \mathbb{E}[Z_\theta(\mathbf{q})] &= r \bar{F}_{v_\theta}((1 + \mu_\theta)r) \frac{\partial F_d}{\partial q_\theta}(\mathbf{q}; r) + \frac{\partial}{\partial q_\theta} \int_r^{\bar{V}} x \bar{F}_{v_\theta}((1 + \mu_\theta)x) \frac{\partial F_d}{\partial x} dx \\ &= r \bar{F}_{v_\theta}((1 + \mu_\theta)r) \frac{\partial F_d}{\partial q_\theta}(\mathbf{q}; r) + \int_r^{\bar{V}} x \bar{F}_{v_\theta}((1 + \mu_\theta)x) \frac{\partial^2 F_d}{\partial q_\theta \partial x} dx \\ &= - \int_r^{\bar{V}} \frac{\partial}{\partial x} (x \bar{F}_{v_\theta}((1 + \mu_\theta)x)) \frac{\partial F_d}{\partial q_\theta} dx \end{aligned}$$

where the second equality follows from exchanging integration and differentiation, which holds because $[\underline{V}, \bar{V}] \times U$ is bounded and integrand is continuously differentiable; and the third from exchanging partial derivatives by Clairaut's theorem, integrating by parts, canceling terms, and using the fact that $\bar{F}_{v_\theta}((1 + \mu_\theta)\bar{V}) = 0$. Using the same notation that in the proof of Lemma 2 of the main paper and canceling terms we can write

$$\frac{\partial H_\theta}{\partial q_\theta} = \sum_{\theta' \neq \theta} \hat{p}_{\theta'} (1 + \mu_{\theta'}) q_{\theta'} \langle f_{\theta'}, \bar{F}_\theta \rangle + \hat{p}_\theta q_\theta \langle \bar{F}_\theta, \bar{F}_\theta \rangle + r \bar{F}_{v_\theta}((1 + \mu_\theta)r) F_d(x; \mathbf{q}). \quad (17)$$

Similarly, the partial derivative of one type's expenditure w.r.t. another type's active probability is

$$\begin{aligned}\frac{\partial H_\theta}{\partial q'_\theta} &= q_\theta \frac{\partial}{\partial q'_\theta} \mathbb{E}[Z_\theta(\mathbf{q})] = -q_\theta \int_r^{\bar{V}} \frac{\partial}{\partial x} (x \bar{F}_{v_\theta}((1 + \mu_\theta)x)) \frac{\partial F_d}{\partial q'_\theta} dx \\ &= -\hat{p}_{\theta'}(1 + \mu_\theta)q_\theta \langle f_\theta, \bar{F}_{\theta'} \rangle + \hat{p}_{\theta'}q_\theta \langle \bar{F}_\theta, \bar{F}_{\theta'} \rangle.\end{aligned}\tag{18}$$

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

$$J_{\mathbf{H}} = \begin{pmatrix} \frac{\partial H_1}{\partial q_1} & \frac{\partial H_1}{\partial q_2} \\ \frac{\partial H_2}{\partial q_1} & \frac{\partial H_2}{\partial q_2} \end{pmatrix}.$$

From (17) one concludes that the principal minors $J|_{\{1\}}$ and $J|_{\{2\}}$ have positive determinant (they are, in fact, positive scalars). The determinant of the remaining minor $J|_{\{1,2\}}$ is that of the whole Jacobian, which is given by

