Competition and Yield Optimization in Ad Exchanges

Santiago R. Balseiro

Submitted in partial fulfillment of the
requirements for the degree
of Doctor of Philosophy
under the Executive Committee
of the Graduate School of Arts and Sciences

COLUMBIA UNIVERSITY

2013

ABSTRACT

Competition and Yield Optimization in Ad Exchanges

Santiago R. Balseiro

Ad Exchanges are emerging Internet markets where advertisers may purchase display ad placements, in real-time and based on specific viewer information, directly from publishers via a simple auction mechanism. The presence of such channels presents a host of new strategic and tactical questions for publishers. How should the supply of impressions be divided between bilateral contracts and exchanges? How should auctions be designed to maximize profits? What is the role of user information and to what extent should it be disclosed? In this thesis, we develop a novel framework to address some of these questions. We first study how publishers should allocate their inventory in the presence of these new markets when traditional reservation-based ad contracts are available. We then study the competitive landscape that arises in Ad Exchanges and the implications for publishers' decisions.

Traditionally, an advertiser would buy display ad placements by negotiating deals directly with a publisher, and signing an agreement, called a guaranteed contract. These deals usually take the form of a specific number of ad impressions reserved over a particular time horizon. In light of the growing market of Ad Exchanges, publishers face new challenges in choosing between the allocation of contract-based reservation ads and spot market ads. In this setting, the publisher should take into account the tradeoff between short-term revenue from an Ad Exchange and the long-term impact of assigning high quality impressions to the reservations (typically measured by the click-through rate). In the first part of this thesis, we formalize this combined optimization problem as a stochastic control problem and derive an efficient policy for online ad allocation in settings with general joint distribution over placement quality and exchange bids, where the exchange bids are assumed to be exogenous and independent of the decisions of the publishers. We prove asymptotic optimality of this policy in terms of any arbitrary trade-off between quality of delivered reservation ads and revenue from the exchange, and provide a bound for its convergence rate to the optimal policy. We also give experimental results on data derived from real publisher inventory, showing that our policy can achieve any Pareto-optimal point on the quality vs. revenue curve.

In the second part of this thesis, we relax the assumption of exogenous bids in the Ad Ex-

change and study in more detail the competitive landscape that arises in Ad Exchanges and the implications for publishers' decisions. Typically, advertisers join these markets with a pre-specified budget and participate in multiple second-price auctions over the length of a campaign. We introduce the novel notion of a Fluid Mean Field Equilibrium (FMFE) to study the dynamic bidding strategies of budget-constrained advertisers in these repeated auctions. This concept is based on a mean field approximation to relax the advertisers' informational requirements, together with a fluid approximation to handle the complex dynamics of the advertisers' control problems. Notably, we are able to derive a closed-form characterization of FMFE, which we use to study the auction design problem from the publisher's perspective focusing on three design decisions: (1) the reserve price; (2) the supply of impressions to the Exchange versus an alternative channel such as bilateral contracts; and (3) the disclosure of viewers' information. Our results provide novel insights with regard to key auction design decisions that publishers face in these markets.

In the third part of this thesis, we justify the use of the FMFE as an equilibrium concept in this setting by proving that the FMFE provides a good approximation to the rational behavior of agents in large markets. To do so, we consider a sequence of scaled systems with increasing market "size". In this regime we show that, when all advertisers implement the FMFE strategy, the relative profit obtained from any unilateral deviation that keeps track of all available information in the market becomes negligible as the scale of the market increases. Hence, a FMFE strategy indeed becomes a best response in large markets.

Table of Contents

1 Introduction **1**

1.1 Display Advertising and Ad Exchanges . . . 1
1.2 Research Questions and Summary . . . 2
1.3 Yield Optimization of Guaranteed Contracts with AdX . . . 4
1.4 Advertiser Competition in AdX: Auction Design . . . 7
1.5 Advertiser Competition in AdX: Approximation Results . . . 10
1.6 Related Literature . . . 11
1.7 Conclusions . . . 15

2 Yield Optimization of Guaranteed Contracts with AdX **19**

2.1 Model Description . . . 20
2.1.1 Objective . . . 22
2.1.2 AdX Model with User Information . . . 23
2.1.3 Discussion of the Assumptions . . . 24
2.2 Problem Formulation . . . 25
2.2.1 Dynamic Programming Formulation . . . 25
2.2.2 Deterministic Approximation Problem (DAP) . . . 27
2.2.3 Our Stochastic Policy . . . 30
2.2.4 Asymptotic Analysis . . . 31
2.2.5 Impact of User Information . . . 33
2.3 Data Model and Estimation . . . 34
2.3.1 Estimation of Placement Qualities . . . 35
2.3.2 Estimation of AdX Bids . . . 36
2.4 Experimental Results . . . 37
2.4.1 Impact of AdX . . . 38
2.4.2 Comparison with Greedy and Static Price Policy . . . 41

2.4.3 Large-scale Instances 43
2.5 Extensions 45
2.5.1 AdX with Multiple Bidders 46
2.5.2 Covering Constraints 47
2.5.3 Target Quality Constraints 47
2.5.4 Handling ties 48
2.6 Conclusions 51

3 Advertiser Competition in AdX: Auction Design 53

3.1 Model Description 53
3.2 Equilibrium Concept 55
3.2.1 Mean Field and Fluid Approximation 56
3.2.2 Fluid Mean Field Equilibrium 58
3.2.3 Publisher's Problem 60
3.3 Fluid Mean Field Equilibrium Characterization 61
3.3.1 Equilibrium Existence and Sufficient Conditions for Uniqueness 61
3.3.2 Equilibrium Characterization 62
3.3.3 Reformulation of the Publisher's Problem 64
3.4 Auction Design and Allocation Decisions 64
3.4.1 The Case of Homogeneous Advertisers 65
3.4.2 Numerical Results for the Case of Heterogeneous Advertisers 70
3.5 Conclusions 73

4 Advertiser Competition in AdX: Approximation Results 75

4.1 Approximation Result for Synchronous Campaigns 76
4.2 Approximation Result for Asynchronous Campaigns 79
4.3 Budget-constrained Mean Field Model 83
4.3.1 Existence of a consistent BMFM 84
4.3.2 Active Bidders 85
4.3.3 Payoff Evaluation in the BMFM 86
4.4 Propagation of Chaos in the BMFM 87

Bibliography 88

A Appendix to Chapter 2 97

A.1 Proofs of Statements 97

A.1.1 Proof of Proposition 2.1 97
A.1.2 Proof of Theorem 2.1 97
A.1.3 Proof of Proposition 2.2 98
A.1.4 Proof of Proposition 2.3 98
A.1.5 Proof of Theorem 2.2 100
A.1.6 Proof of Corollary 2.1 102
A.1.7 Proof of Proposition 2.4 102
A.1.8 Directional Derivatives of the Dual Objective 103
A.1.9 Proof of Proposition 2.5 104
A.2 Comparison with the Primal-Dual Method 106
A.2.1 Numerical Experiments 110
A.3 Incorrect Assignments in the User Type Model 114
A.4 Computation 116
A.5 Fluid Limit 118

B Appendix to Chapter 3 **123**
B.1 Proof of Statements 123
B.1.1 Proof of Proposition 3.1 123
B.1.2 Proof of Theorem 3.1 124
B.1.3 Proof of Theorem 3.2 125
B.1.4 Proof of Proposition 3.2 126
B.1.5 Proof of Theorem 3.3 126
B.1.6 Proof of Theorem 3.4 128
B.1.7 Proof of Corollary 3.2 132
B.1.8 Proof of Theorem 3.5 132
B.1.9 General properties of the expenditure function 135
B.2 Sufficient Conditions for P-matrix Assumption to Hold 138
B.2.1 Proof of Proposition B.1 138

C Appendix to Chapter 4 **141**
C.1 Proof of Statements for Synchronous Campaigns 141
C.1.1 Proof of Theorem 4.1 141
C.1.2 Proof of Proposition 4.1 141
C.1.3 Proof of Proposition 4.2 143
C.1.4 Additional Results 147

C.2 Proof of Statements for Asynchronous Campaigns . . . 150
C.2.1 Proof of Theorem 4.2 . . . 150
C.2.2 Proof of Proposition 4.3 . . . 150
C.2.3 Proof of Proposition 4.4 . . . 151
C.2.4 Proof of Proposition 4.5 . . . 155
C.2.5 Proof of Proposition 4.6 . . . 157
C.2.6 Proof of Proposition 4.7 . . . 158
C.3 Mean-field Model for Systems with a Random Number of Agents . . . 161
C.3.1 Real System . . . 161
C.3.2 Mean-field Model . . . 162
C.3.3 Boltzmann Tree . . . 163
C.3.4 Propagation of Chaos . . . 165
C.3.5 Evaluating Deviations . . . 172
C.3.6 Heterogeneous interaction probabilities . . . 173
C.4 Ad Exchange Market as a Closed System . . . 175
C.5 Auxiliary Results . . . 179

Acknowledgments

Looking back on my years at Columbia, I owe this great experience to many people.

To Omar, for being literally always available whenever I knocked his door, for those great discussions, and for teaching me how to distill a problem to its essence. You were a great advisor.

To Gabriel, for attempting so hard to shape my way of thinking into an economic mindset. You taught me some valuable lessons that I will –hopefully– never forget.

To Awi, for teaching me the values of true scholarship, and for all the support provided during my PhD. You were a great mentor.

To Jon, Muthu and Vahab, for giving me the opportunity to be part of their team and introducing me to the world of Ad Exchanges. I enjoyed some amazing summers with you at Google.

To Gerardo and Irene, for encouraging me to follow my dream of pursuing graduate studies, and helping me get there. This thesis would not have existed without you.

To all my friends during the PhD, for making those long hours in Uris barely bearable. We had lots of fun together.

To Nico, for his continued support and valuable advice all throughout my degree. I will always be in debt to you for your help during my first years.

To my mother, father and brother; for believing in me from day one, and for being the best role models one can hope for. No words can express my love for you.

To my wife, Paulita, for always being cheerful and picking me up when I was down, for not letting me work during the weekends, and for teaching me how to enjoy life. You are my muse, love and inspiration.

Chapter 1

Introduction

1.1 Display Advertising and Ad Exchanges

The market for display ads on the internet, consisting of graphical content such as banners and videos on web pages, has grown significantly in the last decade, generating about 11 billion dollars in the United States in 2011 (Internet Advertising Bureau, 2012). Traditionally, an advertiser would buy display ad placements by negotiating deals directly with a publisher (the owner of the web page), and signing an agreement, called a guaranteed contract. These deals usually take the form of a specific number of ad *impressions* or eyeballs reserved over a particular time horizon. For example, a publisher (such as the New York Times) might sign a contract with an advertiser (such as Macy's) agreeing to deliver one million impressions to females living in New York. A publisher can make many such deals with different advertisers, with potentially sophisticated relationships between the advertisers' targeting criteria.

The growth of display advertising has been accompanied by the emergence of alternative channels for the purchase of display ads. Advertisers may now purchase ad placements through a spot market for the real-time sale of online ad slots, called Ad Exchanges (AdX), which are essentially platforms that operate as intermediaries between online publishers and advertisers. These markets epitomize the vertiginous growth rate of the digital economy: only a few years after their introduction, some exchanges were already running billions of transactions per day vastly exceeding the total number of transactions at financial exchanges worldwide (Mansour et al., 2012). Ad Exchanges allow advertisers to purchase ad placements, in real-time and based on specific viewer information, directly from publishers via a simple auction mechanism. As a result, publishers and advertisers interact in spot markets where decisions must be made automatically in milliseconds. Google's DoubleClick, OpenX, and Yahoo!'s Right Media are prominent examples of such exchanges.

While exchanges differ in their implementations, a generic AdX works as follows. When a user visits a web page, the publisher posts the ad slot in the exchange together with potentially some user information known to her; e.g., the user's geographical location and her *cookies*. Advertisers (or bidders) interested in advertising on the site post bids as a function of their targeting criteria and the user information provided by the publisher. Then, an auction is run to determine the winning advertiser and the ad to be shown to the user. The latter process happens in milliseconds, between the time a user requests a page and the time the page is displayed to her. The publisher repeatedly offers slots to display advertisements on her web-site as users arrive; typically, a given publisher runs millions of these auctions per day. See Mansour et al. (2012) for a more detailed description of Ad Exchanges. A generic AdX model and the timing of events is shown in Figure 1.1.

On its part, advertisers participate in the exchange with the objective of fulfilling marketing campaigns. In practice, such campaigns are commonly based on a given pre-determined budget and extend for a fixed amount of time, over which advertisers participate in a large volume of auctions. Given the large number opportunities and the time scale on which decisions are made, bidding is fully automated. AdX allows advertisers to bid in real-time and pay only for valuable customers, in sharp contrast to the bulk buying of impressions and broader targeting of guaranteed contracts.

1.2 Research Questions and Summary

The emergence of Ad Exchangess presents a rich set of new strategic and tactical questions for publishers, exchanges, and advertisers. How should publishers divide the allocation of impressions between bilateral contracts and exchanges? How should AdX auctions be designed to maximize profits? What is the role of user information and to what extent should it be disclosed? How should advertisers bid in the face of competition? This thesis sheds light on these fundamental issues.

In Chapter 2 of this thesis, we study the problem faced by the publisher, jointly optimizing over AdX and guaranteed contracts. In presence of Ad Exchange, the publisher must quickly decide, for each arriving user, whether to send the inventory to AdX or assign the slot to the best matching reservation from the guaranteed contracts. Hence publishers face the multi-objective problem of maximizing the overall placement quality[1] of the impressions assigned to the reservations together with the total revenue obtained with AdX, while complying with the contractual obligations.

[1]A typical measure of placement quality in the internet advertising industry is the probability that a user clicks on an ad (known as *click-trough rate*).

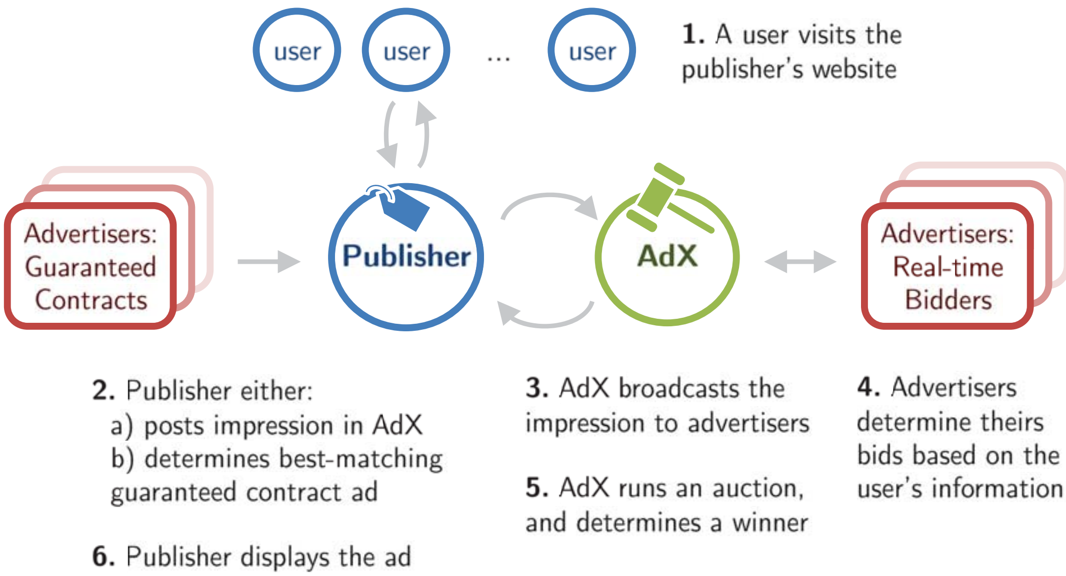

Figure 1.1: A generic AdX model with the timing of the events.

These two objectives are potentially conflicting; in the short-term, the publisher might boost the revenue stream from AdX at the expense of assigning lower quality impressions to the advertisers. In the long term, however, it may be convenient for the publisher to prioritize her advertisers in view of attracting future contracts. We formalize this combined optimization problem as a stochastic control problem and derive an efficient policy for online ad allocation in settings with general joint distribution over placement quality and AdX bids, where the exchange bids are assumed to be exogenous and independent of the decisions of the publishers. We prove asymptotic optimality of this policy in terms of any arbitrary trade-off between quality of delivered reservation ads and revenue from the exchange, and provide a bound for its convergence rate to the optimal policy. We also give experimental results on data derived from real publisher inventory, showing that our policy can achieve any Pareto-optimal point on the quality vs. revenue curve.

In Chapter 3 of this thesis, we study the competitive landscape that arises in Ad Exchanges and the implications for publishers' decisions. In a major departure from previous work, we consider a model in which the competitive landscape is endogenous, that is, advertisers engage in a dynamic game, whose equilibrium outcome determines the market characteristics. Traditional equilibrium concepts from game theory fail to provide insightful or tractable solution concepts given the complex nature of the game. However, leveraging the large number of players and transactions involved in Ad Exchanges, we introduce the novel notion of a Fluid Mean Field Equilibrium (FMFE), which is behaviorally appealing, computationally tractable, and, in some cases of interest, leads to insightful closed form solutions. Furthermore, we use the FMFE framework to provide sharp prescriptions for key auction design decisions that publishers face in these markets, such as the reserve price, the allocation of impressions to the exchange versus an alternative channel, and the disclosure of viewers' information. Notably, we show that proper adjustment of the reserve price is key in (1) making it profitable for the publisher to try selling all impressions in the exchange before utilizing the alternative channel; and (2) compensating for the thinner markets created by greater disclosure of viewers' information. More generally, the FMFE provides a new lens through which one may analyze these markets.

In Chapter 4, we justify the use of the FMFE as an equilibrium concept in this setting by proving that the FMFE provides a good approximation to the rational behavior of agents in large markets. To do so, we consider a sequence of scaled systems with increasing market "sizes", that is, increasing number of agents and auctions. In this regime, we show that when all advertisers implement the FMFE strategy, the relative profit obtained from any unilateral deviation that keeps track of all available information in the market, becomes negligible as the scale of the market increases. Hence, a FMFE strategy indeed becomes a best response in large markets. This result is proven by combining techniques from the revenue management and the mean field literatures; essentially, showing that the Fluid approximation becomes asymptotically optimal as the number of opportunities during the campaign grows, and that the Mean Field approximation becomes asymptotically correct as the number of advertisers competing in the exchange grows.

1.3 Yield Optimization of Guaranteed Contracts with AdX

The contributions of Chapter 2 are as follows.

Firstly, we bring to bear techniques of Revenue Management (RM), and model the publisher's problem as a combination of a *capacity allocation problem* to handle the guaranteed contracts together with a *dynamic pricing problem* to handle the reserve price optimization in the Ad Exchange. We tackle the publisher's multi-objective problem by taking a weighted sum of (i) the revenue from AdX, and (ii) the placement quality of the contracts, and show how to construct the the Pareto efficient frontier of attainable objectives. The Pareto frontier provides managers with a mechanism for visualizing the trade-off between these goals. Using data derived from real publisher inventory, we show empirically that the Pareto efficient is highly concave and that there are significant benefits for the publishers if they jointly optimize over both channels.

The publisher's problem can be thought of as a parallel-flight Network RM problem (see, e.g., Talluri and van Ryzin (1998)) in which users' click probabilities are requests for itineraries, and advertisers are edges in the network. As in the prototypical RM problem, we look for a policy maximizing the *ex-ante* expected revenue, which can be obtained using dynamic programming

	Network RM	**Display Ad**
Resources	Seats	Impressions
Edges	Flight Legs	Contracts
Capacity Constraints	$\leq$	$=$
Objective	Maximize Revenue	Maximize Placement Quality
Decision	Accept/reject itinerary request	Determine best matching contract
Spot market / Dynamic pricing	No	Yes

Table 1.1: Comparison of the Display Ad and Network Revenue Management problems.

(DP). There are three differences, however, with the traditional Network RM problem. First, we aim to satisfy all contracts, or completely deplete all resources by the end of the horizon. Second, in the traditional problem requests are for only one itinerary (which can be accepted or rejected), while in our model each impression can be potentially assigned to any contract and the publisher needs to decide whom to assign the impression based on possibly correlated placement qualities. Finally, the publishers in display advertising may submit impressions to a spot market to increase their revenues, which adds a dynamic pricing dimension to the problem. Table 1.1 summarizes these points.

Secondly, because of the so-called "curse of dimensionality" the optimal policy cannot be computed efficiently in most real-world problems, and instead we aim for a deterministic approximation in which stochastic quantities are replaced by their expectation values and quantities assumed to be continuous (Gallego and van Ryzin, 1994). As a result, we derive a provably good policy that resembles a bid-price control but extended with a pricing function to take into account for AdX. Our policy assigns each guaranteed contract a bid-price (or dual variable), which may be interpreted as the opportunity cost of assigning one additional impression to the reservation. When a user arrives, the pricing function quotes a reserve price to submit to the exchange that depends on the opportunity cost of assigning the impression to an advertiser (and potentially on the impression's attributes). If no AdX bid exceeds this reserve price, the impression is immediately assigned to the advertiser whose placement quality exceeds its opportunity cost by the largest amount. A salient feature of our policy from the managerial standpoint is its simplicity: the publisher only needs to keep track of a single pricing function for the exchange, and one bid-price for each contract that are obtained, in turn, by solving a convex stochastic minimization problem.

The optimal policy always tests the exchange before assigning an impression to a guaranteed contract because the loss of not assigning an impression of high quality to the reservation can be compensated by choosing a high enough reserve price. This result implicitly hinges upon (i) the absence of a fixed cost for accessing the exchange, and (ii) the publisher's ability to dynamically adjust the reserve price for each impression based on the user attributes. In the presence of a fixed cost, the publisher tests the exchange only if the contracts' opportunity cost is less or equal than a fixed threshold; while in the case of static pricing, when the expected revenue from AdX exceeds the contracts' opportunity cost.

Thirdly, we introduce a general model of targeting based on the user's attributes that takes into account the potential correlation between guaranteed contracts' placement quality and exchange's bids. In our model, the publisher first determines the placement quality for the contracts based on the user's attributes, and then discloses some of the these attributes to the exchange, where advertisers bid strategically based on this information. Similar targeting criteria across both channels can potentially introduce positive correlation between the placement quality of the contracts and the bids from the exchange. This positive correlation creates two interdependent effects on the publisher's joint allocation problem – a *diversification loss effect* and a *price discrimination effect*. The former is a negative effect; i.e, both channels competing for the same inventory undermines the publisher's ability to extract a higher rent from the impressions that are less attractive to the reservations by testing the exchange. The latter is a positive effect; i.e, correlation allows the publisher to exploit user attributes as a covariate to predict the bids and price more effectively in the exchange. The sum of these effects is, in most cases, indeterminate.

Fourthly, we provide a rigorous bound on the convergence rate of our policy to the optimal policy (Theorem 2.2). Typically ad allocation research compares to the optimal offline policy in hindsight[2]; instead, we compare our policy with an optimal online policy, obtaining a bound of $O(\sqrt{N})$ on additive regret, where N is the number of impressions in the horizon. Our approximation is suitable when the number of impressions in the horizon is large; which fits well in the context of internet advertising. Moreover, from a computational stand-point we provide an efficient and simple method to compute the dual variables that is applicable to large instances with many contracts. Our procedure combines a Sample Average Approximation together with a Subgradient Descent Method.

Finally, we numerically compare the performance of our policy with two alternative heuristics

[2]While in the absence of the spot market the performance of the offline and online policies are asymptotically equivalent (see, e.g., Talluri and van Ryzin (1998)), in the presence of the spot market this is not longer the case if we assume that the oracle is aware of bids' realizations.

that are common in practice. The first heuristic is a Greedy Policy that disregards the opportunity cost of capacity and assigns the impression to the advertiser with maximum placement quality. The second is a Static Price Policy that sets a constant reserve price for the exchange throughout the horizon. Our results on actual publisher data show that these heuristics significantly underperform when compared to the optimal policy. From a managerial perspective, these results stress the importance of pondering the opportunity cost of capacity in performing the assignment to the guaranteed contracts, and of pricing dynamically in the exchange to react to the users' attributes and the value of the reservations.

1.4 Advertiser Competition in AdX: Auction Design

The contributions of Chapter 3 are as follows.

Firstly, we introduce a model of AdX in which the competitive landscape is endogenous, that is, advertisers engage in a dynamic game, whose equilibrium determines the market characteristics. Our model incorporates advertisers' budget constraints that are prevalent in these markets. These constraints link the different auctions over time, and therefore advertisers require dynamic bidding strategies to optimize the allocation of budget to incoming impressions in order to maximize cumulated profits over the length of the campaign. In many cases, advertisers have similar targeting criteria, and bid for the same inventory of ads. Thus, the dynamic bidding strategy an advertiser adopts impacts the competitive landscape for other advertisers in the market. Moreover, the publisher's auction design decisions, such as the reserve price, also impact these interactions. Thus motivated, we formulate our Ad Exchange model as a game among advertisers and the publisher.[3] First, the publisher defines the parameters of a second-price auction that become common knowledge. Then, given the auction format, advertisers compete in a *dynamic* game. In order to quantify the impact of auction design parameters, we focus first on the competitive landscape that emerges for fixed auction decisions.

An important challenge in our analysis is solving for the equilibrium of the dynamic game among advertisers induced by the auction rules. At one extreme of agent sophistication, a notion of equilibrium that one may consider is Perfect Bayesian Equilibrium (PBE) in which advertisers maintain priors on the states of all other bidders, and update them accordingly using Bayes' rule. Even if priors could be succinctly updated, bidders are left with the problem of computing a best response, which is a high-dimensional dynamic program. Such an approach presents two main drawbacks. First, the analysis of the resulting game is, in most cases, intractable from

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

both analytical and computational stand-points. Second, such sophistication and informational requirements on the part of agents is highly unrealistic.

Secondly, we introduce a novel notion of equilibrium that is tractable, appealing from a behavioral perspective, and provides a good approximation to the strategic interactions among budget-constrained bidders in an Ad Exchange. Our new notion of equilibrium combines two different approximations to address the limitations in PBE. First, we consider a *Mean Field* approximation to relax the informational requirements of agents. The motivation behind the mean field approximation is that, when the number of competitors is large, there is little value in tracking the specific actions of all agents and one may rely on some aggregate and stationary representation of the competitors' bids. This type of approximations have appeared in other auction and industrial organization applications (see, e.g., Iyer et al. (2011); Weintraub et al. (2008)). Moreover, in Ad Exchange markets the number of participants is typically large. Second, borrowing techniques from the revenue management literature (see, e.g., Gallego and van Ryzin (1994)), we consider a stochastic *fluid* approximation to handle the complex dynamics of the advertisers' control problem. Such approximations are suitable when the number of opportunities is large and the payment per opportunity is small compared to the budget; hence, it also fits well in the context of Ad Exchanges. Using these two approximations, we define the notion of a *Fluid Mean Field Equilibrium* (FMFE).

Thirdly, we provide a sharp characterization of the equilibrium strategies under the FMFE. Notably, when a second-price auction is conducted, the resulting FMFE strategy has a simple, yet appealing, form: an advertiser needs to shade her values by a *constant factor.* Intuitively, when budgets are tight, advertisers shade their bids, because there is an option value for future good opportunities. We leverage the latter characterization to analyze properties of FMFE. In particular, we show that an FMFE always exists and provide a broad set of sufficient conditions that guarantee its uniqueness. We also provide a characterization for FMFE that suggests a simple and efficient algorithm for its computation. Lastly, we derive a *closed-form* characterization of the strategies under the FMFE, and of the resulting competitive landscape in the case of homogeneous bidders, i.e., when all advertisers have the same budget and campaign length. Such closed forms for equilibria of dynamic games are remarkably rare and one may significantly leverage such a result when studying the publisher's problem.

Fourthly, we study the auction design problem for a publisher that maximizes expected profits by leveraging our characterization of the outcome of the interactions among advertisers. In particular, we analyze the impact of three different decision variables on the publisher's profits when running second-price auctions: the reserve price, the allocation of impressions, and the disclosure of information. When solving her optimization problem the publisher trades-off the revenues extracted from the auction with the opportunity cost of selling the impressions through an alternative channel. In addition, she needs to consider that changing the auction parameters may change the FMFE strategies played by advertisers. We formulate the publisher's problem as an Mathematical Program with Equilibrium Constraints (MPEC), and numerically analyze the impact of the publisher's decisions on the the advertisers' equilibrium outcome under different scenarios.

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

When the publisher posts an impression in the exchange she can decide which user information to disclose to the advertisers. On the one hand, more information enables advertisers to improve targeting, which results in higher bids conditional on participating in an auction. On the other hand, as more information is provided, fewer advertisers match with each user, resulting in *thinner* markets, which could decrease the publisher's profit. We apply our framework to a stylized model for information disclosure, and show that if the publisher reacts to thinner markets *by setting an appropriate reserve price*, then disclosing more information will always increase the publisher's profits; an appropriately set reserve price allows to extract surplus even in thin markets.

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

1.5 Advertiser Competition in AdX: Approximation Results

The contributions of Chapter 4 are as follows.

Firstly, we show that the FMFE approximates the rational behavior of bidders in large markers by considering the simplified model of *synchronous campaigns*, that is, when all campaigns start at the same time and finish simultaneously. This model captures, for example, the case when advertisers have periodic (daily or weekly) budgets. In this setting 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.

The result is proven by considering a sequence of markets with increasing size. On the demand side, the number of advertisers and their budgets are allowed to increase. On the supply side, the number of impressions is increased so that the expected number of auctions a bidder participates in grows at the same rate as her budget, while the expected number of bidders in each auction remains constant. We impose the additional assumption that the number of advertisers in the market grows slower (in the little-o sense) than the number of auctions an advertiser participates in. This is an appealing regime that applies to most current markets.

The proof is based on the fundamental observation that advertisers bid exactly as prescribed by the FMFE while they have budgets remaining. In addition, the large number of competitors limits the impact of a single advertiser in the market, in the sense that competitors run out of budget –in expectation– close to the end of their campaigns no matter which strategy that advertiser implements. The combination of these two remarks yields that the competitive landscape coincides with that predicted by the FMFE for most of the horizon.

Building on the previous result we bound the performance of an arbitrary strategy by that of a strategy with the benefit of hindsight (which has complete knowledge of the future realizations of bids and values). This is akin to what is typically done in Revenue Management (RM) settings (see, e.g., Talluri and van Ryzin (1998)). The main exception, however, is that here the competitive environment is endogenous and determined through the FMFE consistency requirement. We use this bound to show the approximation result.

Secondly, we prove the approximation result in the general model with *asynchronous campaigns*, that is, when advertisers arrive to the market at random points in time and campaigns overlap. The key simplifications in the FMFE are that (i) all competing advertisers present in the market are allowed to bid for the purpose of determining the competing landscape; and (ii) the actions of an advertiser do not affect the competitors in the market, and competitors' states and the number of matching bidders in successive auctions are independent. The challenges introduced by the asynchronous model are that (i) advertisers may run out of budget during their campaigns and thus the competitive landscape differs from the one predicted by the FMFE, and (ii) competitors' states and the number of matching bidders in successive auctions are not necessarily independent. The complexity of this model precludes the possibility of applying traditional RM techniques, and thus motivated we develop a novel framework based on more elaborate mean field techniques.

The first step of the proof consists of addressing (i) above. 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, 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 are essentially independent, and that 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.

The main limitation of this result is that the scaling is more restrictive than in the synchronous case. Our proof holds under the assumption that the number of advertisers in the market grows exponentially faster than the number of auctions. We conjecture, however, that the family of scalings under which our approximation result is valid is broader.

1.6 Related Literature

Our works draws on four streams of literature, namely, that of Display Advertising with Ad Exchange, Revenue Management, Online Allocation, and Game Theory. Rather than attempting to exhaustively survey the literature on each area, we focus on the work more closely related to ours.

Display Advertising with Ad Exchange. Our work contributes to the growing literature on display advertising, and in particular on that with Ad Exchanges. Muthukrishnan (2009) provides a comprehensive overview of Ad Exchanges.

From the publisher's perspective, there has been recent work on display ad allocation with both contract-based advertisers and spot market advertisers. Ghosh, Papineni, McAfee and Vassilvitskii (2009) focus on "fair" representative bidding strategies in which the publisher bids on behalf of the contract-based advertisers competing with the spot market bidders. This line of work is mainly concerned with computing such fair representative bidding strategies for contract-based advertisers. Yang et al. (2010) studied the problem faced by the publisher of allocating between the two markets using multi-objective programming. As in our work, they consider different objectives for the publisher, such as, minimizing the penalty of under-delivery, maximizing the revenue from the spot market, and the representativeness of the allocation. However, they employ a deterministic model with no uncertainty in which future inventory and contracts are nodes in a bipartite graph. Alaei et al. (2009) proposed an utility model that accounts for two types of advertisers: one oriented towards campaigns and seeking to create brand equity, and the other oriented towards the spot market and seeking to transform impressions to sales. Here impressions are commodities which can be assigned interchangeably to any advertisers. In this setting they look for offline and online algorithms aiming to maximize the utility of their contracts of the allocation. These three papers, however, take the actions of the advertisers as exogenous in the auction design.

Chen (2011) considers the case when the publisher runs the exchange, and employing a mechanism design approach he characterizes, through dynamic programming, the optimal dynamic auction for the spot market. In this model both bids from the spot market and the total number of impressions are stochastic. We focus, instead, on combined yield optimization and present a model and an algorithm taking into account any trade-off between quality delivered to reservation ads and revenue from the spot market. Additionally, in this work the publisher faces short-lived advertisers and budget constraints are not considered. Vulcano et al. (2002) considers a related problem in the context of a single-leg revenue management problem. There, a seller auctions a limited stock to a sequence of buyers separated into different time periods. In contrast to our work, bidders are independent and compete directly against each other within a period, and indirectly with buyers in other periods.

Regarding the disclosure of information, Levin and Milgrom (2010) discuss how targeting can increase efficiency by improving the match between users and advertisers, but at the same time reduce publisher's revenues by creating thinner market. With this motivation, Celis et al. (2011) introduce a new randomized auction mechanism that experimentally performs better than an optimized second-price auction in markets that become thin due to targeting. They present a truthful auction mechanism to handle environments with few bidders with irregular distribution of values in Ad Exchanges. Their mechanism is an extension of a second-price auction with a reserve price, and it is shown to be nearly-optimal in this setting. This work, however, considers a one shot auction and does not take into account the dynamics introduced by budget constraints.

From the advertiser's perspective, Ghosh, Rubinstein, Vassilvitskii and Zinkevich (2009) study the design of a bidding agent that implements a campaign in the presence of an exogenous market.

Revenue Management. Another stream of relevant work is that of RM. Even though RM is typically applied to airlines, car rentals, hotels and retailing (Talluri and van Ryzin, 2004), our problem formulation and analysis is inspired by RM techniques.

In Chapter 2 we propose a modified bid-price control to tackle the publisher's multi-objective problem of allocating impressions between two channels. Bid-price controls are popular method for controlling the sale of inventory in revenue management applications. These were originally introduced by Simpson (1989), and thoroughly analyzed by Talluri and van Ryzin (1998). In this setting, a bid-price control sets a threshold or bid price for each advertiser, which may be interpreted as the opportunity cost of assigning one additional impression to the advertiser. This approach is standard in the context of revenue maximization, e.g. the stochastic knapsack problem by Levi and Radovanovic (2010). From this perspective, our contribution is the inclusion of a spot market, the exchange, as an new sales channel.

In terms of multi-objective optimization in revenue management, Levin et al. (2008) employ a weighted sum approach to determine, in a dynamic pricing setting, the Pareto efficient frontier between revenue and the probability that total revenue falls below a minimum acceptable level. Phillips (2012) uses a similar approach to determine the efficient frontier between any two goals that are linear in load (such as revenue and profits) in a single-leg revenue management problem.

There is some body of literature on display advertising from a revenue management angle that focuses exclusively on guaranteed contracts (see, e.g., Araman and Fridgeirsdottir (2011), Fridgeirsdottir and Najafi (2010), Roels and Fridgeirsdottir (2009), and Turner (2012)), as well as sponsored search advertising (see, e.g, Nazerzadeh et al. (2009)). These papers, however, do not consider the spot market. In the related area of TV broadcasting, Araman and Popescu (2010) study the allocation of advertising space between forward contracts and the spot market when the planner faces supply uncertainty.

Our work in Chapter 3 on competition between advertisers in AdX relates from both methodological and approach standpoints to some stream of work in revenue management. The single agent fluid approximation we use and some of the intuition underlying it is related to that of, e.g., Gallego and van Ryzin (1994). Building on the latter, Gallego and Hu (2011) focusing on price competition, use a notion of fluid, or open-loop, equilibrium. Other papers studying on dynamic games in revenue management (all focusing on price competition) include Farias et al. (2011), de Albéniz and Talluri (2011), and Dudey (1992).

Online Allocation. Our work is closely related to the *Display Ads Allocation (DA)* problem from the Computer Science literature, in which the publisher must assign online impressions to an inventory of ads, optimizing efficiency or revenue of the allocation while respecting pre-specified contracts.

In the DA problem, advertisers demand a maximum number of eligible impressions, and the publisher must allocate impressions that arrive online to them. Each impression has a potentially different value for every advertiser. The goal of the publisher is to assign each impression to one advertiser maximizing the value of all the assigned impressions. The adversarial online DA problem was considered in Feldman et al. (2009), which showed that the problem is inapproximable without exploiting *free disposal*; using this property (that advertisers are at worst indifferent to receiving more impressions than required by their contract), a simple greedy algorithm is $\frac{1}{2}$-competitive, which is optimal. When the demand of each advertiser is large, a $(1-\frac{1}{e})$-competitive algorithm exists (Feldman et al., 2009), and it is tight. The stochastic model of the DA problem is more related to our problem. Following a training-based dual algorithm by Devenur and Hayes (2009), training-based $(1-\epsilon)$-competitive algorithms have been developed for the DA problem and its generalization to various packing linear programs (Feldman et al., 2010; Vee et al., 2010; Agrawal et al., 2009).

Our work differs from all the above in three main aspects: (i) We study both the parametric and non-parametric models, and compare their effectiveness in terms of the size of the sample sizes—both analytically for various distributions and experimentally on real data sets. (ii) Instead of using the framework of competitive analysis and comparing the solution with the optimum solution in hindsight, we compare the performance of our algorithm with the optimal online policy, and present a rate of convergence bound under this model. This is akin to regret bounds found in online Machine Learning; and (iii) None of the above work considers the simultaneous allocation of reservation ads and ads from AdX. In particular, these previous works do not consider the trade-off between the revenue from a spot market based on real-time bidding and the efficiency of reservation-based allocation.

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

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

1.7 Conclusions

In this thesis we develop a novel framework to address some fundamental questions on the design and operation of Ad Exchanges. In the first part of this thesis, we bring to bear techniques of Revenue Management and present an approach to help publishers determine when and how to access AdX to complement their contract sales of impressions. In particular, we model the publishers' problem as a stochastic control program and derive an asymptotically optimal policy with a simple structure: a bid-price control extended with a pricing function for the exchange. We show using data from real inventory that there are considerable advantages for the publishers from jointly optimization over both channels. Publishers may increase their revenue streams without giving away the quality of service of their reservations contracts, which still represents a significant portion of their advertising yield. We also hope our insights here will help understand ad allocation problems more deeply.

In the second part of this thesis, we present a new model for Ad Exchange in which the competitive landscape is endogenous, that is, advertisers engage in a dynamic game, whose equilibrium outcome determines the market characteristics. Overall, our results provide sharp insights on the design of Ad Exchange markets and on the publisher's profit maximization problem. Notably, the structure of the optimal auction design decisions are simple and intuitive, and may have implications on the design of such auctions in practice. At the same time, this work contributes to various streams of literature. By accounting for advertisers' budget constraints and the resulting inter-temporal dependencies and dynamic bidding strategies they induce, we contribute to the internet advertising literature in particular, and more generally, to the literature on auction design in dynamic settings. In fact, we expect that FMFE may have additional applications beyond the one presented in this paper. This work also contributes to the revenue management literature; the publisher's optimization of impression allocation and selling mechanism are core revenue management problems and so is, in some way, the advertisers's scarce resource allocation problem.

On the theoretical arena, we provide two complementary results that confirm that FMFE provides a good approximation to the rational behavior of agents in large markets. First, we study the case of synchronous campaigns, that is, when all campaigns start at the same time and finish simultaneously. In this setting we are able to show our result under a very appealing regime that applies to most current markets. Our main contribution involves extending the asymptotic optimality of RM fluid-based policies to an endogenous environment. Second, we study the more general case of asynchronous campaigns, that is, when advertisers arrive to the market at random points in time and campaigns overlap. The complexity of this model precludes the possibility of applying traditional RM techniques, and thus motivated we develop a novel framework based on more elaborate mean field techniques. The main limitation of our result is that the scaling is more restrictive than in the synchronous case. However, the techniques developed in this work contribute to the growing literature on mean field theory by providing the first result that simultaneously scales the number of players with the number of opportunities to handle the fluid approximation.

Overall, our results provide a new approach to study Ad Exchange markets and the publishers' decisions. The techniques developed build on two fairly distinct streams of literature, revenue management and mean field models and are likely to have additional applications. The sharp results regarding the publisher's decisions could inform how these markets are designed in practice.

At the same time, our framework opens up the door to study a range of other relevant issues in this space. For example, one interesting avenue for future work may be to study the impact of Ad networks, that aggregate bids from different advertisers and bid on their behalf, on the resulting competitive landscape and auction design decisions. Similarly, another interesting direction to pursue is to incorporate common advertisers' values and analyze the impact of cherry-picking and adverse selection. Finally, our framework and its potential extensions can provide a possible structural model for bidding behavior in exchanges, and open the door to pursue an econometric study using transactional data in exchanges.

Chapter 2

Yield Optimization of Guaranteed Contracts with AdX

The material presented in this chapter is based on the working paper Balseiro et al. (2011) co-authored with Jon Feldman, Vahab Mirrokni, and S. Muthukrishnan.

In this chapter we study the problem faced by a publisher who must trade off, in real-time, the short-term revenue from an Ad Exchange with the long-term benefits of delivering good quality spots to the reservation ads. In Section 2.1 we present a stochastic model that captures the dynamics of users' arrivals to the publisher web site, the targeting criteria and constraints imposed by the guaranteed contracts, and the behavior of real-time bidders in AdX. Bringing to bear techniques of Revenue Management, we model the publisher's problem as a combination of a capacity allocation problem to handle the guaranteed contracts together with a dynamic pricing problem to handle the reserve price optimization in the exchange. Here the publisher needs to decide whether to post an arriving impression in the exchange with a proper reserve price, or assign it to the best matching reservation. In Section 2.2.1 we formalize the publisher's combined optimization problem as a stochastic dynamic programming problem and discuss the structure of the optimal policy. Because of the "curse of dimensionality" the dynamic program cannot be solved efficiently in most real-world problems, and instead we propose in Section 2.2.2 a deterministic or fluid approximation. As a result, we derive a bid-price policy extended with a pricing function to take into account for AdX. Additionally, we show that the proposed policy becomes asymptotically close to the optimal one as the capacity of the contracts and the number of impressions in the planning horizon is scaled up. In Section 2.3 we give a parametric model based on our observation of real data, which takes into consideration that advertisers demand for particular user types in their contracts, and estimate the primitives of our model using actual publisher inventory. In Section 2.4 we present three sets of numerical experiments that were conducted with the objective to (i) analyze the impact of introducing an AdX on the publisher's yield, and (ii) compare the performance of our policy with those of two popular heuristics, (iii) evaluate an efficient method to compute the parameters of the policy in large-scale instances. In Section 2.5 we consider a number of extensions of the model and policy, and Section 2.6 concludes with some final remarks.

2.1 Model Description

Consider a publisher displaying ads in a web page. The web page has a single slot for display ads, and each user is shown at most one impression per page. The publisher has signed contracts with a A advertisers under which he agrees to deliver exactly C_a impressions to advertiser $a \in \mathcal{A}$, where we denote by $\mathcal{A} = \{1, \ldots, A\}$ the set of advertisers. Neither over-delivery nor under-delivery is allowed.

Even though the number of users visiting a web page is uncertain, publishers usually have fairly good estimates of the total number of expected users that arrive in a given horizon. In this model we index time based on the arrival of each user, and assume that the total number of users is fixed and equal to N (random number of users can be accommodated in our model by considering dummy arrivals). Each user is identified by a vector of attributes $U_n \in \mathcal{U}$, where $\mathcal{U}$ is some finite subset of $\mathbb{R}^M$, and depending on the attributes, the impression may be more or less attractive to different advertisers. The vector of attributes contains information that is relevant to the advertisers' targeting such as (i) the web address or URL, (ii) keywords related to the content of the web-page; (iii) the dimension and position of the slot in the page; (iv) user's geographical information, that is, where is the user located; (v) user's demographics, such as education level, gender, age or income; (vi) user's device and operating system, and (vii) cookie-based behavioral information, which allows bidders to track the user's past activity in the web. We assume that the vectors of attributes $\{U_n\}_{n=1,\ldots,N}$ are random, independent and identically distributed.

Based on the vector of attributes for the impression, the publisher determines a vector of placement qualities $Q_n = \{Q_{n,a}\}_{a \in \mathcal{A}}$, where $Q_{n,a}$ is the predicted quality advertiser a would perceive if the impression is assigned to her. Qualities lie in some compact space $\Omega \subseteq \mathbb{R}^A$. A typical measure of placement quality is the estimated probability that the user clicks on each ad. In practice, such measure of quality is learned by performing, for example, a logistic regression based on the vector attributes as explanatory variables. Here we abstract from the learning problem and assume that the qualities are deterministically determined from the impression attributes. Hence, the vectors of placement qualities $\{Q_n\}_{n=1,\ldots,N}$ are random, and independent and identically distributed across impressions. We do allow, however, for qualities to be jointly distributed across advertisers. This captures the fact that advertisers might have similar target criteria, and hence the qualities perceived might be correlated. We do not impose any further restrictions on the qualities, other than bounded support. Notice that the publisher observes the realization of the placement quality before showing the ad.

We assume that the number of arriving impressions suffices to satisfy the contracts, or equivalently $\sum_{a\in\mathcal{A}} \rho_a \leq 1$, where $\rho_a = \frac{C_a}{N}$ denotes the capacity-to-impression ratio of an advertiser. An assumption of this general model is that any user can be potentially assigned to any advertiser. In practice each advertiser may be interested in a particular group of user types. It is important to note that this is not a limitation of our results, but rather a modeling choice; in §2.3 we show how to handle targeting criteria by setting $Q_{n,a} = -\tau_a$ for impressions not matching the targeting criteria of an advertiser. This can also be interpreted as forcing the publisher to pay a good will penalty τ_a to the advertisers each time an undesired impression is incorrectly assigned.

Arriving impressions may either be assigned to the advertisers, discarded or auctioned in the Ad Exchange (AdX) for profit. In a general AdX (Muthukrishnan, 2009), the publisher contacts the exchange with a minimum price she is willing to take for the slot. Additionally, the publisher may submit some partial information of the user visiting the website. User information allows advertisers in the exchange to target more effectively, which may in turn result in higher bids (Balseiro et al., 2012*b*). Internally the exchange contacts different ad networks, and in turn they return bids for the slot. The exchange determines the winning bid among those that exceed the reserve price via an auction, and returns a payment to the publisher. In this case we say that the impressions is *accepted*, and the publisher is contractually obligated to display the winning impression. In the case that no bid attains the reserve price, no payment is made and the impression is *rejected*. We present the formal model of the exchange in §2.1.2. The entire operation above is executed before the page is rendered in the user's screen. Thus, in the event that the impression is rejected by the exchange, the publisher may still be able to assign it to some advertiser. Figure 2.1 summarizes the decisions involved.

For notational simplicity we extend the set of advertisers to $\mathcal{A}_0 = \{0\} \cup \mathcal{A}$ by including an outside option 0 that represents discarding an impression. We set the quality of the outside option identically to zero, i.e. $Q_{n,0} = 0$ for all impressions $n = 1, \ldots, N$. In the following, the terms discarding an impression or assigning it to advertiser 0 are used interchangeably. We set $\rho_0 = 1 - \sum_{a\in\mathcal{A}} \rho_a$ to be the fraction of impressions that are not assigned to any advertiser. To wit, a fraction of the ρ_0 impressions will be assigned to the winning impression of AdX, and the remainder effectively discarded.

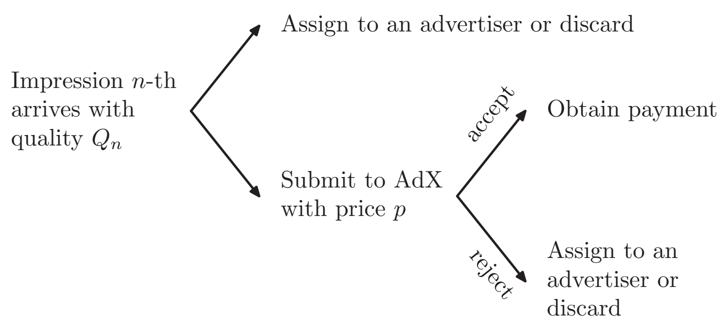

Figure 2.1: Publisher's decision tree for a new impression.

Note. All proofs are presented in the main appendix.