$$\begin{aligned}\det(J) &= \frac{\partial H_1}{\partial q_1} \frac{\partial H_2}{\partial q_2} - \frac{\partial H_1}{\partial q_2} \frac{\partial H_2}{\partial q_1} \\ &= (\hat{p}_2 q_2)^2 (1 + \mu_2) \langle f_2, \bar{F}_1 \rangle \langle \bar{F}_2, \bar{F}_2 \rangle + (\hat{p}_1 q_1)^2 (1 + \mu_1) \langle \bar{F}_1, \bar{F}_1 \rangle \langle f_1, \bar{F}_2 \rangle \\ &\quad + \hat{p}_1 q_1 (1 + \mu_1) \hat{p}_2 q_2 \langle f_1, \bar{F}_2 \rangle \langle \bar{F}_2, \bar{F}_1 \rangle + \hat{p}_1 q_1 \hat{p}_2 q_2 (1 + \mu_2) \langle \bar{F}_1, \bar{F}_2 \rangle \langle f_2, \bar{F}_1 \rangle \\ &\quad + \hat{p}_1 q_1 \hat{p}_2 q_2 \langle \bar{F}_1, \bar{F}_1 \rangle \langle \bar{F}_2, \bar{F}_2 \rangle - \hat{p}_1 q_1 \hat{p}_2 q_2 \langle \bar{F}_1, \bar{F}_2 \rangle \langle \bar{F}_2, \bar{F}_1 \rangle\end{aligned}$$

where the third equation follows from substituting the expressions for the partial derivatives and canceling two terms (here we assumed, without loss of generality, that $r = 0$ since the sum of a positive diagonal matrix with a P-matrix is a P-matrix). Notice that all terms are positive with the exception of the last one. We conclude that the determinant is positive by invoking Cauchy-Schwartz inequality to show that the fifth term dominates the last one. \square

Lemma 2. Consider a sequence of continuous mappings $\{\mathbf{g}^n\}_{n \geq 1}$ with $\mathbf{g}^n : [0, 1]^d \rightarrow [0, 1]^d$ converging uniformly to a continuous mapping \mathbf{g} . Let $X^n = \{\mathbf{x} \in [0, 1]^d : \mathbf{g}^n(\mathbf{x}) = \mathbf{x}\}$ be the set of fixed points of \mathbf{g}^n , and $X = \{\mathbf{x} \in [0, 1]^d : \mathbf{g}(\mathbf{x}) = \mathbf{x}\}$ be the set of fixed points of \mathbf{g} . Then $\lim_{n \rightarrow \infty} \mathbb{D}_\infty(X^n, X) = 0^6$.

Proof. We argue by contradiction. Suppose that $\mathbb{D}_\infty(X^n, X)$ does not converge to zero. Since the set $[0, 1]^d$ is compact, by passing to a subsequence if necessary, we can assume that there exists $\mathbf{x}^n \in X^n$ such that $\text{dist}(\mathbf{x}^n, X) \geq \epsilon$ for some $\epsilon > 0$ and that \mathbf{x}^n converges to a point $\mathbf{x}^* \in [0, 1]^d$. It follows that

⁶We denote the deviation of two sets A and B by $\mathbb{D}(A, B) = \sup_{x \in A} \text{dist}(x, B)$, and the distance between a point x and a set B as $\text{dist}(x, B) = \inf_{y \in B} \|x - y\|_\infty$.

$\mathbf{x}^* \notin X$. But notice that

$$\begin{aligned}\|\mathbf{x}^* - \mathbf{g}(\mathbf{x}^*)\| &\leq \|\mathbf{x}^* - \mathbf{x}^n\| + \|\mathbf{g}^n(\mathbf{x}^n) - \mathbf{g}(\mathbf{x}^n)\| + \|\mathbf{g}(\mathbf{x}^n) - \mathbf{g}(\mathbf{x}^*)\| \\ &\leq \|\mathbf{x}^* - \mathbf{x}^n\| + \sup_{\mathbf{x}} \|\mathbf{g}^n(\mathbf{x}) - \mathbf{g}(\mathbf{x})\| + \|\mathbf{g}(\mathbf{x}^n) - \mathbf{g}(\mathbf{x}^*)\|,\end{aligned}$$

where we used the fact that $\mathbf{x}^n = \mathbf{g}^n(\mathbf{x}^n)$ together with the triangle inequality. We have that the first term of the right-hand side converges to zero from compactness, the second from uniform convergence, and the last from continuity of \mathbf{g} . Thus, we obtain that $\mathbf{x}^* = \mathbf{g}(\mathbf{x}^*)$, a contradiction. \square