2.1.1 Objective

The publisher's problem is to maximize the overall placement quality of the impressions assigned to the advertisers together with the total revenue obtained with AdX, while complying with the contractual obligations. We attack the multi-objective problem by taking a weighted sum of both objectives. The publisher has at her disposal a parameter $\gamma \geq 0$, which allows her to trade-off between these conflicting objectives. The aggregated objective is given by

$$\text{yield} = \text{revenue(AdX)} + \gamma \cdot \text{quality(advertisers)}.$$

A suitable large γ would give priority to assigning high quality impressions to the advertisers; while a small γ would prioritize the revenue from AdX (the publisher may set different values of the parameter for each advertiser). Without loss of generality, we set $\gamma = 1$ for the remainder of this paper, except when noted otherwise.

Two observations are in order. First, by adjusting the tradeoff parameter γ the publisher is able to construct the Pareto efficient frontier of attainable revenue from AdX and quality for the advertisers. In §2.4.1 we study experimentally the impact of the choice of γ on both objectives, and determine the Pareto frontier for real publisher data. Second, the publisher might alternatively impose that the overall quality of the impressions assigned to the advertiser is greater than some threshold, and then maximize the total revenue obtained from AdX; this may have a more natural interpretation for some publishers, and would be simpler than having to set γ. We can model this simply by interpreting γ as the Lagrange multiplier of the quality of service constraint, and our problem as the Lagrange relaxation of the constrained program. In §2.5.3 we analyze the implications of this formulation.

2.1.2 AdX Model with User Information

The publisher submits an impression to AdX with the minimum price it is willing to take, denoted by $p \geq 0$. The impression is accepted if there is a bid of value p or more. We denote by B the winning bid random variable, which may be correlated with the user information $u \in \mathcal{U}$ disclosed by the publisher. In practice, publishers maintain different estimates of the distribution of the maximum bid in the exchange as a function of the attributes of the impression (we discuss this further in §2.2.5). In the following we assume that bids are independent across impression, and identically distributed according to a c.d.f. $F(\cdot; u)$. Hence, when the publisher discloses some information u the impression is accepted with probability $1 - F(p; u) = \bar{F}(p; u)$. For ease of exposition, in this first model the publisher is paid the minimum price p when the impression is accepted. In §2.5.1 we drop this assumption and consider a more general second-price auction.

Suppose the publisher has computed an *opportunity cost* c for selling this inventory in the exchange; that is, the publisher stands to gain c if the impression is given to a reservation advertiser. Given opportunity cost $c \geq 0$ the publisher picks the price that maximizes its expected revenue. Hence, the publisher solves the optimization problem $R(c; u) = \max_{p \geq 0} \bar{F}(p; u)p + F(p; u)c$. Changing variables, we can define $r(s; u) = s\bar{F}^{-1}(s; u)$ to be the expected revenue under acceptance probability s and user information u, and rewrite this as[1]

$$R(c; u) = \max_{s \in [0,1]} r(s; u) + (1 - s)c. \tag{2.1}$$

Also, let $s^*(c; u)$ be the least maximizer of (2.1), and $p^*(c; u) = \bar{F}^{-1}(s^*(c; u); u)$ be the price that verifies the maximum.

Assumption 2.1. *The expected revenue $r(s; u)$ is continuous in s, concave in s, non-negative, bounded, and satisfies $\lim_{s \to 0} r(s; u) = 0$ for every user information $u \in \mathcal{U}$. We call a function that satisfies all of the assumptions above a* regular revenue function.

These assumptions are common in RM literature (see, e.g., Gallego and van Ryzin (1994)). A sufficient condition for the concavity of the revenue is that B has increasing generalized failure rates (Lariviere, 2006). Regularity implies, among other things, the existence of a *null price* $p_\infty(u)$ such that $\lim_{p \to p_\infty} \bar{F}(p; u)p = 0$. Additionally, it allows us to characterize the value function $R(c; u)$. In §2.5.1 we show that the revenue function remains regular in the presence of multiple bidders in AdX by considering the joint density of the highest and second-highest bids. Thus, all our results hold in this case too.

[1] We define the generalized inverse distribution function as $\bar{F}^{-1}(s; u) = \inf\{p \geq 0 \;:\; \bar{F}(p; u) \leq s\}$ to take into account the case where the distribution is not absolutely continuous.

Proposition 2.1. *Suppose that r is regular revenue function. Then, for fixed user information u we have that $R(c;u)$ is non-decreasing in c, convex in c, continuous in c, and $R(c;u) \geq c$. Additionally, $R(c;u) - c$ is non-increasing in c, $s^*(c;u)$ is non-increasing in c, and $p^*(c;u)$ is non-decreasing in c.*

An important consequence of above is that the maximum revenue expected from submitting an impression to AdX is always greater than the opportunity cost. This should not be surprising, since the publisher can pick a price high enough to compensate for the revenue loss of not assigning the impression. Hence, assigning an impression directly to an advertiser (rather than first testing the exchange) is never the right decision, and so in Figure 2.1 the upper branch is never taken.

2.1.3 Discussion of the Assumptions

In the absence of a fixed cost, the optimal policy tests the exchange before assigning the impression to the contracts. Such result depends strongly on the publisher's ability to dynamically adjust the reserve to take into account the opportunity cost of "losing" an impression of high quality to the exchange. If this is not the case, the publisher only tests the exchange when the expected revenue from AdX exceeds the contracts' opportunity cost. We discuss this further in §2.4.2.

Publishers usually receive a revenue share of all impressions sold in the exchange. Under such a *revenue sharing scheme* the exchange keeps a fraction α of the bidder's payment p, and the publisher receives the amount $(1-\alpha)p$ for the impression. Our model can accommodate this scheme by noticing that the publisher only needs to increase the impression's opportunity cost to $c/(1-\alpha)$. It is straightforward to show that the publisher's AdX value function is now given by $R_\alpha(c;u) = (1-\alpha)R(c/(1-\alpha);u)$ and the optimal price is $p^*_\alpha(c;u) = p^*(c/(1-\alpha);u)$, where R and p^* denote the value function and optimal price in the case of no revenue sharing, respectively.

Publishers typically are not charged a fixed cost each time they access AdX. However, a publisher may still assign the exchange a fixed cost $\ell > 0$ to take into account, for example, the negative effect of latency in the user experience or the opportunity cost of capacity when bandwidth is limited. In this case the publisher would access the exchange only if the marginal expected contribution from the exchange exceeds the fixed cost, that is, $R(c;u) - c \geq \ell$. In view of Proposition 2.1, the marginal expected contribution $R(c;u) - c$ is non-increasing in c and one can show that the publisher accesses the exchange only if the opportunity cost is less or equal to the threshold $c^*(\ell;u) = \sup\{c : R(c;u) - c \geq \ell\}$. When the opportunity cost is higher than the threshold the publisher stands to gain little from accessing the exchange, and in the presence of the fixed cost, he decides to bypass the spot market.

Two final assumptions of our model, which are pervasive in the RM literature, are the stationarity and independence of the user arrival process. The former assumption is not entirely realistic because traffic patterns typically vary through the day. For example, a newspaper may observe a spike of traffic in the mornings due to office users and another in the night from home users. Our model can accommodate non-stationary traffic patterns in a straightforward way by allowing the distributions of placement qualities and bids to be time-dependent as done in Talluri and van Ryzin (1998). The latter assumption is not very restrictive because unique user visiting the website arrive essentially at random, so inter-temporal correlation should be expected to be weak.

2.2 Problem Formulation

In this section we start by formulating an optimal control policy for yield maximization based on dynamic programming (DP), where the state of the system is represented by the number of impressions yet to arrive, and a vector of the number of impressions needed to comply with each advertiser's contract. Unfortunately, the state space of the DP has size $O(N^{A+1})$, and in most real-world problems the number of impressions in a single horizon can be in the order of millions. So the DP is not efficiently solvable. We give, instead, an approximation in which stochastic quantities are replaced by their expected values, and are assumed to be continuous. Such "deterministic approximation problems (DAP)" are popular in RM (see, e.g., Talluri and van Ryzin (1998)). In our setting, the approximation we make is to enforce contracts to be satisfied only in expectation. We formulate the problem based on this assumption and obtain an infinite-dimensional program. This DAP is solved by considering its dual problem, which turns out to be a more tractable finite-dimensional convex program. Finally, we wrap a full stochastic policy around it (one that always meets the contracts, not just in expectation), and show that this policy is asymptotically optimal when the number of impressions and capacity are scaled up proportionally.

2.2.1 Dynamic Programming Formulation

Let (m, X) be the state of the system, where we denote by m the total number of impressions remaining to arrive, and by $X = \{x_a\}_{a \in \mathcal{A}}$ the number of impressions needed to comply with each advertiser's contract. Let the value function, denoted by $J_m(X)$, be defined as the optimal expected yield obtainable under state (m, X). Using the fact that is optimal to first test the exchange, we obtain the following Bellman equation

$$J_m(X) = \mathbb{E}_{U_n}\left[\max_{p\geq 0}\left\{\bar{F}(p;U_n)(p+J_{m-1}(X)) + (1-\bar{F}(p;U_n))\max_{a\in\mathcal{A}_0}\{Q_{n,a}+J_{m-1}(X-\mathbf{1}_a)\}\right\}\right]$$
$$= J_{m-1}(X) + \mathbb{E}_{U_n}\left[R\left(\max_{a\in\mathcal{A}_0}\{Q_{n,a}-\Delta_a J_{m-1}(X)\};U_n\right)\right], \tag{2.2}$$

where we defined $\mathbf{1}_a$ as a vector with a one in entry a and zero elsewhere, $\mathbf{1}_0 = 0$, and $\Delta_a J_m(X) = J_m(X) - J_m(x - \mathbf{1}_a)$ as the expected marginal yield of one extra impression for advertiser a. In (2.2) the objective accounts for the yield obtained from attempting to send the impression to AdX. The first term in the maximand accounts for the expected revenue from the exchange, while the second term accounts for the decision of assigning the impression to a reservation or discarding it (when $a = 0$). In (2.2) we used the fact that assigning an impression directly to an advertiser is never the right decision (except in boundary conditions, see below). The publisher, however, may choose to discard impressions with low quality after being rejected by AdX.

Our objective is to compute $J^*_N = J_N(C)$. Let M be an upper-bound on the expected yield.[2] The boundary conditions are

$$J_m(x) = -M, \qquad \forall X \text{ s.t. } x_a < 0 \text{ for some } a \in \mathcal{A},$$
$$J_m(x) = -M, \qquad \forall m < \sum_{a\in\mathcal{A}} x_a.$$

Recall that when the contract with an advertiser is fulfilled, no extra yield is obtained from assigning to her more impressions. This is the case of the first boundary condition, which guarantees that advertisers whose contract is fulfilled are excluded from the assignment. In particular, when $X = 0$ all remaining impressions are sent to AdX with the yield maximum price $p^*(0)$ when $x = 0$. The second boundary condition guarantees that the contracts with the advertisers are always fulfilled. When $\sum_{a\in\mathcal{A}} x_a = m$ AdX must be bypassed, and impressions should be assigned directly to the advertisers. The optimal policy is described in Policy 1.

In the above policy, when the impression is submitted to AdX, the optimal price ponders an opportunity cost of $Q_{n,a^*_n} - \Delta_{a^*_n} J_{n-1}(X)$. This opportunity cost, when positive, is just the value of the impression adjusted by the loss of potential yield from assigning the impression right now. Note that the two boundary conditions are implicit in the optimal policy. This guarantees that the policy complies with the contracts. It is routine to check that the value function $J_n(X)$ is finite for all feasible states and that Policy 1 is optimal for the dynamic program in (2.2). It is worth noting that in order to implement the optimal policy one needs to pre-compute the value function, which is intractable in most real instances.

---

[2] One could set, e.g., $M \triangleq N\max\{p_\infty, \bar{Q}\}$ where $p_\infty = \max_{u\in\mathcal{U}} p_\infty(u)$ and $\bar{Q}$ is an upper-bound on the placement quality

Policy 1 Optimal dynamic programming policy.

1: Observe state (m, X), and the impression's vector of attributes U_n.
2: Determine the vector of placement qualities Q_n.
3: Let $a_n^* = \arg\max_{a \in \mathcal{A}_0} \{Q_{n,a} - \Delta_a J_{n-1}(X)\}$.
4: Submit to AdX with price $p^*\left(Q_{n,a_n^*} - \Delta_{a_n^*} J_{n-1}(X); U_n\right)$.
5: **if** impression rejected by AdX and $a_n^* \neq 0$ **then**
6:     Assign to advertiser a_n^*.
7: **end if**

2.2.2 Deterministic Approximation Problem (DAP)

We aim for an approximation in which (i) the policy is independent of the history but dependent on the realization of the vector of attributes U_n (recall that placement qualities are deterministically determined based on the attributes), (ii) capacity constraints are met in expectation, and (iii) controls are allowed to randomize. These approximations turn out to be reasonable when the number of impressions is large. When an impression arrives, the publisher controls the reserve price submitted to AdX, and the advertiser to whom the impression is assigned, if rejected by AdX. Alternatively, in this formulation we state the controls in terms of total probabilities, where each control is a function from the attribute space $\mathcal{U}$ to $[0, 1]$. Let $\vec{s} = \{s_n(\cdot)\}_{n=1,\dots,N}$ and $\vec{\imath} = \{i_n(\cdot)\}_{n=1,\dots,N}$ be vectors of functions from $\mathcal{U}$ to $\mathbb{R}$, such that when the n^{th} impression arrives with attributes u the impression is accepted by AdX with probability $s_n(u)$, and with probability $i_{n,a}(u)$ it is assigned to advertiser a. From this controls, one can determine the conditional probability of an impression being assigned to advertiser a given that it has been rejected by AdX by $I_{n,a}(u) = i_{n,a}(u)/(1 - s_n(u))$, and the reserve price to be posted in the exchange by $\bar{F}^{-1}(s_n(u); u)$. When it is clear from the context, we simplify notation by eliminating the dependence on u from the controls.

A control is feasible for the DAP if (i) it satisfies the contractual constraint in expectation, (ii) the individual controls are non-negative, and (iii) for every realization of the qualities the probabilities sum up to at most one. We denote by $\mathcal{P}$ the set of controls that satisfy the latter two conditions. That is, $\mathcal{P} = \{(s, i) \in \mathcal{U} \to [0, 1]^{A+1} : \sum_{a \in \mathcal{A}} i_a + s \leq 1, s \geq 0, i \geq 0\}$. The objective of the DAP is to find a sequence of real-valued measurable functions that maximize the expected yield, or equivalently

$$
\begin{aligned}
J_N^D = \max_{(s_n,i_n)\in\mathcal{P}} & \sum_{n=1}^{N} \mathbb{E}\left[r(s_n;U_n) + \sum_{a\in\mathcal{A}} i_{n,a} Q_{n,a} \right] \\
\text{s.t.} & \sum_{n=1}^{N} \mathbb{E}\left[i_{n,a} \right] = N\rho_a, \qquad \forall a\in\mathcal{A}.
\end{aligned} \tag{2.3a}
$$

The first term of the objective accounts for the revenue from AdX, while the second accounts for the quality perceived by the advertisers. Notice that in the DAP we wrote the total capacity as $N\rho_a$ instead of C_a to allow the problem to be scaled.

Alas, the problem is still hard to solve since the number of functions is linear in N. However, exploiting the regularity of the revenue function, we can show that in the optimal solution to DAP, we can drop the dependence on n in the controls. This follows from the linearity of the constraints together with the concavity of the objective. We formalize this discussion in the following proposition.

Proposition 2.2. *Suppose that the revenue function is regular. Then, there exists a time-homogenous optimal solution to the DAP, i.e. where $s_n(\cdot) = s(\cdot)$ for all $n = 1,\ldots,N$ and $i_n(\cdot) = i(\cdot)$ for all $n = 1,\ldots,N$.*

The previous proposition allows us scale the problem so that $N = 1$, and consider the maximum expected revenue of one impression, denoted by J_1^D. The total revenue for the whole time horizon is then $J_N^D = N J_1^D$. In order to compute the DAP's optimal solution, we consider its dual problem, which we informally derive next.

Derivation of the Dual to DAP. To find the dual, we introduce Lagrange multipliers $v = \{v_a\}_{a\in\mathcal{A}}$ for the capacity constraints (2.3a). The Lagrangian, denoted by $\mathcal{L}(s,i;v)$ is

$$
\mathcal{L}(s,i;v) = \mathbb{E}\left[r(s;U) + \sum_{a\in\mathcal{A}} i_a Q_a - \sum_{a\in\mathcal{A}} v_a \left(i_a - \rho_a \right) \right].
$$

The dual function, denoted by $\psi(v)$, is the supremum of the Lagrangian over the set $\mathcal{P}$. Thus, we have that

$$
\begin{aligned}
\psi(v) &= \sup_{(s,i)\in\mathcal{P}} \mathcal{L}(s,i;v) \\
&= \sup_{s\ge 0} \left\{ \mathbb{E}\left[r(s;U)\right] + \sup_{i\ge 0, \sum_{a\in\mathcal{A}} i_a \le 1-s} \mathbb{E}\left[\sum_{a\in\mathcal{A}} i_a(Q_a - v_a)\right]\right\} + \sum_{a\in\mathcal{A}} v_a\rho_a \\
&= \sup_{s\ge 0} \mathbb{E}\left[r(s;U) + (1-s)\max_{a\in\mathcal{A}_0}\{Q_a - v_a\}\right] + \sum_{a\in\mathcal{A}} v_a\rho_a \\
&= \mathbb{E}\left[R\big(\max_{a\in\mathcal{A}_0}\{Q_a - v_a\};U\big)\right] + \sum_{a\in\mathcal{A}} v_a\rho_a
\end{aligned}
$$

where the first equation follows from partitioning the optimization between AdX acceptance and the assignment probability controls, the second from optimizing over the advertiser assignment controls i, and the last equation from solving the AdX variational problem. Note that R is convex and non-decreasing in its first argument and the maximum is convex w.r.t v, hence the composite function within the expectation is convex. Using the fact that expectation preserves convexity, we obtain that the objective $\psi(v)$ is convex in v.

Next, the dual problem is $\min_v \psi(v)$. When the revenue function is regular, the DAP's objective is concave and bounded from above. Moreover, the constraints of the primal problem are linear, and the feasible set $\mathcal{P}$ convex. Hence, by the Strong Duality Theorem (p.224 in Luenberger (1969)) the dual problem attains the primal objective value. So, we have that dual problem is given by the following convex stochastic problem

$$
J_1^D = \min_v \left\{ \mathbb{E}\left[R\big(\max_{a\in\mathcal{A}_0}\{Q_a - v_a\};U\big)\right] + \sum_{a\in\mathcal{A}} v_a\rho_a \right\}. \tag{2.4}
$$

Deterministic optimal control. Once the optimal dual variables v are known, the primal solution can be constructed from plugging the optimal Lagrange multipliers in $\mathcal{L}(s,i;v)$. Following the derivation of the dual, we obtain that the optimal survival probability is $s^*\left(\max_{a\in\mathcal{A}_0}\{Q_a - v_a\};U\right)$. Hence, the impression has a value of $\max_{a\in\mathcal{A}_0}\{Q_a - v_a\}$ for the publisher, and she picks the reserve price that maximizes her revenue given that value. From the optimization over the assignment controls, we see that an impression is assigned to an advertiser a only if she maximizes the *contract adjusted quality* $Q_a - v_a$. Thus the dual variables v_a act as the *bid-prices* of the guaranteed contracts. Additionally, the impression can be discarded only if the maximum is not verified by an advertiser (i.e. all contract adjusted qualities are non-positive).

Notice that optimizing the Lagrangian states that the impression should be assigned to an advertiser maximizing the contract adjusted quality, but does not specify how the impression should be assigned when –multiple– advertisers attain the maximum. In the case when the probability of a tie occurring is zero, the problem admits a simple solution: assign the impression to the unique maximizer of $Q_a - v_a$. We formalize this discussion in the the following theorem.

Theorem 2.1. *Suppose that the revenue function is regular, and there is zero probability of a tie occurring, i.e.* $\mathbb{P}\{Q_a - v_a = Q'_a - v'_a\} = 0$ *for all distinct* $a, a' \in \mathcal{A}_0$*. Then, the optimal controls for the DAP are* $s(U) = s^*\left(\max_{a \in \mathcal{A}_0}\{Q_a - v_a\}; U\right)$*, and* $I_a(U) = \mathbf{1}\{Q_a - v_a > Q_{a'} - v_{a'}\ \forall a' \in \mathcal{A}_0\}$*, that is, the impression is assigned to the unique advertiser maximizing the contract adjusted quality. Furthermore, the optimal dual variables solve the equations*

$$\mathbb{E}\left[\Big(1 - s^*(Q_a - v_a; U)\Big)\mathbf{1}\left\{Q_a - v_a > Q_{a'} - v_{a'}\ \forall a' \in \mathcal{A}_0\right\}\right] = \rho_a, \qquad \forall a \in \mathcal{A}.$$

2.2.3 Our Stochastic Policy

The solution of the DAP suggests a policy for the stochastic control problem, but we must deal with two technical issues: (i) when more than one advertiser maximizes $Q_a - v_a$ we need to decide how to break the tie, and (ii) we are only guaranteed to meet the contracts in expectation, whereas we must meet them exactly. We defer the first issue until §2.5.4, where we give an algorithm for generalizing the controls to the case where ties are possible.

We propose a bid-price control extended with a pricing function for AdX given by p^*. The policy, which we denote by μ^B, is defined in Policy 2. In there we let $x_{n,a}$ be the total number of impressions left to assign to advertiser a to comply with the contract, $m = N - n$ the total number of impressions remaining to arrive, and v to be the optimal solution of (2.4).

Policy 2 Bid-Price Policy with Dynamic Pricing μ^B.

```
Observe state (m, X), the attributes U_n, and the realization Q_n.
Let A_n = {a ∈ A : x_{n,a} > 0} be the set of ads yet to be satisfied.
if Σ_{a∈A} x_{n,a} < m then
    Let a*_n = arg max_{a∈A_n∪{0}} {Q_{n,a} − v_a}.
    Submit to AdX with price p_n = p*(Q_{n,a*_n} − v_{a*_n}; U_n).
    if impression rejected by AdX and a*_n ≠ 0 then assign to advertiser a*_n, else discard.
else
    Assign to advertiser a*_n = arg max_{a∈A_n} {Q_{n,a} − v_a}.
end if
```

Notice that impressions are only assigned to advertisers with contracts that have yet to be fulfilled. When all contracts are fulfilled, impressions are sent to AdX with the revenue maximizing price $p^*(0; U_n)$. Moreover, when the total number of impressions left is equal to the number of impressions necessary to fulfill the contracts, the exchange is bypassed and all incoming impressions are directly assigned to advertisers (no impression is discarded). Hence, the stochastic policy μ^B satisfies the contracts for every sample path.

The proposed stochastic policy shares some resemblance with the optimal dynamic programming policy. The intuition is that, when the number of impressions is large, the actual state of the system becomes irrelevant because $\Delta_a J_{m-1}(x)$ is approximately constant (for states in likely trajectories), and equal to v_a. In that case both policies are equivalent.

The policy can be alternatively interpreted as the publisher bidding on behalf of the guaranteed contracts in a sequence of repeated auctions run by the exchange as in Ghosh, Papineni, McAfee and Vassilvitskii (2009). The pricing function and the bid-prices determine a reserve price or "bid" for the contracts that takes into account the value of assigning the impression to a reservation together with option value of future opportunities. In this dual interpretation the spot market lies in the spotlight while the guaranteed contracts are pushed to the background, in sharp contrast to the current practice of first aiming to fulfill the reservations and then submitting the remnant inventory to AdX. Our original interpretation is more appealing because it does not rely so heavily on the publisher always testing the exchange, which may not be optimal, for example, in the presence of a fixed-cost.

2.2.4 Asymptotic Analysis

In this section we show that the heuristic policy constructed from the DAP is asymptotically optimal for the stochastic problem when the number of impressions and capacity are scaled up proportionally. Following a similar analysis to Gallego and van Ryzin (1994) and Talluri and van Ryzin (1998), we first show that the optimal objective value of the DAP provides an upper bound on the objective value of the dynamic program. Then, we show that the upper bound is asymptotically tight when compared to the relative expected performance of our policy.

For the first result we proceed as follows. First, we formulate the problem as a stochastic control problem (SCP). Though not practical, this abstract and equivalent formulation is useful from a theoretical point of view. Second, we proceed by taking the optimal stochastic control policy, and construct a feasible solution for the DAP by taking expectations over the history. Later, we exploit the concavity of the objective and apply Jensen's inequality to show that this new solution attains a greater revenue in the DAP. In the following we denote by $J_N^* = J_N(C)$ the optimal objective value of the equivalent stochastic control problem.

Proposition 2.3. *The optimal objective value of the DAP provides an upper bound on the objective value of the optimal policy, i.e.* $J_N^* \leq J_N^D$.

Now we complete the analysis by lower bounding the yield of the stochastic policy in terms of the DAP objective. In proving that bound, we look at N^*, the first time that any advertisers contract is fulfilled or the point is reached where all arriving impressions need to be assigned to the advertisers. We refer to the time after N^* as the *left-over regime*. The first key observation in the proof is that before time N^*, the controls of the stochastic policy behave exactly as the optimal deterministic controls. The second key observation is that the expected number of impressions in the left-over regime is $O(\sqrt{N})$, and the left-over regime has a small impact on the objective.

Theorem 2.2. *Let J_N^B be the expected yield under the stochastic policy μ^B. Then,*

$$\frac{J_N^B}{J_N^*} \geq \frac{J_N^B}{J_N^D} \geq 1 - \frac{1}{\sqrt{N}} K(\rho),$$

where $K(\rho) = \sqrt{\frac{A}{A+1} \sum_{a \in \mathcal{A}_0} \frac{1-\rho_a}{\rho_a}}$, and $\rho = \{\rho_a\}_{a \in \mathcal{A}}$.

In terms of yield loss, our previous bound can be written as $J_N^* - J_N^B \leq \sqrt{N} K(\rho) J_1^D$, achieving an $O(\sqrt{N})$ loss w.r.t the optimal online policy. In particular, we may fix the capacity-to-impression ratio of each advertiser, and consider a sequence of problems in which capacity and impressions are scaled up proportionally according to ρ. Then, the yield under policy μ^B converges to the yield of the optimal online policy as N goes to infinity.

A key observation in proving the last theorem was that the number of impressions in the left-over regime is $O(\sqrt{N})$. In fact, using a Chernoff bound, we may show that the probability that the number of impressions in the left-over regime exceeds a fraction of the total impressions decays exponentially fast.

Corollary 2.1. *The probability that the number of impressions in the left-over regime exceeds a fraction $\epsilon > 0$ of the total impressions decays exponentially fast, as given by*

$$\mathbb{P}\{N - N^* \geq \epsilon N\} \leq \sum_{a \in \mathcal{A}_0} \exp(-2\epsilon^2 \rho_a N).$$

The policy described in §2.2.3 is static in the sense that it does not react to changes in supply: the dual variables v are computed at the beginning and remain fixed throughout the horizon. To address this issue, in practice, one would periodically resolve the deterministic approximation (2.4). Recently, Jasin and Kumar (2010) showed that carefully chosen periodic resolving schemes together with probabilistic allocation controls can achieve bounded yield loss w.r.t. the optimal online policy. It is worth noting that those results do not directly apply to our setting: they consider a network RM problem with discrete choice, while our model deals with jointly distributed (and possibly continuous) placement qualities and AdX. Nevertheless, by periodically resolving the DAP one should be able to obtain similar performance guarantees for the yield loss of the control.

2.2.5 Impact of User Information

Publishers typically disclose some of the impressions' attributes to the exchange, which allows advertisers to bid strategically based on this information. Similar targeting criteria across both channels can potentially introduce *positive correlation* between the placement quality of the contracts and the bids from the exchange. To obtain some managerial insights on the impact of user information, in the remaining of this section we discuss the effect of correlation on the publisher's joint allocation problem.

Positive correlation creates two interdependent effects – a *diversification loss effect* and a *price discrimination effect*. The former is a negative effect. The benefit of jointly optimizing over both channels is derived, to a great extent, from the publisher's ability to exploit the exchange to extract a higher rent from the impressions that are less attractive for the guaranteed contracts. This diversification effect is severely undermined when the targeting criteria in both channels are in perfect synchrony, and advertisers compete for the same inventory. The latter is a positive effect. Because the publisher determines the reserve price before bids are revealed, he can not fully extract the AdX surplus. However, in the presence of correlation, the publisher can exploit the impression's attributes as a covariate to predict the bids, adjust the reserve price accordingly, and extract a higher surplus from the exchange.

The total contribution of these antagonistic effects is indeterminate, and in some cases, the price discrimination effect may even dominate, resulting in yield increasing with correlation. Figure 2.2 plots the expected yield as a function of correlation for a publisher with one contract and one bidder in the exchange. Notably, when the publisher assigns a higher priority to the revenue from AdX (γ is low), yield increases with correlation. However, when the publisher prioritizes the contracts (γ is high), yield is decreasing with correlation. Additionally, to isolate the effects we plot the yield when the publisher is able to perfectly price discriminate, i.e., set the reserve price equal to the highest bid. In this case the publisher extracts all the surplus from the exchange, and the AdX variational problem is $\bar{R}(c;u) = \mathbb{E}\big[\max\{B,c\} \mid U = u\big] \geq R(c;u)$. The optimal yield under perfect pricing, denoted by $\bar{J}_1^D$, dominates the yield under imperfect pricing, and decreases with correlation since only the diversification loss effect is present here. The difference $\bar{J}_1^D - J_1^D$ can be understood as the sum of AdX surplus and the loss incurred when the highest bid falls between the reserve price and the opportunity cost, which decreases with correlation as the publisher's ability to price discriminate improves. As a result, the detrimental consequences of positive correlation on channel diversification are compensated, to some extent, by the publisher's ability to price more effectively in the exchange by exploiting the user information.

As a final remark, there are several reason why the correlation between these channels might

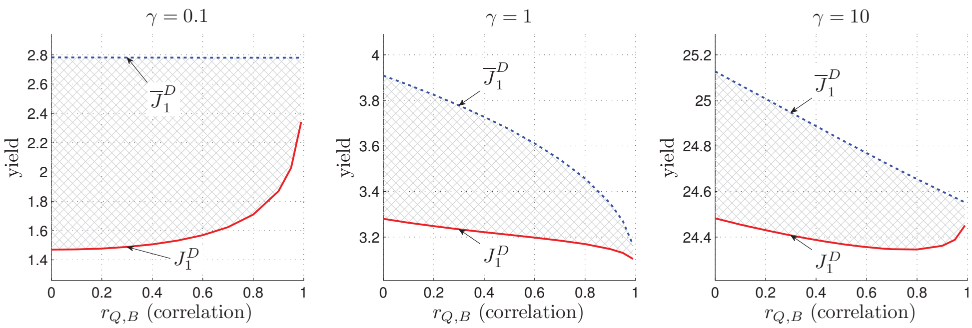

Figure 2.2: Publisher's expected yield as a function of the correlation $r_{q,b}$ for tradeoff parameters $\gamma = 0.1, 1, 10$. The publisher in consideration has signed one contract with capacity $\rho = 0.6$ and has one bidder in the exchange. The distribution of the placement quality and the bid from the exchange is a bivariate log-normal distribution with means $\mu_q = \mu_b = 1$, variances $\sigma_q^2 = \sigma_b^2 = \frac{1}{2}$, and correlation $r_{q,b}$. The solid curve denotes the actual publisher's yield J_1^D and the dashed curved denotes the publisher's under perfect price discrimination $\bar{J}_1^D$.

not be perfect. First, publishers usually do not disclose all user attributes to the spot market, thus rendering the targeting in the latter coarser. For example, registered users disclose personal information that the publisher exploits, due to privacy issues, solely to improve the targeting of guaranteed contracts. Second, advertisers in the spot market are increasingly targeting users based on *cookies*, which are private bits of information stored in the users' computers that allows them to track the past activity of the user on the web. Cookies are dropped by advertisers when users visit their own web-sites, and are only accessible to them. Thus, a strong component of the spot market bids is based on private information. In §2.3.2 we empirically explore the correlation between the bids from the exchange and the placement quality, and show that the dependence is statistically weak.

2.3 Data Model and Estimation

We have thus far assumed that any user could be potentially assigned to any advertiser. In practice, however, advertisers have specific targeting criteria. For instance, a guaranteed contract may demand females with certain age range living in New York, while other contract may demand males in California. In this section we give a parametric model based on our observation of real data, which takes into consideration that advertisers demand for particular *user types* in their contracts.

Instead of grouping user types according to their attributes, we aggregate user types that match the criteria of the same subset of advertisers. This has the advantage of reducing the space of types to a function of the number of advertisers (which is typically small in practice) rather then the number of possible types (which is potentially large). Hence, a user type is characterized by the subset of advertisers $T \subseteq \mathcal{A}$ that are interested in it. In the following, we let $\mathcal{T}$ be the support of the type distribution, and $\pi(T)$ the probability of an arriving impression being of type T. As before we assume that, across different impressions, types are independent and identically distributed. Given a particular type T, the predicted quality perceived by the advertisers within the type is modeled by the non-negative random vector $Q(T) = \{Q_a(T)\}_{a \in T}$. Thus, the ex-ante distribution of quality is given by the mixture of the types distribution with mixing probabilities $\pi(T)$. All our previous results hold for the mixture distribution.

Even if the total number of impressions suffices to satisfy the contracts, i.e. $\sum_{a \in \mathcal{A}} \rho_a \leq 1$, the inventory may not be enough to satisfy the contracts targeting criteria. Our algorithm guarantees that the total number of impressions C_a is always respected, yet some advertisers may be assigned impressions outside of their criteria. If an impression of type T happens to be assigned to an advertiser $a \notin T$, the publishers pays a nonnegative goodwill penalty τ_a. These penalties allow the publisher to prioritize certain reservations, specially when the contracts are not feasible.

2.3.1 Estimation of Placement Qualities

We study the performance of our algorithm on display ads data sets from two anonymous publishers; one online gaming website, and one news website. The data set is collected over a period of one week during March of 2010. The number of advertisers in these two instances are 3 and 6, respectively and the number of impressions in a data set range from 200 thousands to 3 millions. The targeting criteria of the guaranteed contracts is based on the URL, the geographic location, the type of browser or operating system used by the users, time of the day, and contextual features of web pages. Although the number of types may be exponential in A, in practice we observe that a linear number of them suffice to characterize 98% of the inventory. The first publisher has 4 user types while the second has 10 user types. The capacities of the reservations were used to compute the ratios ρ.

We associate a predicted click-through-rate to each impression, which is learned via a system that uses the impression's attributes as explanatory variables. We observe that the predicted quality perceived by the advertisers within a type is approximately log-normal. This can be seen in Figure 2.3, where the empirical distribution of log-quality is graphically represented for a type with two advertisers (data is log-transformed). The histograms on the off-diagonal show the marginal log-quality of each advertiser, which approximately resemble a normal curve. On

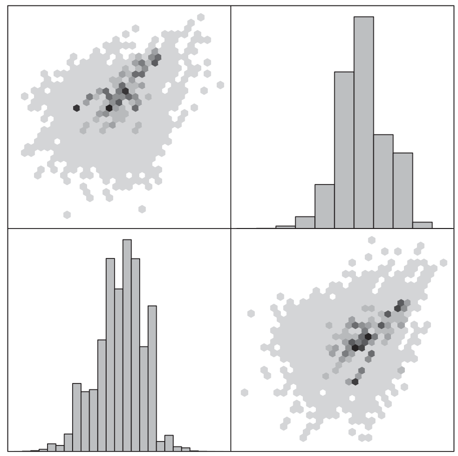

Figure 2.3: Graphical representation of the empirical distribution of log-quality for a type with two advertisers (data is log-transformed). The off-diagonals contain histograms of the marginal distributions, while the diagonals contain scatter plots of the joint distribution (the darker the bin, the higher the frequency).

the diagonals, scatter plots show the correlation between advertisers, which is strongly positive. In some sense this is expected, since many advertisers have similar targeting criteria. Given a particular type T, we assume that quality follows a multivariate log-normal with mean vector μ_T and covariance matrix Σ_T for the advertisers in the type, and takes a value of $-\tau_a$ for advertisers not in the type. The total distribution of quality is given by the mixture of these types distribution with mixing probabilities $\pi(T)$. Thus, we have that

$$Q \sim \begin{cases} \ln \mathcal{N}(\mu_T, \Sigma_T), & \text{for } a \in T, \\ -\tau_a, & \text{for } a \notin T, \end{cases} \quad \text{w.p. } \pi(T).$$

Logs were analyzed to estimate the types' frequencies, and the parameters of the underlying log-normal distributions (using maximum likelihood estimation).

2.3.2 Estimation of AdX Bids

Bidding data from the same period of time was used to estimate the primitives of AdX. With multiple bidders, AdX runs a sealed bid second-price auction. We analyze the first and second highest bids for the inventory submitted to AdX *independently* of the impression's attributes and placement qualities. Sample data is used to compute the two primitives of our model: (i) the complement of the quantile of the highest bid $p(s)$, and (ii) the revenue function of $r(s)$. Both functions are estimated on a uniform grid $\{s_j\}_1^{100}$ of survival probabilities in the $[0,1]$ range as follows. Let $\{(b_{1,m}, b_{2,m})\}_{m=1,\ldots,M}$ be the sampled highest and second highest bids from the exchange. First, for each point in the grid j, the price $p_j = p(s_j)$ is estimated as the $(1 - s_j)$-th population quantile of the highest bid. Then, using sampled bids, we estimate the revenue function w.r.t. to prices at the grid points as

$$r(p_j) = \frac{1}{M}\sum_{m=1}^{M} \mathbf{1}\{b_{1,m} \geq p_j\} \max\{b_{2,m}, p_j\} \tag{2.5}$$

Finally, the revenue function is obtained by composing (2.5) and $p(s)$. Figure 2.4a describes Instance 1, a publisher with 4 types and 3 advertisers. The estimated survival probability and revenue function for the publisher is shown in Figure 2.4b. The parameters for the remaining publishers are available at the web-page of the first author.

We provide some insight into the dependence structure between the guaranteed contracts' placement qualities and AdX bids by studying the Pearson's correlation between these two quantities over two publishers. The setup is as follows. First, we aggregate impressions by *ad slot*, where an ad slot refers to a given position in a publisher's web-page and is defined by the triple $l = (\text{position}, \text{web-page}, \text{publisher})$. The total number of ad slots L is in the order of thousands. Second, we compute the average value of maximum bids and average value of maximum placement qualities (predicted click-through-rates) over all impressions corresponding to each ad slot. Letting M_l be the number of impressions in ad slot l, the average value of maximum bids for the slot is $b_l = \frac{1}{M_l}\sum_{m\in l} b_{1,m}$, and the average value of maximum placement quality is $q_l = \frac{1}{M_l}\sum_{m\in l}\max_{a\in\mathcal{A}}\{q_{m,a}\}$. Finally, we compute the sample correlation coefficient, denoted by $r_{q,b}$, between the vectors of placement qualities $\{q_l\}_{l=1}^{L}$ and slot bids $\{b_l\}_{l=1}^{L}$. We find that the correlation of these two vectors is $r_{q,b} \approx -2\%$, and therefore conclude that correlation between the highest bid of an ad slot and the average placement quality is weak. As discussed earlier, this lack of correlation may be the result of advertisers' in AdX determining their bids using different signals from the publisher. The usual "lack of correlation does not imply independence" warning must apply here, and this finding should not be interpreted as a statement between independence of these two channels.

2.4 Experimental Results

In this section we present three numerical experiments conducted to study our model. First, we analyze the impact of introducing an AdX on the publisher's yield using actual publisher data. Second, we compare the performance of our policy with those of two popular heuristics. Finally, we discuss an efficient method to compute the dual variables and present results for large-scale instances.

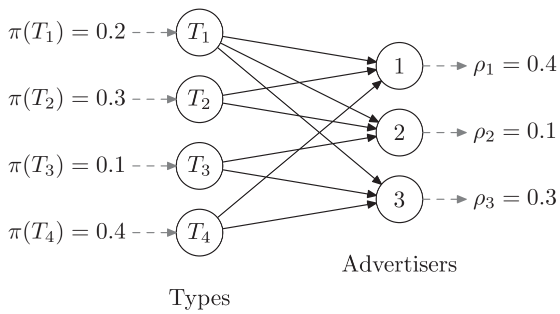

(a) User type-advertiser graph.

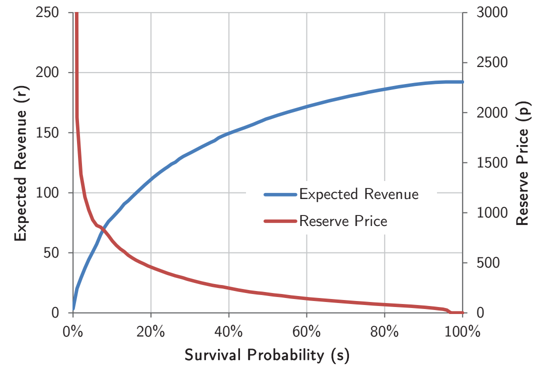

(b) Estimated survival probability and revenue function for AdX.

Type	Ads	$\pi(T)$	μ_T	Σ_T
T_1	$\{1,2,3\}$	0.2	$\begin{pmatrix} 7.8155 \\ 7.8155 \\ 7.8155 \end{pmatrix}$	$\begin{pmatrix} 0.3 & 0.1 & 0.1 \\ 0.1 & 0.3 & 0.1 \\ 0.1 & 0.3 & 0.1 \end{pmatrix}$
T_2	$\{1,2\}$	0.3	$\begin{pmatrix} 6.6755 \\ 7.0655 \end{pmatrix}$	$\begin{pmatrix} 0.3180 & 0.1649 \\ 0.1649 & 0.3602 \end{pmatrix}$
T_3	$\{2,3\}$	0.1	$\begin{pmatrix} 6.6355 \\ 7.8055 \end{pmatrix}$	$\begin{pmatrix} 0.4347 & 0.2357 \\ 0.2357 & 0.4367 \end{pmatrix}$
T_4	$\{1,3\}$	0.4	$\begin{pmatrix} 7.2155 \\ 6.9155 \end{pmatrix}$	$\begin{pmatrix} 0.23 & 0.05 \\ 0.05 & 0.40 \end{pmatrix}$

(c) Parameters of the distribution of log-quality.

Figure 2.4: Description of Instance 1.

2.4.1 Impact of AdX

This first experiment explores the potential benefits of introducing an AdX, and how the publisher can take advantage of it. We study the impact of the trade-off parameter γ on both objectives, that is, the quality of the impressions assigned to the advertisers, and the revenue from AdX. The limiting choices of $\gamma = 0$, and $\gamma = \infty$ are of particular interest. The first choice represents the case where the publisher disregards the quality of the impressions assigned to the advertisers, and strives to maximize the revenue extracted from AdX. Here the publisher strategically picks the reserve price so that just enough impressions are rejected to satisfy the contracts. In the second choice, the publishers prioritizes the quality of the impressions assigned, and submits the remanent inventory to AdX. We use this case as the baseline to which we compare our method.

The experiment was conducted as follows. First, we set up a grid on the trade-off parameter γ. Then, we solve the publisher's dual problem as given in (2.4) exactly (see A.4 for details on solving the stochastic program for our data model). Table 2.1 reports the expected quality

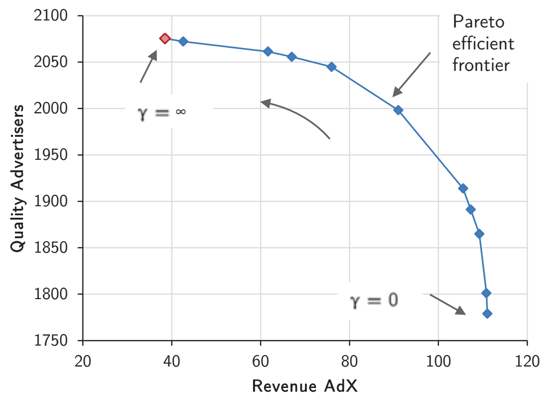

(a) Pareto Frontier (Publisher 1)

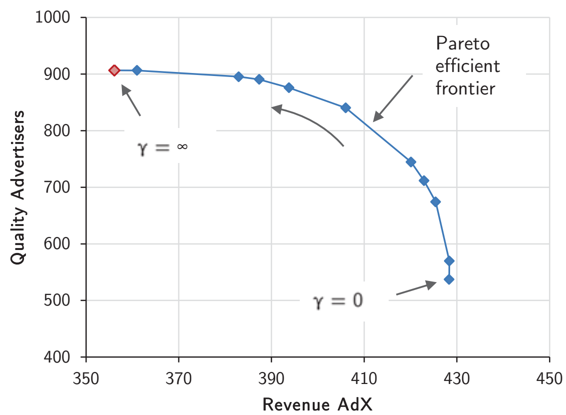

(b) Pareto Frontier (Publisher 2)

Figure 2.5: Plots, in a quality vs. revenue graph, of the objective values of the optimal solutions for the different choices of γ, together with the Pareto frontier.

and revenue for different choices of γ. In Figure 2.5 we plot, in a quality vs. revenue graph, the objective values of the optimal solutions for the different choices of γ, together with the Pareto frontier.

Discussion. Results confirm that, as we increase the trade-off parameter γ, the quality of the impressions assigned to the advertisers increases, while the revenue from AdX subsides. Interestingly, starting from the baseline case that disregards AdX ($\gamma = \infty$), we observe that the revenue from AdX can be substantially increased by sacrificing a small fraction of the overall quality of the impressions assigned. For instance, by exploiting strategically the AdX, the second publisher can increase AdX's revenue by 8% by giving up only 1% quality. Conversely, starting from the case that disregards the advertiser's quality ($\gamma = 0$), the publisher can raise the quality in a large

Instance 1

γ	0	0.001	0.01	0.05	0.075	0.1
Yield	110.94	112.78	128.88	202.46	249.11	296.95
Quality	1770.60	1779.07	1800.73	1864.95	1891.19	1913.78
Revenue	110.94	111.00	110.87	109.21	107.27	105.57

γ	0.25	0.5	0.75	1	10	∞
Yield	591.03	1098.58	1611.74	2127.22	20794.92	∞
Quality	1998.34	2045.17	2059.35	2065.05	2075.23	2075.52
Revenue	91.45	76.00	67.23	62.17	42.61	38.48

Instance 2

γ	0	0.001	0.01	0.05	0.075	0.1
Yield	428.73	429.02	434.22	459.57	477.39	495.66
Quality	544.69	545.27	573.23	676.34	720.31	752.35
Revenue	428.73	428.47	428.49	425.75	423.37	420.42

γ	0.25	0.5	0.75	1	10	∞
Yield	617.04	834.39	1056.05	1279.88	9433.13	∞
Quality	843.47	880.69	891.49	896.89	906.46	907.05
Revenue	406.17	394.05	387.43	382.99	368.53	356.11

Table 2.1: Expected yield, advertisers' quality and revenue from AdX for two instances, and different choices of γ.

amount at the expense of a small decrease in AdX's revenue.

Alternatively, the previous analysis can be understood in terms of the Pareto frontier. The Pareto frontier is highly concave, relatively horizontally flat around $\gamma = \infty$, and vertically flat around $\gamma = 0$. This explains the huge marginal improvements at the extremes. There are several advantages to the quality vs. revenue representation. First, the Pareto frontier allows for quick grasp of the nature of the operation. When the publisher's current operation is sub-optimal, its performance point should lie in the interior of the frontier. In this case, the Pareto frontier allows the publisher to measure its efficiency, and quantify the potential benefits an optimal policy may introduce. Second, when the choice of the trade-off parameter is not clear, the publisher may impose a lower bound on the overall quality of the impressions, and instead maximize the total revenue from AdX. The efficient frontier provides the maximum attainable revenue, and the proper γ to achieve the quality constraint.

2.4.2 Comparison with Greedy and Static Price Policy

This second experiment compares the performance of our policy with the following heuristics:

- *The Greedy Policy*, which disregards the opportunity cost of capacity and assigns the impression to the advertiser with maximum quality. The policy is allowed to dynamically price and test the exchange before the assignment. Similar corrections to the ones in the original stochastic policy are introduced to guarantee that all contracts are satisfied almost surely. Note that the Greedy Policy is equivalent to setting the dual variables to $v_a = 0$ in the bid-price policy.

- *The Static Price Policy*, which sets a constant reserve price for the exchange throughout the horizon, and commits to the exchange even if the impression is rejected. The policy is allowed to adjust the contracts' qualities by choosing optimal bid prices. In this case the optimal reserve price for the exchange is $p^*(0)$ and the publisher access the exchange only if the maximum contract adjusted quality is below $R(0)$.

Table 2.2 compares the expected yield of the optimal policy with the expected yield of the Greedy and Static Price policy for two instances, and different choices of γ. In order to objectively assess the performance of the different policies we employ fluid limit (see §A.5), that is, we report the limiting yield per impression as the capacity and number of impressions are simultaneously scaled to infinity.

Instance 1

γ	Yield	Yield Greedy	(Gap%)	Yield Static Price	(Gap%)
0.001	122.78	41.91	-65.9%	40.52	-59.3%
0.01	128.88	58.60	-54.5%	59.10	-53.9%
0.05	202.46	145.61	-28.1%	140.66	-30.5%
0.075	249.11	202.60	-18.7%	194.01	-22.1%
0.1	296.95	259.00	-12.8%	245.58	-17.3%
0.25	591.03	540.39	-8.6%	556.59	-5.8%
0.5	1098.58	920.43	-16.2%	1075.33	-2.1%
0.75	1611.74	1324.34	-17.8%	1595.06	-1.0%
1	2127.22	1736.93	-18.3%	2113.44	-0.6%
10	20794.92	16991.09	-18.3%	20776.54	-0.1%

Instance 2

γ	Yield Optimal	Yield Greedy	(Gap%)	Yield Static	(Gap%)
0.001	429.02	361.74	-15.7%	359.72	-16.2%
0.01	434.22	375.00	-13.6%	368.54	-15.1%
0.05	459.57	442.76	-3.7%	401.90	-12.5%
0.075	477.39	370.68	-22.4%	424.68	-11.0%
0.1	495.66	315.47	-36.4%	447.44	-9.7%
0.25	617.04	347.77	-43.6%	582.97	-5.5%
0.5	834.39	469.29	-43.8%	809.74	-3.0%
0.75	1056.05	605.31	-42.7%	1036.59	-1.8%
1	1279.88	745.45	-41.8%	1263.60	-1.3%
10	9433.13	5938.63	-37.0%	9429.19	0.0%

Table 2.2: Comparison of the expected yield of the optimal policy with the expected yield of the Greedy and Static Price policy for two instances, and different choices of γ.

Discussion. Results confirm that the Greedy Policy underperforms in the given instances with losses in yield of up to 65%. From a managerial perspective, the sub-optimality of the Greedy Policy stresses the importance of pondering the opportunity cost of capacity in performing the assignment of the impressions to the guaranteed contracts. If the publisher fails to take into account the opportunity cost of capacity, then some contracts are fulfilled early in the horizon and the opportunity to assign the top impressions is missed. To fix ideas, consider a publisher that agrees to split his inventory equally between two contracts with independent qualities distributed uniformly in [1,2] and [0,1], respectively. On the one hand, the Greedy Policy would first assign all impressions to the high-quality contract and the remainder to the low-quality contract, which results in total expected yield of $\frac{1}{2}(\frac{3}{2}+\frac{1}{2})=1$. On the other hand, the optimal policy sets the dual variables to $v_1=1$ and $v_2=0$ guaranteeing that only the top impressions are assigned to each contract, which results in a total expected yield of $\frac{7}{6}$.

The Static Price Policy tends to underperform when the trade-off parameter γ is close zero, that is, when the publisher strives to maximize the revenue extracted from AdX. If the publisher fails to dynamically adjust the auctions' reserve price to take into account the opportunity cost of the impression, he can incur losses in yield of up 60%. From a managerial perspective, these results show that the ability to dynamically price plays a key role in the joint optimization between the guaranteed contracts and the spot market. When the trade-off parameter γ is large, the exchange's revenue contribution to the yield is negligible, and the Static Price Policy is nearly optimal.

2.4.3 Large-scale Instances

In this section we present a solution method to efficiently solve large-scale instances. A difficulty of solving the stochastic optimization problem (2.4) is that the involved multidimensional integral cannot be computed with high accuracy when the publisher has many contracts. We tackle this problem by performing a Sample Average Approximation (SAA), which relies on approximating the underlying stochastic program via sampling, and then solving the approximate problem via a Sub-gradient Descent Method (SDM).

The basic idea of the SAA is simple: a random sample of placement qualities is generated and the expectation is approximated by the sample average function (Shapiro et al., 2009). Letting $\{q_m\}_{m=1}^{M}$ be an i.i.d. sample of M vectors of placement qualities, the SAA is given by

$$\min_{v} \frac{1}{M}\sum_{m=1}^{M} R\Big(\max_{a\in\mathcal{A}_0}\{q_{m,a}-v_a\}\Big)+\sum_{a\in\mathcal{A}}\rho_a v_a, \tag{2.6}$$

which is non-differentiable convex minimization problem. One can show that optimal solution and the objective value of SAA problem are consistent estimators of the optimal solution and objective value of the stochastic program, respectively (see, e.g., Shapiro et al. (2009)).

We solve the deterministic SAA problem via a SDM, which involves iterating the dual variables by taking steps on the opposite direction of any sub-gradient of the approximated objective with a proper step-size (see, e.g., Boyd and Mutapcic (2008) for a review on the topic). Starting from an initial solution $v^{(0)}$, our algorithm computes the new dual variables using the formula $v^{(k+1)} = v^{(k)} - \alpha_k g(v^{(k)})$, where $g(v) \in \mathbb{R}^A$ is a sub-gradient of the SAA objective at point v, and α_k is the step-size (we employ a constant step-length rule, that is, $\alpha_k = \alpha / \|g(v^{(k)})\|_2$). A sub-gradient of the SAA objective function is readily given by $g(v) \triangleq -\frac{1}{M}\sum_{m=1}^{M}\left(1 - s^*(q_{m,a_m^*} - v_{a_m^*})\right)\mathbf{1}_{a_m^*} + \rho$, where $a_m^* \in \arg\max_{a \in \mathcal{A}_0} q_{m,a} - v_a$ is any advertiser achieving the maximum in the m^{th} sample.

A certificate of sub-optimality of the dual solution can be established by constructing a feasible solution to the primal problem of the SAA, and then invoking weak duality to obtain a lower bound on the optimal value of the dual problem. The primal problem of (2.6) is given by

$$\begin{aligned} \max_{(s_m, i_m) \in \bar{\mathcal{P}}} \ & \frac{1}{M}\sum_{m=1}^{M} r(s_m) + \sum_{a \in \mathcal{A}} i_{m,a} q_{m,a} \\ \text{s.t.} \ & \frac{1}{M}\sum_{m=1}^{M} i_{m,a} = \rho_a, \forall a \in \mathcal{A}, \end{aligned}$$

where s_m and $i_{m,a}$ denote the probability that the m^{th} impression is accepted by AdX or assigned to advertiser a^{th}, respectively; and $\bar{\mathcal{P}} = \left\{(s,i) \in [0,1]^{A+1} : \sum_{a \in \mathcal{A}} i_a + s \le 1\right\}$. A feasible primal solution based on a dual solution v can be constructed by setting initially $s_m = s^*\left(\max_{a \in \mathcal{A}_0}\{q_{m,a} - v_a\}\right)$ and $i_{m,a} = (1 - s_m)\mathbf{1}\{a = a_m^*\}$; and then guaranteeing feasibility of the primal solution by applying corrections similar to the ones in Policy 2.

We test the algorithm by generating random instances with $A = 100$ advertisers.[3] The SAA is solved on a *training set* of $M = 10,000$ samples with 2000 iterations of the SDM, which in total take an average of 2 minutes in a personal computer[4]. Additionally, we measure the impact of the SAA on the stochastic program by evaluating the performance of the policy on a independent *test set* of $M' = 1,000,000$ impressions. Table 2.3 reports the aggregate results for a total of 100 random instances.

[3]The capacities of the contracts are drawn from $\rho \sim U_{[0,1]} \times U_{\Delta^A}$, where $U_{[0,1]}$ is a uniform random variable on $[0,1]$ and U_{Δ^A} is a uniform random vector over the probability simplex $\Delta^A = \{x \in \mathbb{R}^A : x_a \ge 0, \sum x_a = 1\}$. The instances have a single type with log-normal qualities with random mean vector $\mu = U_{[6,8]^A}$, and random covariance matrix $\Sigma = D'D$ with $D = U_{[0.4,0.6]^{A \times A}}/\sqrt{A}$.

[4]The algorithm is implemented in Matlab 7.11 and executed on a Windows PC with an Intel 2.0GHz CPU, and 4GB of RAM.

Training Set ($M = 10,000$ **impressions**)

	Median	Median Absolute Deviation
Dual yield	218.62	8.58
Primal yield	218.19	8.54
Duality Gap	0.25%	0.097%
CPU Time	83.4 seconds	15.1 seconds

Test Set ($M' = 1,000,000$ **impressions**)

	Median	Median Absolute Deviation
Primal yield	217.4	14.54
Error w.r.t. training	0.25%	0.17%

Table 2.3: Aggregate results of the sub-gradient descent method applied to the sample average approximation problem for a total of 100 random instances.

Discussion. Results confirm that our method provides an efficient procedure to solve the original stochastic program. Since both the SAA objective value and gradient can be computed in $O(MA)$ time, the SDM is able to quickly obtain the dual variables with a median duality gap of 0.2%. Additionally, by simulating the resulting policy in a larger sample we obtain that the median relative error incurred by the SAA is only of 0.25%, concluding that the SAA provides a good approximation to the original stochastic program.

One advantage of the SAA is that it provides a non-parametric approach to estimate the dual variables when the distribution of the placement quality is unknown. In order to solve the original stochastic minimization problem in practice, first, one needs to postulate a parametric model for placement quality (as done in §2.3), and then use a sample of data to learn the parameters of the underlying model. The SAA is powerful *because* it makes no distributional assumptions about the placement qualities, and directly learns the dual variables by replacing the expectation by a sample average function. Finally, in order to improve the performance of the policy the publisher can periodically resolve the dual problem, while delivering the impressions, by taking a few sub-gradient descent steps with the updated capacity-to-impression ratios.

2.5 Extensions

In this section we consider a number of extensions of the model and policy from the previous section.

2.5.1 AdX with Multiple Bidders

Here we generalize our results to the case where multiple buyers participate in the Ad Exchange. We model AdX as an auction with K risk neutral buyers with individual valuations drawn independently from the same distribution with c.d.f $F(\cdot)$, density $f(\cdot)$, and support $[p_0, p_\infty]$ (to simplify the notation we drop the dependence on the user attributes). Moreover, we assume that the distribution of the values have increasing failure rates, are absolutely continuous and strictly monotonic.

Myerson (1981) argued that under our assumptions the optimal mechanism is a Vickrey or second-price sealed-bid auction. Moreover, it is known that in such auctions bidding the true valuation is a dominant strategy for the buyers, and that the optimal reservation price $p^*(c)$ is independent of the number of buyers (Laffont and Maskin, 1980).

Let $B_{1:K}$ and $B_{2:K}$ be the order statistics which denote the highest and the second highest bid respectively. Given a reserve price p, the item is sold if $B_{1:K} \geq p$, i.e., there is some bid higher than the reserve price. The winning buyer pays the second highest bid, or alternatively $\max\{B_{2:K}, p\}$, since the seller should receive at least the reserve price p. Therefore, the publisher's maximization problem is

$$R(c) = \max_{p \geq 0} \mathbb{E}\left[\mathbf{1}\{B_{1:K} \geq p\} \max\{B_{2:K}, p\} + \mathbf{1}\{B_{1:K} < p\} c\right].$$

The setup of §2.1.2 can be consider as a particular case of a second-price auction in which we have only one bidder and $B_{2:K} = 0$.

As done previously, we cast our problem in terms of survival or winning probabilities. Then, letting s be the probability than the impression is sold, we have that $s = \mathbb{P}\{B_{1:K} \geq p\} = 1 - F^K(p)$ since valuations are i.i.d. Conversely, the reserve price as a function of the survival probability is given by $p(s) = \bar{F}^{-1}(1 - (1-s)^{1/K})$, which is well-defined due to the strict monotonicity of the c.d.f. In terms of survival probabilities, the problem is now

$$R(c) = \max_{0 \leq s \leq 1} r(s) + (1-s)c,$$

where we defined the revenue function as $r(s) = r(p(s))$, and $r(p) = \mathbb{E}\left[\mathbf{1}\{B_{1:K} \geq p\} \max\{B_{2:K}, p\}\right]$.

The next proposition shows that the revenue function is regular, and as a consequence all previous results hold for the case with multiple bidders.

Proposition 2.4. *Under the previous assumption, the revenue function $r(s)$ is regular. Moreover, the optimal reserve price $p^*(c)$ solves*

$$\frac{\bar{F}(p)}{f(p)} = p - c,$$

when $c \in [p_0 - 1/f(p_0), p_\infty]$. When the opportunity cost is higher than the null price ($c > p_\infty$), the publisher bypasses the exchange ($p^*(c) = p_\infty$). Finally, when the opportunity cost is low enough ($c < p_0 - 1/f(p_0)$), the impression is kept by the highest bidder ($p^*(c) = p_0$).

2.5.2 Covering Constraints

Guaranteed contracts typically specify a lump-sum amount in return for a fixed number of impressions and the publisher is not be monetarily rewarded for delivering impressions beyond these targets. In some settings, however, the publisher may seek to exceed these contractual targets in view of attracting feature business, at the expense of reducing the revenue from the exchange.

Our model is quite general and allows to easily accommodate *covering constraints*, that is, the case where the number of impressions assigned to each contract should be greater or equal to the capacity. In this case the capacity constraint of the DAP is relaxed to $\sum_{n=1}^{N} \mathbb{E}\left[i_{n,a}\right] \geq N\rho_a$, for all $a \in \mathcal{A}$. The analysis proceeds as before with the only difference that now the dual variables are non-negative, that is, the publisher should solve its dual problem under the constraint that $v_a \geq 0$. Additionally, when implementing the stochastic policy the publisher should now allow contracts to exceed their capacity. This amounts to determining the maximum contract-adjusted quality between all contracts $a \in \mathcal{A}$ (when the total number of impressions left is greater than the number of impressions necessary to fulfill the contracts), or equivalently changing Line 4 of Policy 2 to $a_n^* = \arg\max_{a \in \mathcal{A} \cup \{0\}} \{Q_{n,a} - v_a\}$. Regarding the performance the bid-price control μ^B, Theorem 2.2 still holds in this setting.

2.5.3 Target Quality Constraints

Some publishers might feel more comfortable specifying target quality constraint than picking a Lagrange multiplier to weight the impact of quality in the objective. In other settings the advertisers themselves might demand that certain level of quality is guaranteed. In the following, we consider the case where the publisher strives to maximize the revenue from AdX, while complying with target quality constraints and capacity constraints.

The publisher imposes that the average quality of the impressions assigned to advertiser a is larger or equal than a threshold value ℓ_a. The one-impression DAP is similar, except that the objective only accounts for AdX's revenue, and for the inclusion of the constraints

$$\mathbb{E}\left[i_a(Q)Q_a\right] \geq \ell_a, \qquad \forall a \in \mathcal{A}. \tag{2.7}$$

Let $\gamma_a \geq 0$ be the Lagrange multiplier associated to (2.7). Problem (2.3) can be interpreted as the Lagrange relaxation of our new problem w.r.t. the target quality constraints, and the dual variables γ_a as the shadow prices of the target quality constraints. The new constraints preserve the convexity of the primal program, and strong duality still holds. Following the same steps, we obtain the new dual problem

$$\min_{\gamma \geq 0, v} \left\{ \mathbb{E}R\left(\max_{a \in \mathcal{A}_0} \{\gamma_a Q_a - v_a\} \right) + \sum_{a \in \mathcal{A}} v_a \rho_a - \gamma_a \ell_a \right\},$$

which still is a convex minimization problem. The publisher now jointly optimizes over v, and γ to determine the bid-prices of the stochastic policy.

Regarding the performance of bid-price control, one can reproduce the steps of Theorem 2.2's proof to show that the policy asymptotically attains the optimal revenue from AdX, while complying with the delivery targets. Additionally, from the same asymptotic analysis one obtains that the expected average quality is lower bounded by

$$\mathbb{E}\left[\frac{1}{N} \sum_{n=1}^{N} \left[1 - s_n^{\mu}(Q_n)\right] I_{n,a}^{\mu}(Q_n) Q_{n,a} \right] \geq \left(1 - \frac{1}{\sqrt{N}} K(\rho)\right) \mathbb{E}\left[i_a^*(Q) Q_a\right].$$

Hence, for advertisers with binding constraint (2.7), albeit not feasible, the expected average quality becomes arbitrary close to the threshold value as the number of impressions in the horizon increases. For the remaining advertisers whose target quality constraint is not binding, the expected average quality will surpass the threshold for suitably large N.

2.5.4 Handling ties

Theorem 2.1 had an assumption that there would be no ties between advertisers verifying the maximum $Q_a - v_a$. In this section we show how to construct a primal optimal solution to the DAP and the corresponding stochastic policy in the general case (for example, when the distribution of placement quality is discrete or has atoms). Devenur and Hayes (2009) proposed introducing small random and independent perturbations to the qualities, or smoothing the dual problem to break ties. We provide an alternate method that directly attacks ties, and provides a randomized tie-breaking rule. Computing the parameters of the tie-breaking rule requires solving a flow problem on a graph of size $2^{|\mathcal{A}|}$; thus in some settings it may not be possible. In section A.4, we show that in practice ties do not occur frequently. However, for completeness we provide a full characterization of the problem.

For any non-empty subset $S \subseteq \mathcal{A}_0$, we define a *S-tie* as the event when the maximum is verified exactly by all the advertisers $a \in S$, and the impression is rejected by AdX. Note that the tie may be a singleton, in the case that exactly one advertiser verifies the maximum. Since the dual variables v are known, the probability of such event can be written as

$$\mathbb{P}(S\text{-tie}) = \mathbb{E}\Big[\big(1 - s^*(\lambda(Q);U)\big)\mathbf{1}\big\{Q_a - v_a = \lambda(Q)\ \forall a \in S,\ Q_a - v_a < \lambda(Q)\ \forall a \notin S\big\}\Big],$$

where $\lambda(Q) = \max_{a\in\mathcal{A}_0}\{Q_a - v_a\}$. With some abuse of notation we define the $\emptyset$-*tie* as the event when the impression is accepted by AdX, that is, $\mathbb{P}(\emptyset\text{-tie}) = \mathbb{E}[s^*(\lambda(Q);U)]$. Note that the tie events induce a partition of the quality space, and we have that $\sum_{S\subseteq\mathcal{A}_0}\mathbb{P}(S\text{-tie}) = 1$.

We look for a random tie-breaking rule that assigns an arriving impression to advertiser $a \in S$ with conditional probability $I_a(S)$ given that a S-tie occurs. Hence, the routing probabilities depend on which advertisers tie, and not on the particular realization of the qualities (they are independent of $\lambda(Q)$). Therefore, under such policy the total probability, originating from S-ties, of an impression being assigned to advertiser a is $y_a(S) = \mathbb{P}(S\text{-tie})I_a(S)$. We can interpret $y_a(S)$ as the normalized flow of impression assigned to the advertiser originating from S-ties. We will show that, in terms of $y_a(S)$ as decision variables, finding the tie-breaking rule amounts to solving a transportation problem.

First, in order for the publisher to fulfill the contract with an advertiser $a \in \mathcal{A}$ the incoming flow of impressions over all possible ties sums up to ρ_a. The previous constraint can be written as

$$\sum_{S\subseteq\mathcal{A}_0:a\in S} y_a(S) = \rho_a, \qquad \forall a \in \mathcal{A}. \tag{2.8}$$

Notice that we impose no constraints for $a = 0$ since any number of impressions can be discarded. Alternatively, we could set $\rho_0^{\text{eff}} = 1 - \mathbb{P}(\emptyset\text{-tie}) - \sum_{a\in\mathcal{A}}\rho_a$ because the impressions effectively discarded are those that are rejected by AdX and not assigned to an advertiser. Second, the outgoing flow of impressions originating from a particular tie should sum up to the actual probability of that tie occurring. Then, we have that

$$\sum_{a\in S} y_a(S) = \mathbb{P}(S\text{-tie}), \qquad \forall S \subseteq \mathcal{A}_0. \tag{2.9}$$

Third, we require that $y_a(S) \geq 0$ for all $S \subseteq \mathcal{A}_0$ and $a \in S$. Finally, in order to obtain the tie-breaking rule we need a non-negative flow satisfying constraints (2.8) and (2.9). Once a such solution is found, the optimal controls can be computed as

$$I_a(U) = \begin{cases} y_a(S)/\mathbb{P}(S\text{-tie}) & \text{if } a \in S, \text{ and } Q \text{ is an } S\text{-tie},\\ 0 & \text{otherwise,}\end{cases}$$

with the pricing function as before.

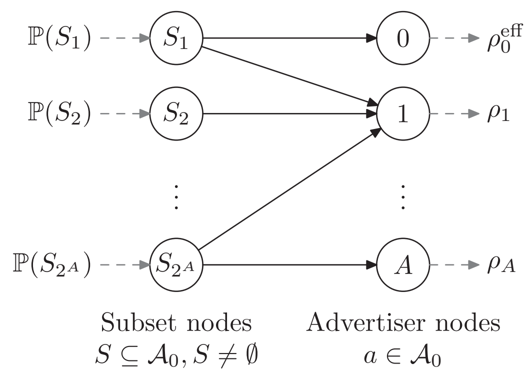

Figure 2.6: Bipartite flow problem solved to obtain the tie-breaking rule. On the left-hand side of the graph we include one node for each non-empty subset $S \subseteq \mathcal{A}_0$ (subset nodes), and in the right-hand side we add one node for each advertiser $a \in \mathcal{A}_0$ (advertiser nodes). The supply for subset nodes is $\mathbb{P}(S\text{-tie})$, while the demand for advertiser nodes is ρ_a. Arcs in the graph represent the membership relation, i.e., the subset node S and advertiser node a are connected if and only if $a \in S$. Arc capacities are set to infinity.

It is not hard to see that the previous problem can be stated as a feasible flow problem in a bipartite graph. We briefly describe how to construct such graph next. On the left-hand side of the graph we include one node for each non-empty subset $S \subseteq \mathcal{A}_0$, and in the right-hand side we add one node for each advertiser $a \in \mathcal{A}_0$. In the following we refer to nodes in the left-hand side as *subset nodes*, and to those in the right-hand side as *advertiser nodes*. The supply for subset nodes is $\mathbb{P}(S\text{-tie})$, while the demand for advertiser nodes is ρ_a. Arcs in the graph represent the membership relation, i.e., the subset node S and advertiser node a are connected if and only if $a \in S$. Moreover, arc capacities are set to infinity. In Figure 2.6 the resulting bipartite graph is shown.

An important question is whether the flow problem admits a feasible solution. The next result proves that the answer is affirmative when the dual variables v are optimal for the dual problem (2.4). The proof proceeds by casting the feasible flow problem as a maximum flow problem, and then exploiting the optimality conditions of v to lower bound every cut in the bipartite graph.

Proposition 2.5. *Suppose that $v \in \mathbb{R}^A$ is an optimal solution for the dual problem* (2.4)*. Then, there exists a non-negative flow satisfying constraints* (2.8) *and* (2.9)*.*

We conclude this section by showing that the solution constructed is optimal for the primal problem. Notice that the solution is feasible because it satisfies constraints (2.8) and (2.9). In order to prove optimality it suffices to show that it attains the dual objective value, or that it satisfies the complementary slackness conditions. The latter follows trivially.

Once the optimal controls are calculated, we construct our stochastic policy as follows. We let $\mathcal{A}_n^* = \arg\max_{a \in \mathcal{A}_n \cup \{0\}} \{Q_{n,a} - v_a\}$, be the set of advertisers that attain the maximum. Now, if the impression is rejected by AdX and $\mathcal{A}_n^* \neq \{0\}$, we assign it to advertiser a in $\mathcal{A}_n^*$ with probability $I_a(U_n)/\sum_{a' \in \mathcal{A}_n^*} I_{a'}(U_n)$. Notice that impressions are only assigned to advertisers with contracts that have yet to be fulfilled. Additionally, as the contracts of some advertisers are fulfilled, these are excluded of the assignment, and the routing probabilities $I_a(\cdot)$ of the remaining advertisers are scaled-up and normalized.

2.6 Conclusions

Ad Exchanges are an emerging market for the real-time sale of online ad slots on the Internet. In this work we present an approach to help publishers determine when and how to access AdX to complement their contract sales of impressions. In particular, we model the publishers' problem as a stochastic control program and derive an asymptotically optimal policy with a simple structure: a bid-price control extended with a pricing function for the exchange. We show using data from real inventory that there are considerable advantages for the publishers from jointly optimization over both channels. Publishers may increase their revenue streams without giving away the quality of service of their reservations contracts, which still represents a significant portion of their advertising yield. We also hope our insights here will help understand ad allocation problems more deeply.

Internet advertising, and in particular AdX, is likely to prove to be a fertile area of research. There are several promising directions of research stemming from this work. One intuitive approach to improve the performance of a control, which is appealing for its simplicity, consist on resolving the deterministic approximation periodically throughout the horizon. In a follow-up work we intend to show that one can indeed improve on the static control and obtain sharper bounds by resolving the DAP. Another problem that needs further study is that of learning in the case of unknown distributions, which is of great importance given the fast-paced and changing nature of the Internet. There exists independent research on online algorithms for capacity allocation and online pricing for repeated auctions, but none on the joint optimization problem. Finally, as more publishers reach out for AdX, advertisers will have the opportunity to buy their inventory from either market. The existence of two competing channels, the exchange as a spot market and the reservations as future market, introduces several interesting research questions. How should publishers price their contracts and allocate their inventory? How should advertisers hedge their campaign between these two markets? We hope that this work pave the way for further research on this important topic.

Chapter 3

Advertiser Competition in AdX: Auction Design

The material presented in this chapter is based on the working paper Balseiro et al. (2012*b*) co-authored with Omar Besbes and Gabriel Weintraub.

In this chapter we study the competitive landscape that arises in Ad Exchanges and the implications for publishers' decisions. In Section 3.1 we introduce a stochastic model that captures the main characteristics of advertiser competition in AdX. In this model advertisers join the market with a pre-specified budget and participate in multiple second-price auctions over the length of a campaign. As a result advertisers engage in a dynamic game, whose equilibrium determines the market characteristics. In Section 3.2 we introduce the novel notion of a Fluid Mean Field Equilibrium (FMFE) that is behaviorally appealing, computationally tractable, and in some important cases yields a closed-form characterization. In Section 3.3 we provide a sharp characterization of the equilibrium strategies under the FMFE. In particular, we show that an FMFE always exists and provide a broad set of sufficient conditions that guarantee its uniqueness. In Section 3.4 we use this framework to provide sharp prescriptions for key auction design decisions that publishers face in these markets, such as the reserve price, the allocation of impressions to the exchange versus an alternative channel, and the disclosure of viewers' information. We conclude with some final remarks in Section 3.5.

3.1 Model Description

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

Advertisers. Advertisers arrive to the exchange according to a Poisson process $\{K(t)\}_{t\geq 0}$ with intensity λ. We index the sequence of arriving advertisers by $k \geq 1$, and denote by $\{\tau_k\}_{k\geq 1}$ the arrival times. Advertiser k is characterized by a type vector $\theta_k = (b_k, s_k, \alpha_k, \gamma_k) \in \mathbb{R}^4$. The first component of the type, b_k, denotes the budget and the second component, s_k, denotes the campaign length. That is, the k^{th} advertiser's campaign takes place over the time horizon $[\tau_k, \tau_k + s_k)$ and her total expenditure cannot exceed b_k. Once the advertiser leaves the exchange she never comes back.

When the publisher contacts the exchange she submits some partial information about the user visiting the website, that for example, could include *cookies*. This information, in turn, may heterogeneously affect the value an advertiser perceives for the impression, and the amount she is willing to bid. The last two components of the type, α_k and γ_k, determine the targeting criteria and the valuation distributions as we now explain. When the n^{th} user arrives, the advertisers in the exchange observe the user information disclosed by the publisher, and determine whether they will participate or not in the auction based on their targeting criteria. We assume that the k^{th} advertiser matches a user with probability α_k independently and at random (both across impressions and advertisers). Conditional on a match, advertisers have independent private values for an impression. In particular, all values for advertiser k are independent and identically distributed random variables with a continuous cumulative distribution $F_v(\cdot;\gamma_k)$, parametrized by $\gamma_k \in \mathbb{R}$. The distributions have compact support $[\underline{V}, \bar{V}] \subset \mathbb{R}_+$ and continuously differentiable density.[1] Later, we will explain how the publisher can affect the value of α_k and the distribution parameter γ_k by changing her user information disclosure policy.

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

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

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

Publisher. On the supply side, the publisher has an opportunity cost for selling her inventory of impressions in the exchange; that is, the publisher obtains some fixed amount $c > 0$ for each impression not won by some advertiser in the exchange. The publisher's payoff is given by the long-run average profit rate generated by the auctions, where the profit is measured as the difference between the payment from the auction and the lost opportunity cost c when the impression is won by an advertiser in the exchange. The publisher's objective is to maximize its payoff. We analyze three levers that the publisher may use to do so: (i) the reserve price r to set for the auctions, (ii) the volume of ads to send to the exchange given the positive external opportunity cost c, and (iii) the amount of information she discloses to the advertisers that we will represent with a real number ι.[2]

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

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

3.2 Equilibrium Concept

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

A standard solution concept used for dynamic games of incomplete information is that of weak perfect Bayesian equilibrium (WPBE) (Mas-Colell et al., 1995). Roughly speaking, in such a game, a pure *strategy* for advertiser k is a mapping from histories to bids, where the histories represent past observations. A strategy specifies, given a history and assuming the advertiser participates in an auction at time t, an amount to bid. A strategy profile in conjunction with a belief system constitute a WPBE if the following holds. First, given a belief system and the com-

[2] To simplify notation, we will not make this dependence explicit until needed.

petitors' strategies, an advertiser's bidding strategy maximizes expected future payoffs. Second, beliefs must be consistent with the equilibrium strategies and Bayes' rule whenever possible.

WPBE and commonly used refinements, such as perfect Bayesian equilibrium and sequential equilibrium, require advertisers to hold beliefs about the entire future dynamics of the market, including the future market states. With more than few competitors in the market this imposes a very strong rationality assumption over advertisers as these belief distributions are high-dimensional. Moreover, to find a best response, advertisers need to solve a dynamic programming problem that optimizes over history-dependent strategies. This optimization problem can be high-dimensional and be intractable both analytically and computationally. Hence, solving for WPBE for most markets of interest is not possible. More importantly, WPBE imposes informational requirements and a level of sophistication on the part of agents that seems highly unrealistic. This motivates the introduction of alternative equilibrium concepts. After some background in §3.2.1, we introduce such an alternative in §3.2.2.

3.2.1 Mean Field and Fluid Approximation

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

To win an auction, an advertiser competes against all other bidders and against the reserve price r. We denote by D the steady-state maximum of the "competitors' bids", where we assume the publisher is one more competitor that submits a bid equal to r. Assume for a moment that D is i.i.d. across different auctions and distributed according to a c.d.f $F_d(\cdot)$. (Note that $F_d(\cdot)$ will be endogenously determined in equilibrium in §3.2.2).

In this setting, the advertiser's dynamic bidding problem in the repeated auctions can be casted as a revenue management-type stochastic dynamic programming problem, in which bidding decisions across periods are coupled through the budget constraint. However, the resulting Hamilton-Jacobi-Bellman is a partial differential equation that, in general, does not have a closed-form solution. To get a better handle on the bidder's dynamic optimization problem we introduce a second level of approximation motivated by the fact that a given advertiser has a very large number of bidding opportunities (campaigns span for weeks or months, and thousands of impressions arrive per day). In such an environment, the advertiser's stochastic dynamic programming problem can be well approximated through a stochastic fluid model. In particular, the approximation we focus on is predicated on assuming that bidders solve a control problem in which the budget constraint need only be satisfied in expectation, and bidding strategies ignore the individual state and are only dependent on the actual realization of the bidder's value. This can be shown to provide advertisers with provably good policies when both the number of impressions and budgets are large, so the number of bidding opportunities over the campaign length also grows large.

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

$$J_\theta^{\mathrm{F}}(F_d) = \max_{w(\cdot)} \ \alpha\eta s \mathbb{E}\Big[\mathbf{1}\{D \le w(V)\}(V - D)\Big] \tag{3.1a}$$

$$\text{s.t. } \alpha\eta s \mathbb{E}\Big[\mathbf{1}\{D \le w(V)\} D\Big] \le b, \tag{3.1b}$$

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

Proposition 3.1. *Suppose that* $\mathbb{E}[D] < \infty$*. An optimal bidding strategy that solves* (3.1) *for type* θ *is given by*

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

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

The optimal bidding strategy has a simple form: an advertiser of type θ needs to *shade* her values by the *constant factor* $1 + \mu_\theta^*$, and this factor guarantees that the advertiser's expenditure does not exceed the budget. In the previous expression the multiplier μ_θ^* is the shadow price of the budget constraint and gives the marginal utility in the advertiser's campaign of one extra unit of budget. Intuitively, when budgets are tight, advertisers shade their bids, because there is an option value for future good opportunities. When budgets are not tight, the optimal dual multiplier is equal to zero and advertisers bid truthfully. The proof of the result relies on an analysis of the dual of problem (3.1). Note that (3.1) is not a convex program and hence standard strong duality arguments do not apply. As a result, the proof establishes from first principles that no duality gap exists in the present case.

3.2.2 Fluid Mean Field Equilibrium

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

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

A difficulty with the consistency check is that the number of *active* bidders, those that match the target criteria and have remaining budgets, depends on the market dynamics. In particular, the budget dynamics depend on who wins and how much the winner pays in each auction. Hence, in principle, characterizing the resulting steady-state distribution of the maximum bid of active competitors' (that have remaining budgets) is complex. However, it is reasonable to conjecture that, when the number of opportunities during the campaign length is large, rational advertisers would deplete their budgets close to the end of their campaign with high probability. For analytical tractability we impose that, in our equilibrium concept, any bidder currently in the exchange that matches the targeting criteria, without regard of her budget, gets to bid. Under this assumption, the number of bidders in an auction equals the number of advertisers matching the targeting criteria, denoted by $M(t)$, which is just an independent sampling from the process $Q(t)$. We rigorously justify that this additional layer of approximation is asymptotically correct in large markets (see Appendix 4.2 and Balseiro et al. (2012*a*)).[3]

Since arrival and departures of advertisers are governed by an $M/G/\infty$ queue and campaign lengths are bounded, it is not hard to show that under fluid-based strategies the market process Ω is Harris recurrent, so it is ergodic and admits a unique invariant steady-state distribution (see, e.g., Asmussen (2003) p. 203). Let $\left(\hat{M}, \{\hat{\Theta}_k\}_{k=1}^{\hat{M}}\right)$ be a random vector that describes the number of matching bidders, together with their respective types when sampling a market state according to the invariant distribution. Notice that advertisers with longer campaign lengths and higher matching probability are more likely to participate in an auction, and thus the law of a type sampled from the invariant distribution does not coincide with the law of the types in the population. Indeed, by exploiting the fact that arrival-time and service-time pairs constitute a Poisson random measure on the plane (see, e.g., Eick et al. (1993)), one can show that $\hat{M}$ is Poisson with parameter $\mathbb{E}[\alpha_\Theta \lambda s_\Theta]$, and that each component of the vector of types is independently and identically distributed as $\mathbb{P}_{\hat{\Theta}}\{\theta\} = \frac{\alpha_\theta s_\theta}{\mathbb{E}[\alpha_\Theta s_\Theta]} \mathbb{P}_\Theta\{\theta\}$ for each type $\theta \in \mathbf{\Theta}$, and independent of $\hat{M}$.[4]

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

$$\max\left(\left\{\beta_{\hat{\Theta}_k}(V_{\hat{\Theta}_k})\right\}_{k=1}^{\hat{M}}, r\right), \tag{3.2}$$

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

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

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

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

that the distribution of the maximum bid of other advertisers is given by $F_d(\boldsymbol{\beta})$ (equation (3.2)).

Essentially a FMFE is a set of bidding strategies such that (i) these strategies induce a given competitive landscape as represented by the steady-state distribution of the maximum bid, and (ii) given this landscape, advertisers' optimal fluid-based bidding strategies coincide with the initial ones. We focus on symmetric equilibria in the sense that all bidders of a given type adopt the same strategy. Note that in the FMFE, upon arrival to the system an advertiser is assumed to compete against the market steady-state maximum bid D.[5] Finally, for a given set of auction design parameters (r, η, ι), we denote by $\mathcal{C}_{\text{FMFE}}(r, \eta, \iota)$ the set of corresponding FMFEs.

We introduced FMFE by heuristically arguing that it should be a sensible equilibrium concept for large markets when the number of bidding opportunities per advertiser are also large. In Theorem 4.2 in Appendix 4.2, we show that when all advertisers implement the FMFE strategy, the relative profit increase of any unilateral deviation to a strategy that keeps track of all information available to the advertiser becomes negligible as the scale of the market increases. This provides a theoretical justification for using FMFE as an approximation of advertisers' behavior.

3.2.3 Publisher's Problem

We model the grand game played between the publisher and advertisers as a Stackelberg game in which the publisher is the leader and the advertisers the followers. In particular, the publisher first selects the reserve price in the second-price auction r, the rate of impressions η, and the extent of information disclosed ι. Then, after observing these design parameters, the advertisers react and play the induced dynamic game among them. In our analysis we assume that the solution concept for the game played between advertisers is that of FMFE.

The publisher's optimization problem may be written as

$$\max_{r,\eta,\iota} \quad \eta \mathbb{E}\Big[\mathbf{1}\{\beta_{\hat{\Theta}}(V_{\hat{\Theta}})_{1:\hat{M}} \geq r\}\Big(\max\{\beta_{\hat{\Theta}}(V_{\hat{\Theta}})_{2:\hat{M}}, r\} - c\Big)\Big] \tag{3.3a}$$

$$\text{s.t.} \quad \boldsymbol{\beta} \in \mathcal{C}_{\text{FMFE}}(r, \eta, \iota) \tag{3.3b}$$

where $\beta_{\hat{\Theta}}(V_{\hat{\Theta}})_{1:\hat{M}}$ is the highest bid among the bids $\left\{\beta_{\hat{\Theta}_k}(V_{\hat{\Theta}_k})\right\}_{k=1}^{\hat{M}}$ submitted in the auction and $\beta_{\hat{\Theta}}(V_{\hat{\Theta}})_{2:\hat{M}}$ the second-highest bid. The publisher's problem amounts to maximizing her long run average profit rate (3.3a) considering the opportunity cost of the alternative channel, subject to the constraint that β is an FMFE. For given (η, r, ι), a priori, a FMFE may not exist or even if one exists it may not be unique. Hence, the optimization problem above is not well posed in

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

general. However, we will prove in the coming sections that a FMFE exists and is unique for an important class of models.

Note that both r and η directly affect the publisher's objective. In addition, it is clear from equations (3.1), that they could affect bidders' FMFE strategies. Furthermore, recall that ι affects the matching probabilities $\alpha_k(\iota)$ and the valuation parameters $\gamma_k(\iota)$. Hence, changing ι also directly affects the publisher's objective and the FMFE being played. In the next sections, we make these dependencies more explicit and analyze the optimal selection of each of those design decisions.

3.3 Fluid Mean Field Equilibrium Characterization

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

3.3.1 Equilibrium Existence and Sufficient Conditions for Uniqueness

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

$$\mu_\theta^* \in \arg\min_{\mu \geq 0} \Psi_\theta(\mu; \boldsymbol{\mu}^*), \qquad \text{for all types } \theta \in \Theta. \tag{3.4}$$

One may establish that the system of equations (3.4) always admits a solution to obtain the following.

Theorem 3.1. *There always exists an FMFE.*

The proof shows that the dual strategy space can be reduced to a compact set, and that the dual objective function is jointly continuous in its arguments, and convex in the first argument.

Then, a standard result that relies on Kakutani's Fixed-Point Theorem implies existence of an FMFE.

We now turn to sufficient conditions for uniqueness. Let $\mathbf{G} : \mathbb{R}_+^{|\boldsymbol{\Theta}|} \times \mathbb{R}_+ \rightarrow \mathbb{R}_+^{|\boldsymbol{\Theta}|}$ be a vector-valued function that maps a profile of multipliers and a reserve price to the steady-state expected expenditures per auction of each bidder type. The expected expenditure of a θ-type bidder in a second-price auction when advertisers implement a profile of multipliers $\boldsymbol{\mu}$ is given by

$$G_\theta(\boldsymbol{\mu}, r) \triangleq \mathbb{E}\left[\mathbf{1}\{(1+\mu_\theta)D \leq V\}D\right],$$

where the maximum competing bid is given by $D = \max\left(\left(V_{\hat{\Theta}}/\left(1+\mu_{\hat{\Theta}}\right)\right)_{1:\hat{M}}, r\right)$ [6].

Assumption 3.1 (P-matrix)**.** *The Jacobian of $-\mathbf{G}$ with respect to $\boldsymbol{\mu}$ is a P-matrix for all $\boldsymbol{\mu}$ in $\mathbb{R}_+^{|\boldsymbol{\Theta}|}$*

A matrix is P-matrix if all its principal minors are positive (Horn and Johnson, 1991, p.120). Assumption 3.1 can be shown to hold for various cases of interest. For example, it is easy to see that it always holds for the case of homogeneous advertisers, i.e., when the space of types $\boldsymbol{\Theta}$ is a singleton. In Appendix B.2, we provide examples of settings with heterogeneous advertisers in which it also holds. The next theorem shows that the equilibrium is unique under the P-matrix assumption.

Theorem 3.2. *Suppose Assumption 3.1 holds. Then, there is a unique FMFE of the form $\beta_\theta(v) = v/(1+\mu_\theta)$, θ in Θ.*

We prove the result by formulating the FMFE conditions as a Non-linear Complementarity Problem (NCP) as presented in Corollary 3.1 below, and employing a Univalence Theorem to show that the expected expenditure mapping is injective (Facchinei and Pang, 2003*a*). Moreover, it is possible to establish under further mild regularity conditions that any set of continuous increasing bidding functions that constitute an FMFE necessarily yield the same outcome (in terms of auctions' allocations and payments) as that of the FMFE in Theorem 3.2. In the rest of the paper, we focus on the simple and intuitive FMFE strategies that can be described by a vector of dual multipliers.

3.3.2 Equilibrium Characterization

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

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

Corollary 3.1. *Any FMFE characterized by a vector of multipliers $\boldsymbol{\mu}^*$, such that $\beta_\theta(v) = v/(1+\mu_\theta^*)$ for all $v \in [\underline{V}, \bar{V}]$ and $\theta \in \boldsymbol{\Theta}$, solves*

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

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

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

We conclude this subsection by refining the result for the case of homogeneous bidders, in which one can provide a quasi-closed form characterization for FMFE. Suppose that Θ is a singleton. In this case, we shall see that Assumption 3.1 always holds and there exists a unique FMFE. Let

$$G_0(r) = G_\theta(0, r) \tag{3.5}$$

denote the steady-state unconstrained expected expenditure-per-auction of a single bidder for a second price auction with reserve price r when all advertisers (including herself) bid their own values. Note that the expected expenditure for a bidder over its campaign when all bidders are truthful is given by $\alpha \eta s G_0(r)$. This quantity plays a key role in the FMFE characterization.

Proposition 3.2. *Suppose $\boldsymbol{\Theta}$ is a singleton. Then a Fluid Mean Field Equilibrium exists and is unique. In addition, the equilibrium may be characterized as follows: $\beta_\theta(v) = v/(1+\mu^*)$ for all $v \in [\underline{V}, \bar{V}]$, where $\mu^* = 0$ if $\alpha \eta s G_0(r) < b$, and μ^* is the unique solution to $\alpha \eta s G_0(r(1+\mu)) = b(1+\mu)$ if $\alpha \eta s G_0(r) \geq b$.*

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

3.3.3 Reformulation of the Publisher's Problem

When Assumption 3.1 holds, given the existence and uniqueness of a FMFE, one may now properly and explicitly formulate the publisher's problem. Let $I(\boldsymbol{\mu}, r) = 1 - F_d(r; \boldsymbol{\mu})$ denote the probability that the impression is won by some advertiser in the exchange when advertisers shade according to the profile $\boldsymbol{\mu}$ and the publisher sets a reserve price r. Using the characterization of an FMFE in Corollary 3.1, one may alternatively write the publisher's problem in terms of multipliers rather than bidding functions, and obtain the following Mathematical Program with Equilibrium Constraints (MPEC)

$$\begin{aligned} \max_{r,\eta,\iota} \quad & \lambda \sum_{\theta \in \boldsymbol{\Theta}} p_\theta \alpha_\theta \eta s_\theta G_\theta(\boldsymbol{\mu}, r) - \eta c I(\boldsymbol{\mu}, r) \\ \text{s.t.} \quad & \mu_\theta \geq 0 \quad \perp \quad \alpha_\theta \eta s_\theta G_\theta(\boldsymbol{\mu}, r) \leq b_\theta, \quad \forall \theta \in \Theta, \end{aligned} \tag{3.6}$$

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

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

3.4 Auction Design and Allocation Decisions

In this section we focus on the publisher's profit maximization problem. The framework developed in the previous sections can be applied to general inputs with regard to advertiser heterogeneity and market parameters as one could solve the resulting MPEC profit maximization problem. However, to gain further insights on the publisher's problem, and the trade-offs at hand, we first focus on the case of homogeneous bidders. In this case, quasi-closed form solutions may be obtained for the publisher's optimal decisions. In §3.4.2, we illustrate numerically how some of these insights may generalize to the case of heterogeneous bidders, and help better understand the latter.

3.4.1 The Case of Homogeneous Advertisers

We first consider the case in which all advertisers have a fixed budget b, stay in the market for a deterministic time s and share the same matching probability α and valuation parameter γ. That is, bidders are homogeneous and the space of types $\boldsymbol{\Theta}$ is a singleton. By Proposition (3.2) we know that in this case a unique FMFE exists and we can characterize it in quasi-closed form. We leverage this result to study the publisher's decisions. Throughout this section, we drop the dependence on θ. In the following we denote by $h_v(x) = f_v(x)/\bar{F}_v(x)$ the *failure rate* of the advertisers values (who have a common distribution), and by $\xi_v(x) = xh_v(x)$ the *generalized failure rate* of the values. We assume that values possess strictly increasing generalized failure rates (IGFR). This assumption is common in the pricing and auction theory literature, and many distributions satisfy this condition (see, e.g., Myerson (1981) and Lariviere (2006)).[7]

3.4.1.1 Optimal Reserve Price

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

Theorem 3.3 (Optimal reserve price)**.** *Suppose that η and ι are fixed. If $\alpha\eta s G_0(r_c^*) < b$, then r_c^* is the unique optimal reserve price. If $\alpha\eta s G_0(r_c^*) \geq b$, then the unique optimal reserve price is $\bar{r} = \sup \mathcal{R}^*$, where $\mathcal{R}^* = \{r : \alpha\eta s G_0(r) \geq b\}$. Furthermore, in the FMFE induced by the optimal reserve price, advertisers bid truthfully.*

The optimal reserve price admits a closed-form expression that highlights how it balances various effects. The expected expenditure for a bidder over its campaign when all bidders are truthful evaluated at r_c^*, $\alpha\eta s G_0(r_c^*)$, plays a key role in the result. In fact, when the budget is large in the sense that advertisers do not deplete their budget in expectation when the reserve price is r_c^* ($\alpha\eta s G_0(r_c^*) \leq b$), then it is expected that r_c^* should still be optimal in our setting. Intuitively, if the budget does not bind, the auctions decouple into independent second price auctions. When, however, $\alpha\eta s G_0(r_c^*) > b$, advertisers shade their values when the reserve price is r_c^*. In the proof, we show that in this case the optimal reserve price must be in $\mathcal{R}^*$, that is,

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

it must induce bidders to deplete their budgets in expectations. For all such reserve prices, the revenue rate for the publisher is given by λb and this is the maximum revenue rate she can extract from advertisers. Hence, recalling the objective value (3.6) of the publisher, the optimal reserve price must be the value $r \in \mathcal{R}^*$ that minimizes the probability of selling an impression in the exchange, and therefore the opportunity cost. Increasing the reserve price has two effects on this probability: (1) the *direct effect*: assuming advertiser's strategies do not change, an increase of the reserve price decreases the probability of selling an impression in the exchange; and (2) the *indirect effect*: a change in the reserve price also alters the strategies of the advertisers through the induced FMFE. In the proof we show that the direct effect is dominant, implying that $\bar{r} = \sup \mathcal{R}^*$, that minimizes the opportunity cost in $\mathcal{R}^*$, is optimal.

It is worthwhile to compare the result above with the work in one-shot auctions with budget constraints. In the case of a common budget for all bidders, authors have typically found that budget constraints decrease the optimal reserve price relative to the setting without budget constraints (see Laffont and Robert (1996) and Maskin (2000)). The reason is that with budget constraints the reserve price is less effective in extracting rents of higher valuation types; hence, when trading-off higher revenues conditional on a sale taking place with an increase in the probability of a sale, the latter has more weight than in the absence of budgets. In our case, instead, the optimal reserve price with budget constraints is larger or equal than r_c^*. In fact, the optimal reserve price is $\max\{\bar{r}, r_c^*\}$, because one can show that $\bar{r} \geq r_c^*$ if and only if $\alpha\eta s G_0(r_c^*) \geq b$. The difference with the one-shot auction is that the budget constraint is imposed over a large set of auctions as opposed to having a constraint per auction, leading to a different trade-off for the publisher. Indeed, when the budget constraint binds, the reserve price does not affect expected revenues, the publisher is already extracting all budgets from the bidders. Therefore, the only role of the reserve price becomes one of reducing the opportunity cost by decreasing the probability of a sale. As we saw, this is achieved by increasing the reserve price while still extracting the maximum amount of revenues.

3.4.1.2 Optimal Allocation of Impressions

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

Theorem 3.4 (Optimal allocation of impressions)**.** *Suppose that r and ι are fixed. If $cI_0(r) \geq \alpha\lambda s G_0(r)$, then the publisher is better off not participating in the exchange, and the optimal rate of impressions is $\eta^* = 0$. If $cI_0(r) < \alpha\lambda s G_0(r)$, then the publisher stands to gain from participating in the exchange, and the optimal allocation of impressions to the exchange is $\eta^*(r) = \min\{\eta^0(r), \bar{\eta}\}$ where $\eta^0(r) = b/(\alpha s G_0(r))$.*

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

Initially, when the supply is sufficiently small, bidders do not deplete their budget and hence are truthful (cf. Proposition 3.2). In such a region, increasing the allocation of impressions yields larger revenues for the publisher, which is in line with intuition. However, when the rate of impressions is high enough, all advertisers deplete their budgets by the end of their campaign and no further revenue may be extracted by allocating more impressions to the exchange, which corresponds to the market being "saturated". The smallest rate at which saturation takes place is exactly given by $\eta^0(r) = b/(\alpha s G_0(r))$. At that rate, advertisers are truthful; beyond that rate, bidders start shading their bids. Allocating further impressions to the exchange does not yield additional revenues since advertisers are already spending all their budget and hence the key resides in understanding the impact of an increase in supply on the opportunity cost. When the publisher increases the impression rate above market saturation, there are again two effects to consider; (1) the direct effect: sending more impressions to the exchange directly increases the publisher's opportunity cost if these impressions are won; and (2) the indirect effect: as more impressions are available, advertisers shade their bids more and in the presence of a reserve price, the probability of a sell and the opportunity cost decrease. In the proof we show that the direct effect dominates and increasing the rate above market saturation is suboptimal. Thus, the optimal rate of impression is the minimum of $\eta^0(r)$ and $\bar{\eta}$.

Next we characterize the optimal decision of the publisher when she optimizes over both the allocation of impressions and the reserve price. In contrast to Theorem 3.4, when jointly optimizing over the reserve price and the rate of impressions, the publisher is always better off allocating $\bar{\eta}$ impressions to the exchange. In this case, because the reserve price optimization considers the alternative channel, the exchange becomes a "free option" that is always worth testing.

Corollary 3.2 (Joint allocation and reserve price optimization)**.** *Suppose that ι is fixed. The optimal decision for the publisher is to send all impressions to the exchange, and set the reserve price according to Theorem 3.3. That is, the unique optimal rate of impressions is $\eta^* = \bar{\eta}$, and the optimal reserve price is equal to $\max\{r_c^*, \bar{r}(\bar{\eta})\}$, where $\bar{r}(\eta) = \sup \mathcal{R}^*(\eta)$ and $\mathcal{R}^*(\eta) = \{r : \alpha\eta s G_0(r) \geq b\}$.*

We study the joint optimization problem by partitioning it in two stages. In the first stage the publisher looks for the allocation of impressions that maximizes the optimal value of the second-stage problem, obtained from optimizing over reserve prices. Exploiting Theorem 3.3 to characterize the maximum profit over reserve prices, we show that the second-stage objective value is increasing with the rate of impressions. Therefore, when jointly optimizing over reserve prices and the rate of impressions, the publisher is better off allocating $\bar{\eta}$ impressions to the exchange.

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

3.4.1.3 Optimal Disclosure of Information

When the publisher posts an impression in the exchange she can decide which user information (if any) to disclose to the advertisers. The publisher may decide to disclose, e.g., the content of the web-page, user geographical location, user demographics, or cookie-based behavioral information (which allows bidders to track the user's past activity in the web). On the one hand, each additional level of targeting may reduce the probability that an advertiser matches with a given user, because more criteria need to be satisfied to do so. This may lead to *thinner* markets and could decrease the publisher's revenue. On the other hand, more information provides better targeting that results in higher valuations and higher bids, conditional on a match. Our FMFE framework can be used to analyze (numerically or analytically) different settings regarding the impact of information disclosure on the publisher's profit. Below, we illustrate this through a particular stylized model for information disclosure.[8]

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

$$\begin{aligned} F_{v(\alpha)}(x) &= F_v(x/\sigma(\alpha)), \\ \alpha\sigma(\alpha) &= 1, \text{ for all } \alpha \in (0, 1]. \end{aligned} \tag{3.7}$$

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

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

Theorem 3.5 (Joint information disclosure and reserve price optimization)**.** *Suppose that η is fixed and that advertisers' valuations follow* (3.7). *When the publisher reacts to thinner markets by setting an appropriate reserve price, then disclosing more information improves the profit, that is, the publisher's profit $\Pi(\alpha) = \max_{r\geq 0} \Pi(\mu(r, \alpha), r, \alpha)$ is non-increasing in α.*

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

We provide the main ideas behind the argument. First, in view of Theorem 3.3, advertisers bid truthfully at the optimal reserve price, and there is no need to take into account the shading of bids. Hence, when changing α, there is no *indirect effect*, and it is enough to show that the *direct effect* of decreasing α (corresponding to the impact of having thinner markets but larger valuations) increases the publisher's profits. Now, there are two cases to consider. When the expenditure at $r_c^*(\alpha)$ does not exceed the budget, we have that $r_c^*(\alpha)$ is the optimal reserve price. In the proof, we show that in this case profits increase as α decreases. When the expenditure at $r_c^*(\alpha)$ exceeds the budget, then the publisher prices at $\bar{r}(\alpha) = \sup\{r \geq 0 : \alpha\eta s G_0(r, \alpha) \geq b\}$, and advertisers deplete their budgets in expectation. Here, we show that as the matching probability decreases, the optimal reserve price $\bar{r}(\alpha)$ changes, resulting in more impressions returned to the publisher, and a lower opportunity cost, increasing publishers' profits.

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

Suppose that the publisher is disclosing an initial level of information that attains a matching probability α_0, she is pricing at the optimal reserve price $r_c^*(\alpha_0)$, and the advertiser's expenditure does not exceed the budget. Consider the publisher's profit as a function of the matching probability when the reserve price is *not* adjusted, which is given by $\Pi(r_c^*(\alpha_0), \alpha)$ (we dropped the dependence on $\mu(r, \alpha)$ to simplify the notation). One can show that $\Pi(r_c^*(\alpha_0), \alpha)$ is locally non-increasing near α_0, that is, a small increment in the disclosure of information actually increases profits.[9] Nonetheless, it is possible to show that disclosing more information and further decreasing α may cause profits to decrease.

3.4.2 Numerical Results for the Case of Heterogeneous Advertisers

While it was possible to obtain essentially closed-form solutions for the publisher's decisions in the case of homogeneous advertisers, it is not, in general, possible to derive such results for the case of heterogeneous advertisers. However, one may always numerically analyze the impact of the publisher's decisions on the advertisers' equilibrium outcome under different scenarios by solving for the FMFE using the characterization in Corollary 3.1 for different auction parameters. In this

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

section, we conduct a series of numerical experiments to explore the impact of the allocation of impressions and the reserve price on the publisher's profit to conduct some robustness check on some of the conclusions of §3.4.1.

The setup is as follows. We consider randomly generated instances with a heterogeneous population of advertisers with five types. Budgets for each type are sampled from a discrete uniform distribution with support $\{1, 2, ..., 10\}$. Additionally, we experiment with the proportion of these types by choosing the probabilities p_θ of an arriving advertiser being of type θ uniformly from the probability simplex. Throughout the experiments we fix the matching probability $\alpha = 0.1$ and the campaign length to $s = 10$, but select the arrival rate λ uniformly in $[1, 5]$ so that the average number of matching bidders in an auction $\alpha\lambda s$ varies from 1 to 5. Finally, we assume that values are exponentially distributed with each type's mean independently and uniformly sampled from $[1, 5]$ (the supports of valuations are truncated to $[0, 10]$). From the perspective of the publisher, we study scenarios with different opportunity costs c for the alternative channel, by choosing the cost uniformly from $[1, 5]$. In total, we consider 150 different scenarios. Additionally, for each scenario, we experiment with increasing levels of impressions allocated to the exchange, as given by η.[10]

For each combination of parameters and value of η, we optimize problem (3.6) over the reserve price; and compute the optimal reserve price r^* and the optimal profit for the publisher Π^*. Then, we compute the latter quantities for a grid of values of the rate of impressions η. An important conclusion in §3.4.1 was that it is optimal for the publisher to send all the impressions to the exchange as long as the reserve price was properly adjusted. For all sampled parameters, we indeed find that, when the publisher accounts for the optimal reserve price, her profit is increasing with the allocation of impressions. As before, the publisher is better off sending all impressions to the exchange, even in the presence of an alternative channel.

To explore an illustrative example in further detail, Figure 3.1 depicts, for a given set of parameters with two types, the optimal publisher's profit (a), the equilibrium advertisers' expenditures (b) and multipliers (c) at the optimal solution, as well as the optimal reserve price (d) (all as a function of the allocation of impressions to the exchange η). Notice that when the publisher prices optimally, the high-budget type always bids truthfully. However, this is not necessarily true for the low-budget type. As opposed to the homogeneous case, now the publisher cannot perfectly discriminate between types and for some levels of supply, low-type advertisers will shade their bids under the optimal reserve price.

Focusing on the optimal reserve price, as expected, we observe that advertisers do not have a

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

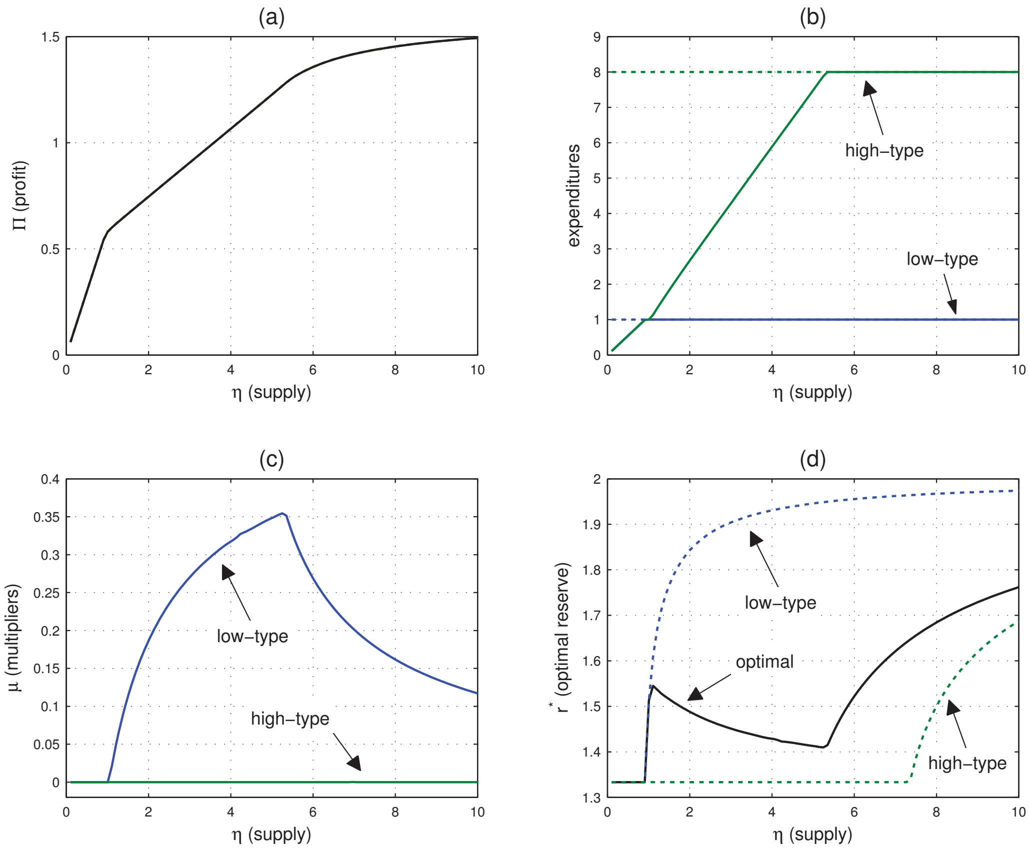

Figure 3.1: Optimal profit, expenditures, multipliers and reserve price as a function of the rate of impressions for an instance with $\alpha = 0.1$, $\lambda = 1$, $s = 40$, Unif$[0, 2]$ distribution, $c = \frac{2}{3}$, $b = (1, 8)$, and $p = (\frac{1}{5}, \frac{4}{5})$. For illustration purposes we only consider two types and different parameters than above. (a) the solid line denotes the optimal profit. (b) the solid lines correspond to the campaign expenditures for each type, while the dashed lines denote the budgets. (c) equilibrium multipliers as a function of the allocation of impressions. (d) the solid line corresponds to optimal reserve price, while the dashed lines denote the optimal prices when all advertisers share the same budget (either b_1 or b_2).

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

In our numerical experiments, a similar structure and tradeoff appears when there are more than two types of advertisers with different budgets in the population, with one new kink in the optimal reserve price for each additional type. As previously mentioned, in these experiments, we also observed that the publisher's profit rate is increasing in the rate of impressions. As an illustration, in Figure 3.2 we provide similar plots than the ones above for a market with five advertiser types.

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

3.5 Conclusions

Overall, our results provide a new approach to study Ad Exchange markets and the publishers' decisions. The techniques developed build on two fairly distinct streams of literature, revenue management and mean field models and are likely to have additional applications. The sharp results regarding the publisher's decisions could inform how these markets are designed in practice.

At the same time, our framework opens up the door to study a range of other relevant issues in this space. For example, one interesting avenue for future work may be to study the impact of Ad

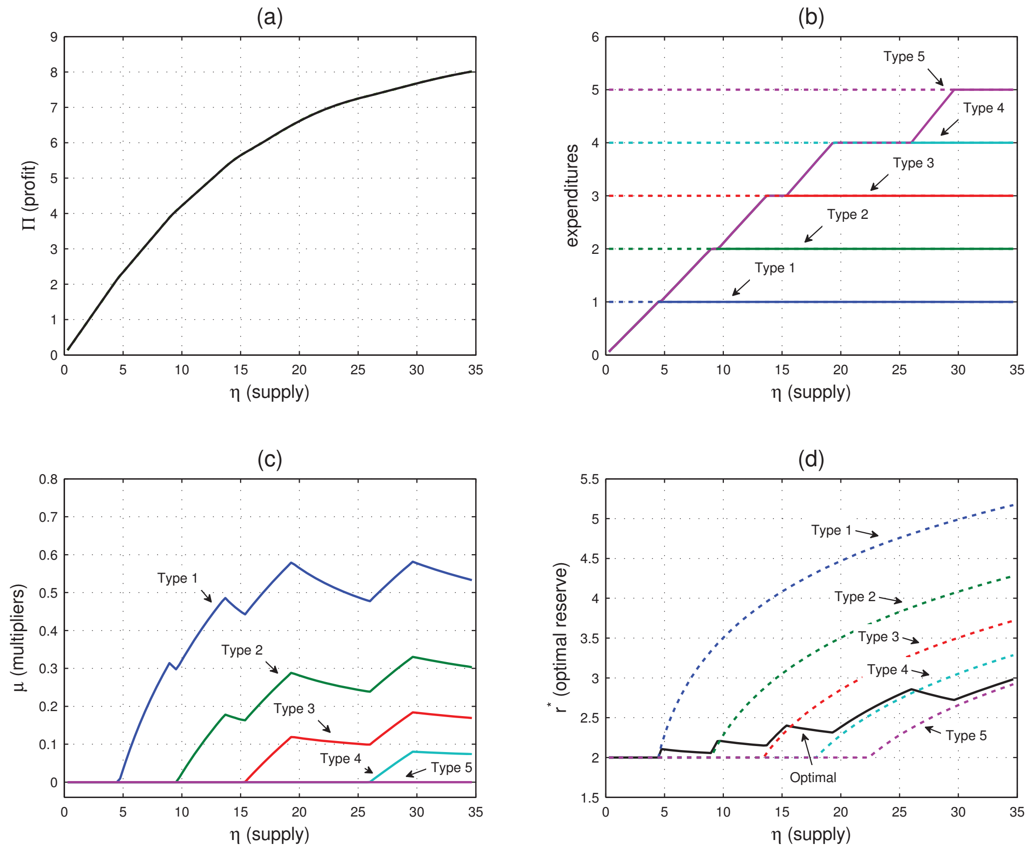

Figure 3.2: Optimal profit, expenditures, multipliers and reserve price as a function of the rate of impressions for an instance with $\alpha = 0.1$, $\lambda = 4$, $s = 10$, Exponential(1) distribution of values, $c = 1$, $b_\theta = 1, 2, 3, 4, 5$, and types probabilities are equal to $1/5$ each.

networks, that aggregate bids from different advertisers and bid on their behalf, on the resulting competitive landscape and auction design decisions. Similarly, another interesting direction to pursue is to incorporate common advertisers' values and analyze the impact of cherry-picking and adverse selection. Finally, our framework and its potential extensions can provide a possible structural model for bidding behavior in exchanges, and open the door to pursue an econometric study using transactional data in exchanges, a direction we are currently exploring.

Chapter 4

Advertiser Competition in AdX: Approximation Results

> The material presented in this chapter is based on the technical report Balseiro et al. (2012a) co-authored with Omar Besbes and Gabriel Weintraub.

In this chapter 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.

We provide two complementary proofs for the approximation result under different demand models. First, we study the case of *synchronous campaigns*, that is, when all campaigns start at the same time and finish simultaneously. This simpler model captures, for example, the case when advertisers have periodic (daily or weekly) budgets. In this setting we are able to show our result under the assumption that the number of advertisers in the market grows slower (in the little-o sense) than the number of matching auctions, which is an appealing regime that applies to most current markets. Second, we study the more general case of *asynchronous campaigns*, that is, when advertisers arrive to the market at random points in time and campaigns overlap. This is the model presented in Chapter 3. The complexity of this model precludes the possibility of applying traditional RM techniques, and thus motivated we develop a novel framework based on more elaborate mean field techniques. The main limitation of our result is that the scaling is more restrictive than in the synchronous case. To wit, our proof holds under the assumption that the number of advertisers in the market grows exponentially faster than the number of auctions.

We conjecture, however, that the family of scalings under which our approximation result is valid is broader.

4.1 Approximation Result for Synchronous Campaigns

In this section we show that the FMFE approximates the rational behavior of bidders in large markets by considering the simpler model of *synchronous campaigns*, that is, when all campaigns start at the same time and finish simultaneously. This model captures, for example, the case when advertisers have periodic (daily or weekly) budgets. We start by describing the synchronous model and defining the FMFE concept for this setting. We continue by formally stating the result, and describing the intuition and main steps of the proof.

Synchronous Campaign Model. There is a fixed number of agents in the market, which we denote by K. Campaigns are synchronous, that is, they all start at time 0, finish at a common time s, and neither arrivals nor departures are allowed during the time horizon $[0, s]$. Agents are indexed by $k = 1, \ldots, K$. The k^{th} agent is characterized by a type vector $\theta_k = (b_k, \alpha_k, \gamma_k) \in \mathbb{R}^3$, where b_k denotes the budget, α_k the probability that the agent participates in an auction, and $F_v(\cdot; \gamma_k)$ the cumulative distribution function of valuations. Types are publicly known and revealed at the beginning of the horizon. On the supply side, impressions arrive according to a Poisson process with intensity η. Following the notation in Chapter 3 we index the sequence of arriving impressions by $n \geq 1$, and we denote the sequence of arrival times by $\{t_n\}_{n \geq 1}$. Additionally, we let $M_{n,k} = 1$ indicate that the k^{th} agent participates in the n^{th} auction.

In this setting the expected expenditure function of the k^{th} advertiser of a single auction when advertisers shade their bids according to a vector of multipliers $\boldsymbol{\mu} \in \mathbb{R}^K_+$, denoted by $G_k(\boldsymbol{\mu})$, is given

$$G_k(\boldsymbol{\mu}) \triangleq \mathbb{E}\left[\mathbf{1}\{(1 + \mu_k) D_{-k} \leq V_k\} D_{-k}\right],$$

with the maximum competing bid given by $D_{-k} = \max_{i \neq k, M_i = 1} \left\{V_i/(1 + \mu_i)\right\} \vee r$, where we ignored the index n to simplify the notation. A similar analysis to the one performed in the case of asynchronous campaigns yields that the vector of multipliers in the FMFE can be characterized as the solution of the following Non-linear Complementarity Problem (NCP)

$$\alpha_k \eta s G_k(\boldsymbol{\mu}) \leq b_k \quad \perp \quad \mu_k \geq 0, \quad \forall k = 1, \ldots, K. \tag{4.1}$$

We shall prove the approximation result by considering a sequence of markets with increasing size. On the demand side, the number of advertisers and their budgets are allowed to increase. On the supply side, the number of impressions is increased so that the expected number of auctions a bidder participates in grows at the same rate as her budget, while the expected number of bidders in each auction remains constant. Instead of explicitly indexing the different markets in the sequence, we prove our result for any market satisfying the following set of assumptions on the primitives.

Assumption 4.1. *There exists non-negative bounded constants $\underline{g}, \bar{g}, \underline{z}, \bar{a}$ independent of the scaling such that:*

1. *The ratio of budget to number of matching auctions is bounded from above and below across advertisers, i.e., for all advertiser k we have that $b_k/(\alpha_k \eta s) \in [\underline{g}, \bar{g}]$.*

2. *The ratio of matching probabilities of any two advertisers is uniformly bounded across advertisers, that is, for every pair of advertisers $k \neq i$ we have that $\alpha_k/\alpha_i \leq \bar{a}$.*

3. *The expected expenditure per auction when advertisers bid truthfully is uniformly bounded from below across advertisers, i.e., for all advertiser k we have that $G_k(\mathbf{0}) \geq \underline{z}$.*

The first two assumptions guarantee that no advertiser has an excessive market power by limiting budgets and the number for matching auctions. The third assumption implies that, in equilibrium, all advertisers have a positive expected expenditure per auction so that no advertiser is systematically outbid in equilibrium. Thus, these assumptions guarantee that there is no dominant advertiser in the market. It does allow, however, for heterogeneity across the advertisers types.

Approximation Result. We denote the k^{th} advertiser *history* up to time t by $h_k(t)$. The history encapsulates all available information up to time t including the advertisers' types; the realizations of her values up to that time; her bids; the budgets of all advertisers; and the result of every auction. 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.

We study the expected payoff of advertiser k when she implements a strategy $\beta \in \mathbb{B}$ and all other advertisers follow the FMFE strategies $\boldsymbol{\beta}^{\mathrm{F}}$. This expected payoff is denoted by $J_k(\beta, \boldsymbol{\beta}^{\mathrm{F}}_{-k})$, where the expectation is taken over the actual market process. In this notation, $J_k(\beta^{\mathrm{F}}_k, \boldsymbol{\beta}^{\mathrm{F}}_{-k})$ measures the *actual* expected payoff of the FMFE strategy for the advertisers in the exchange. We have the following result.

Theorem 4.1. *Suppose that Assumption 4.1 holds. Consider a market in which all bidders, except the k^{th} bidder, follow FMFE strategies $\boldsymbol{\beta}^{\mathrm{F}}$. Suppose that the k^{th} advertiser deviates and* *implements a non-anticipating and adaptive strategy $\beta \in \mathbb{B}$. The expected payoff of this deviation with respect to the FMFE strategy β^{F} satisfies*

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

where $O(\cdot)$ stands for Landau's big-O notation.

The bound given by the result converges to zero if the number of advertisers grows slower than the number of auctions the advertiser participates in, or more formally if $K = o(\alpha\eta s)$. This is a natural regime because in most markets the number of auctions a bidder participates in typically much larger than the number of competitors. The factor $(\alpha\eta s)^{-1/2}$ in the second error term is introduced by the open-loop nature of the fluid-based strategies in the FMFE. This term is reminiscent of error bounds on the performance of a fluid-based strategy competing against an exogenous environment typically obtained in the RM literature. The main difference with the RM results is the factor $\sqrt{K}$ in the error term that accounts for the impact of an endogenous competitive landscape.

Though not stated, the result holds even when the advertiser might deviate to hindsight strategies that have perfect knowledge of future realizations of bids and values. Hindsight strategies can underbid competitors and deplete their budgets without incurring any cost, and thus have significant advantages over non-anticipating policies. We conjecture, however, that the actual gap with respect to the optimal non-anticipating policy is smaller in practice.

Outline. We prove the result in two steps. First, we lower bound the expected performance of the k^{th} advertiser when all advertisers (including herself) implement the FMFE strategy in terms of the objective value of the fluid problem (3.1). Second, we upper bound the expected payoff of *any* strategy the k^{th} advertiser may implement when the remaining implement the FMFE strategies via a hindsight bound.

Proposition 4.1 (Lower Bound)**.** *Suppose that Assumption 4.1 holds and all advertisers implement FMFE strategies $\boldsymbol{\beta}^{\mathrm{F}}$. The expected payoff of the k^{th} advertiser is lower bounded by*

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

where $\bar{J}^{\mathrm{F}}_k \triangleq J^{\mathrm{F}}_k / (\alpha_k \eta s)$ is the normalized objective value of the fluid problem (3.1)*.*

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

Proposition 4.2 (Upper Bound). *Suppose that Assumption 4.1 holds and all advertisers implement FMFE strategies $\boldsymbol{\beta}^{\mathrm{F}}$, and the k^{th} advertiser implements an alternative strategy $\beta \in \mathbb{B}$. The expected payoff of the k^{th} advertiser is upper bounded by*

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

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

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

4.2 Approximation Result for Asynchronous Campaigns

In this section we prove the approximation result in the general model with *asynchronous campaigns*, that is, when advertisers arrive to the market at random points in time and campaigns overlap. This is the model presented in Chapter 3. 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. (3.1) and (3.2) in Chapter 3). Thus, irrespectively of the scaling, the FMFE strategy is given by $\boldsymbol{\beta}^{\mathrm{F}} = \{\beta_\theta^{\mathrm{F}}\}_{\theta \in \boldsymbol{\Theta}}$ 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 $\boldsymbol{\beta}^{\mathrm{F}}$. This expected payoff is denoted by $J_\theta^\kappa(\beta, \boldsymbol{\beta}^{\mathrm{F}})$, 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, $J_\theta^\kappa(\beta_\theta^{\mathrm{F}}, \boldsymbol{\beta}^{\mathrm{F}})$ measures the *actual* expected payoff of the FMFE strategy for the advertisers in the exchange.[1] We have the following result.

Theorem 4.2. *Suppose that $r \in (0, \overline{V})$ 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 $\boldsymbol{\beta}^{\mathrm{F}}$. Suppose that a θ-type advertiser (the zeroth bidder), upon arrival to the market, deviates and*

[1] Note that this performance metric may differ from the FMFE value function, given by the objective value of the approximation problem $J_\theta^{\mathrm{F}}(F_d)$ given in (3.1) in the main paper.

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 $\beta_\theta^{\mathrm{F}}$ satisfies

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

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 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

[2] We discuss the nature of this distribution in Section 4.4 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.

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 4.2 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 4.2) 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 (3.1) 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 $\mathbb{P}_{\hat{\Theta}}\{\theta\}$ 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 4.2 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 § 4.3 and § 4.4. The proofs of the results are provided in § C.

4.3 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 $\boldsymbol{\beta}^{\mathrm{F}}$. 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 $X_\theta^{\mathrm{MF}}(t) = (b_\theta(t), s_\theta(t)) \in \mathbb{R}_+^2$ 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 $X_\theta^{\mathrm{MF}} = \{X_\theta^{\mathrm{MF}}(t)\}_{t\in[0,s_\theta]}$ is Markov.

We next describe the evolution of the continuous time Markov process X_θ^{MF}. Initially, we have that $X_\theta^{\mathrm{MF}}(0) = (b_\theta, s_\theta)$. 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 Unifrom 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 $w_{n,k} = \beta_{\theta_{n,k}}^{\mathrm{F}}(v_{n,k})\mathbf{1}\{b_{n,k} > 0\}$, 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 $w_{n,0} = \beta_\theta^{\mathrm{F}}(v_{n,0})\mathbf{1}\{b(t_{\theta,n}^-) > 0\}$, 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

[3]In 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 $w_{n,k} = \min\{b_{n,k}, \beta_{\theta_{n,k}}^{\mathrm{F}}(v_{n,k})\}$.

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 $X_\theta^{\mathrm{MF}}(\mathbb{P}_e) = \{X_\theta^{\mathrm{MF}}(t;\mathbb{P}_e)\}_{t\in[0,s_\theta]}$. 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 $X_\theta^{\mathrm{MF}}(\mathbb{P}_e)$ 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 2 (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*

$$\mathbb{P}_e\{\mathcal{X},\theta\} = \mathbb{P}\{X_\theta^{\mathrm{MF}}(U[0,s_\theta];\mathbb{P}_e) \in \mathcal{X}; \hat{\Theta} = \theta\} \tag{4.2}$$

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.

4.3.1 Existence of a consistent BMFM

The consistency equation (4.2) 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 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 $\mathbf{q} = \{q_\theta\}_{\theta\in\mathbf{\Theta}}$ be a vector of active probabilities. Since values and types are independent, equation (4.2) 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,

$$q_\theta = \mathbb{P}\{b_\theta(U[0,s_\theta];\mathbf{q}) > 0\} = \frac{1}{s_\theta}\int_0^{s_\theta} \mathbb{P}\{b_\theta(t;\mathbf{q}) > 0\}\mathrm{d}t = \mathbb{E}\left[\frac{1}{s_\theta}\int_0^{s_\theta} \mathbf{1}\{b_\theta(t;\mathbf{q}) > 0\}\mathrm{d}t\right]. \tag{4.3}$$

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 $\mathbf{q}^\kappa$ satisfying the fixed-point equation (4.3).

Proposition 4.3. *For every scaling κ, there exists a vector of active probabilities $\mathbf{q}^\kappa$ satisfying the consistency equation* (4.3)*. Moreover, the consistent probability distribution of extended states is given by $\mathbb{P}_e^\kappa\{\mathcal{X}, \theta\} = \mathbb{P}\{X_\theta^{\mathrm{MF}(\kappa)}(U[0, s_\theta^\kappa]; \mathbf{q}^\kappa) \in \mathcal{X}; \hat{\Theta} = \theta\}$.*

To prove this proposition we first show that the right-hand side of equation 4.3 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.

4.3.2 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 4.4. *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 $\lim_{\kappa \to \infty} \|1 - \mathbf{q}^\kappa\|_\infty = 0$. Additionally, 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.*

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.

[4] We note that this result and the results in Subsection 4.3.3 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 4.4.

4.3.3 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 $J_\theta^{\mathrm{MF}(\kappa)}(\beta, \boldsymbol{\beta}^{\mathrm{F}})$ 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 $\boldsymbol{\beta}^{\mathrm{F}}$.

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 (3.1) of the main paper as $\bar{J}_\theta^{\mathrm{F}}(F_d) \triangleq J_\theta^{\mathrm{F}}(F_d)/(\alpha_\theta \eta s_\theta)$. We also define F_d as the distribution of the maximum bid in the FMFE.

Proposition 4.5 (Lower Bound)**.** *Consider any consistent BMFM with scaling κ in which all competitor bidders follow the FMFE strategy $\boldsymbol{\beta}^{\mathrm{F}}$. Suppose that the zeroth advertiser of type θ implements the FMFE strategy $\beta_\theta^{\mathrm{F}}$. The expected payoff of the zeroth advertiser in the BMFM, denoted by $J_\theta^{\mathrm{MF}(\kappa)}(\beta^{\mathrm{F}}, \boldsymbol{\beta}^{\mathrm{F}})$, is lower bounded by*

$$\liminf_{\kappa\to\infty} \frac{1}{\alpha_\theta^\kappa \eta^\kappa s_\theta} J_\theta^{\mathrm{MF}(\kappa)}(\beta_\theta^{\mathrm{F}}, \boldsymbol{\beta}^{\mathrm{F}}) \geq \bar{J}_\theta^{\mathrm{F}}(F_d).$$

The intuition underlying the proof of this result relies heavily on Proposition 4.4. 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 4.4 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 $\bar{J}_\theta^{\mathrm{F}}(F_d^\kappa)$, which itself is close to $\bar{J}_\theta^{\mathrm{F}}(F_d)$ (by the second part of Proposition 4.4).

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

Proposition 4.6 (Upper Bound)**.** *Consider any consistent BMFM with scaling κ in which all competitor bidders follow the FMFE strategy $\boldsymbol{\beta}^{\mathrm{F}}$. Suppose that the zeroth advertiser of type θ implements an alternative strategy β^κ. The expected payoff of the zeroth advertiser, denoted by $J_\theta^{\mathrm{MF}(\kappa)}(\beta^\kappa, \boldsymbol{\beta}^{\mathrm{F}})$, is upper bounded by*

$$\limsup_{\kappa\to\infty} \frac{1}{\alpha_\theta^\kappa \eta^\kappa s_\theta} J_\theta^{\mathrm{MF}(\kappa)}(\beta^\kappa, \boldsymbol{\beta}^{\mathrm{F}}) \leq \bar{J}_\theta^{\mathrm{F}}(F_d).$$

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 4.4 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.

4.4 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 $J_\theta^\kappa(\beta^\kappa, \boldsymbol{\beta}^{\mathrm{F}})$, to the performance of the same strategy in the BMFM, denoted by $J_\theta^{\mathrm{MF}(\kappa)}(\beta^\kappa, \boldsymbol{\beta}^{\mathrm{F}})$. 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 $\mathbb{P}_e^\kappa$ of a consistent BMFM.

Proposition 4.7. *Consider a $\kappa-$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 $\mathbb{P}_e^\kappa$. 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*

$$\lim_{\kappa\to\infty} \frac{1}{\alpha_\theta^\kappa \eta^\kappa s_\theta} \left| J_\theta^\kappa(\beta^\kappa, \boldsymbol{\beta}^{\mathrm{F}}) - J_\theta^{\mathrm{MF}(\kappa)}(\beta^\kappa, \boldsymbol{\beta}^{\mathrm{F}}) \right| = 0.$$

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 C.3 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.4 we formally show that our AdX market can be modeled as such system.

Bibliography

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

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

Agrawal, S., Wang, Z. and Ye, Y. (2009), A dynamic near-optimal algorithm for online linear programming. Working paper posted at http://www.stanford.edu/ yyye/.

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

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

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

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

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

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

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

Balseiro, S. R., Besbes, O. and Weintraub, G. Y. (2012*b*), Auctions for online display advertising exchanges: Approximations and design. Working Paper.

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

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

Bertsekas, D. P. (2000), *Dynamic Programming and Optimal Control*, 2nd edn, Athena Scientific.

Billingsley, P. (1968), *Convergence of Probability Measures*, John Wiley & Sons Inc.

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

Boyd, S. and Mutapcic, A. (2008), Subgradient methods. Notes for EE364b, Stanford University, Winter 2006-07.

Casella, G. and Berger, R. L. (2002), *Statistical Inference; 2nd ed.*, Duxbury advanced series, Wadsworth, New Delhi.

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

Chauvin, B. and Giroux, G. (1990), 'A geometric rate of convergence to the equilibrium for boltzmann processes with multiple particle interactions', *Journal of Applied Probability* **27**(3), 510–520.

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

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

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

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

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

Devenur, N. R. and Hayes, T. P. (2009), The adwords problem: online keyword matching with budgeted bidders under random permutations, *in* 'Proceedings of the 10th ACM conference on Electronic commerce', EC '09, ACM, pp. 71–78.

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

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

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

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

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

Feldman, J., Henzinger, M., Korula, N., Mirrokni, V. S. and Stein, C. (2010), Online stochastic packing applied to display ad allocation, *in* 'Proceedings of the 18th annual European conference on Algorithms: Part I', ESA'10, Springer-Verlag, pp. 182–194.

Feldman, J., Korula, N., Mirrokni, V., Muthukrishnan, S. and Pál, M. (2009), Online ad assignment with free disposal, *in* 'Proceedings of the 5th International Workshop on Internet and Network Economics', WINE '09, Springer-Verlag, pp. 374–385.

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

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

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

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

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

Graham, C. (1992), 'Nonlinear diffusion with jumps', *Annales de l'I.H.P. Probabilits 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.

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

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

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

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

Jasin, S. and Kumar, S. (2010), A re-solving heuristic with bounded revenue loss for network revenue management with customer choice. Working paper posted at http://www.stanford.edu/ skumar/.

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

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

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

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

Levi, R. and Radovanovic, A. (2010), 'Provably near-optimal lp-based policies for revenue management in systems with reusable resources', *Operations Research* **58**(2), 503–507.

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

Levin, Y., McGill, J. and Nediak, M. (2008), 'Risk in revenue management and dynamic pricing', *Operations Research* **56**(2), 326–343.

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

Luenberger, D. G. (1969), *Optimization by Vector Space Methods*, 1st edn, John Wiley & Sons, Inc., New York, NY, USA.

Mansour, Y., Muthukrishnan, S. and Nisan, N. (2012), Doubleclick ad exchange auction. Working Paper.

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

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

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

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

Nazerzadeh, H., Saberi, A. and Vohra, R. (2009), Dynamic cost-per-action mechanisms and applications to online advertising. Working Paper.

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

Phillips, R. (2012), 'Efficient frontiers in revenue management', *Journal of Revenue and Pricing Management* **11**(4), 371–385.

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

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

Shapiro, A. (1991), 'Asymptotic analysis of stochastic programs', *Annals of Operations Research* **30**, 169–186.

Shapiro, A., Dentcheva, D. and Ruszczynski, A. (2009), *Lectures on Stochastic Programming: Modeling and Theory.*, MOS-SIAM Series on Optimization 9, Society for Industrial and Applied Mathematics (SIAM).

Simpson, R. (1989), Using network flow techniques to find shadow prices for market demands and seat. Flight Transportation Laboratory Memorandum M89-1.

Talluri, K. T. and van Ryzin, G. J. (2004), *The Theory and Practice of Revenue Management*, International Series in Operations Research & Management Science, Vol. 68, Springer.

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

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

Vee, E., Vassilvitskii, S. and Shanmugasundaram, J. (2010), Optimal online assignment with forecasts, *in* 'Proceedings of the 11th ACM conference on Electronic commerce', EC '10, ACM, pp. 109–118.

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

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

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

Appendices

Appendix A

Appendix to Chapter 2

A.1 Proofs of Statements

A.1.1 Proof of Proposition 2.1

In this proof we drop the dependence on the user attributes to simplify the notation. First, observe that for all c the objective function of (2.1) is concave and continuous in s, and the feasible set is compact. Hence, by Weierstrass Theorem the set of optimal solutions is non-empty and compact. Thus, both $R(c)$ and $s^*(c)$ are well-defined.

Second, $R(c) \geq c$ follows from letting $s = 0$. To see that $R(c)$ is non-increasing, let $c < c'$, and s^* be the optimal solution under cost c. Then, $R(c) = r(s^*) + (1 - s^*)c \leq r(s^*) + (1 - s^*)c' \leq R(c')$ where the first inequality follows because $s^* \leq 1$, and the second because no solution is better than the optimal. To see that $R(c) - c$ is non-increasing, let $c' < c$, and s^* be the optimal solution under cost c. Then, by a similar argument we get that $R(c') - c' \geq r(s^*) - s^*c' \geq r(s^*) - s^*c = R(c) - c$. Convexity follows in a similar way (this is a standard result).

Third, observe that the objective function of (2.1) is jointly continuous in s and c. Thus, by the Maximum Theorem $R(c)$ is continuous in c, and $s^*(c)$ is upper-hemicontinuous.

Finally, because $r(s) + (1 - s)c$ has decreasing differences in (s, c) and the feasible set is a lattice, by Topkis's Theorem $s^*(c)$ is non-increasing in c. The result for $p^*(c)$ follows from the fact that $\bar{F}^{-1}(s)$ is non-increasing in s. $\square$

A.1.2 Proof of Theorem 2.1

The optimality conditions of v for problem (2.4) imply that the directional derivative of $\psi(v)$ along any direction is greater or equal to zero. In particular for each advertiser $a \in \mathcal{A}$ it should the case that $\nabla_{\mathbf{1}_a} \psi(v) \geq 0$, and $\nabla_{-\mathbf{1}_a} \psi(v) \geq 0$. Applying proposition A.1 to both directions, together with the fact that there is zero probability of a tie occurring, we get that

$$\mathbb{E}\left[(1 - s^* (Q_a - v_a)\,; U)\, \mathbf{1}\{Q_a - v_a > Q_{a'} - v_{a'}\ \forall a' \in \mathcal{A}_0\}\right] = \rho_a,$$

and the result follows. □

A.1.3 Proof of Proposition 2.2

Let $\vec{s} = \{s_n(\cdot)\}_{n=1,\ldots,N}$ and $\vec{\imath} = \{i_n(\cdot)\}_{n=1,\ldots,N}$ be any feasible vectors of controls. Let $\bar{s}$ be the mean of the controls (in terms of prices $\bar{p}$ would be the generalized $\bar{F}$-mean), which is defined point-wise $\bar{s} = \frac{1}{N}\sum_{n=1}^{N} s_n$. Similarly, let $\bar{\imath}$ be such that $\bar{\imath}_a = \frac{1}{N}\sum_{n=1}^{N} i_{n,a}$ point-wise for all $a \in \mathcal{A}$. We will show that solution in which $(\bar{s}, \bar{\imath})$ are used for all impressions is a feasible control with greater or equal revenue than the original one.

First, for the feasibility of $(\bar{s}, \bar{\imath})$ observe that for each advertiser $a \in \mathcal{A}$

$$C_a = \sum_{n=1}^{N} \mathbb{E}_{U_n}\left[i_{n,a}(U_n)\right] = \mathbb{E}_U\left[\sum_{n=1}^{N} i_{n,a}(U)\right] = N\mathbb{E}_U\left[\bar{\imath}_a(U)\right],$$

where the first equation follows from the feasibility of $(\vec{p}, \vec{\imath})$, the second from the linearity of expectation and the stationarity of user attributes, and the third from substituting $\bar{\imath}_a$ pointwise for U. Clearly, from the convexity of $\mathcal{P}$ it follows that $(\bar{s}, \bar{\imath}) \in \mathcal{P}$.

Second, we denote by $J^D(\vec{s}, \vec{\imath})$ and $J^D(\bar{s}, \bar{\imath})$ the objective value of the solutions $(\vec{s}, \vec{\imath})$ and $(\bar{s}, \bar{\imath})$ respectively. We have that

$$\begin{aligned} J^D(\vec{s}, \vec{\imath}) &= \sum_{n=1}^{N} \mathbb{E}_{U_n}\left[r\left(s_n(U_n); U_n\right) + \sum_{a\in\mathcal{A}} Q_{a,n} i_{n,a}(U_n)\right] = \mathbb{E}_U\left[\sum_{n=1}^{N} r\left(s_n(U); U\right) + \sum_{a\in\mathcal{A}} Q_a \sum_{n=1}^{N} i_{n,a}(U)\right] \\ &\le \mathbb{E}_U\left[N r\left(\frac{1}{N}\sum_{n=1}^{N} s_n(U); U\right) + \sum_{a\in\mathcal{A}} Q_a \sum_{n=1}^{N} i_{n,a}(U)\right] = N\mathbb{E}_U\left[r\left(\bar{s}(U)\right) + \sum_{a\in\mathcal{A}} Q_a \bar{\imath}_a(U)\right] = J^D(\bar{s}, \bar{\imath}), \end{aligned}$$

where the second from the linearity of expectation and the stationarity of user attributes, the inequality follows from the concavity of the revenue function, and the thrid equality from substituting $\bar{s}$ and $\bar{\imath}$ pointwise for all U. □

A.1.4 Proof of Proposition 2.3

First, we formulate the problem as a stochastic control problem (SCP). Second, we show that the optimal objective value of the DAP provides an upper bound on the objective value of the SCP.

Step 1. A stochastic control policy maps states of the system to control actions (prices and target advertiser), and is adapted to the history up to the decision epoch. We restrict our attention to policies that always submit the impression to AdX, which were argued to be optimal. Recall that given the reserve price, the publisher knows the actual probability that the impression is accepted by AdX. As before, we recast the problem in terms of the survival probability control. Hence, the publisher picks the probability that the impression is accepted. Conversely, given a survival probability the reserve price can be easily computed using $\bar{F}^{-1}(\cdot)$. We denote by $s_n^\mu(U) \in [0,1]$ the target survival probability under policy μ at time n when an impression with attributes U arrives. Similarly, we let $I_{n,a}^\mu(U) \in \{0,1\}$ indicate whether the n^{th} impressions is assigned to advertiser a or not when policy μ is used. In particular, $I_{n,a}^\mu(U) = 1$ indicates that the impression should be assigned to the advertiser if rejected by AdX.

We let the binary random variable $X_n(s_n^\mu)$ indicate whether the n^{th} impression is accepted by AdX or not when policy μ is used. Specifically, $X_n(s_n^\mu) = 1$ indicates that the impression is accepted by AdX, and when $X_n(s_n^\mu) = 0$ the impression is rejected by AdX. Notice that, conditioning on impression's attributes and the history, $X_n(s_n^\mu)$ is a Bernoulli random variable with success probability s_n^μ.

We denote by $\mathcal{M}$ the set of admissible policies, i.e. policies that are non-anticipating, adapting and feasible. A feasible policy should satisfy the contractual obligations with each advertiser, or equivalently $\sum_{n=1}^{N} [1 - X_n(s_n^\mu)]\, I_{n,a}^\mu = C_a$ in an almost sure sense. Additionally, the target advertiser controls should satisfy that $\sum_{a\in\mathcal{A}} I_{n,a}^\mu \leq 1$, since the impression should be assigned to at most one advertiser. Finally, the equivalent stochastic optimal control problem is

$$J_N^* = \max_{\mu\in\mathcal{M}} \mathbb{E}\left[\sum_{n=1}^{N} r(s_n^\mu) + (1 - s_n^\mu) \sum_{a\in\mathcal{A}} I_{n,a}^\mu Q_{n,a}\right], \tag{A.1}$$

where J_N^* denotes the optimal expected revenue over the set of admissible policies $\mathcal{M}$. The objective follows from conditioning on the quality of the impression and the history. By the Principle of Optimality it is the case that the dynamic program described in §2.2.1 provides an optimal solution to the SCP (Bertsekas, 2000) and $J_N(C) = J_N^*$.

Step 2. Let μ^* be the optimal policy for the stochastic control problem. Let $\hat{s} = \{\hat{s}_n(\cdot)\}_{n=1,\ldots,N}$ and $\hat{\imath} = \{\hat{\imath}_n(\cdot)\}_{n=1,\ldots,N}$ be deterministic vectors of controls defined as

$$\hat{s}_n(U) = \mathbb{E}_{\mathcal{F}_n}\left[s_n^{\mu^*}(U) \mid U\right] \qquad \forall U \text{ pointwise},$$
$$\hat{\imath}_{n,a}(U) = \mathbb{E}_{\mathcal{F}_n}\left[(1 - s_n^{\mu^*}(U)) I_{n,a}^{\mu^*}(U) \mid U\right] \qquad \forall U \text{ pointwise}, a \in \mathcal{A},$$

where the expectation is taken over the history of the system until n, which is denoted by $\mathcal{F}_n$, and conditional on a particular realization of U. The resulting controls are independent of the history, and dependent only on the realization of U and the impression number n. Thus, they fulfill the first approximation and are valid deterministic vectors of controls. We will show that $(\hat{s}, \hat{\imath})$ is feasible for the DAP, and that its objective value (in the DAP) dominates the optimal objective value of the SCP. Then, we may conclude that $J_N^* \le J_N^D(\hat{s}, \hat{\imath}) \le J_N^D$, because no feasible solution is better than the optimal.

First, for the contract fulfillment constraint we have that for each advertiser $a \in \mathcal{A}$

$$\begin{aligned} C_a &= \mathbb{E}\left[\sum_{n=1}^{N}(1 - X_n(s_n^{\mu^*}(U_n)))I_{n,a}^{\mu^*}(U_n)\right] \\ &= \sum_{n=1}^{N} \mathbb{E}\left[\mathbb{E}_{\mathcal{F}_n}\left[(1 - s_n^{\mu^*}(U_n))I_{n,a}^{\mu^*}(U_n) \mid U_n\right]\right] = \sum_{n=1}^{N} \mathbb{E}\left[\hat{\imath}_{n,a}(U)\right], \end{aligned}$$

where the first equality follows from taking expectations to the almost sure contract fulfillment constraint of μ^*, the second from the tower rule, and the third from substituting $\hat{s}$ and $\hat{\imath}$ pointwise for all U and the fact that impressions are i.i.d. Non-negativity of the controls follows trivially. Additionally, is it not hard to show that $\sum_{a\in\mathcal{A}} \hat{\imath}_{n,a}(\cdot) + s_n(\cdot) \le 1$ for all n. Thus, $(\hat{s}, \hat{\imath})$ is a feasible deterministic control.

Second, the objective value of the optimal stochastic control is bounded by

$$\begin{aligned} J_N^* &= \mathbb{E}\left[\sum_{n=1}^{N} \mathbb{E}_{\mathcal{F}_n}\left[r\left(s_n^{\mu^*}(U_n); U_n\right) \mid U_n\right] + \sum_{a\in\mathcal{A}} Q_{n,a}\mathbb{E}_{\mathcal{F}_n}\left[(1 - s_n^{\mu^*}(U_n))I_{n,a}^{\mu^*}(U_n) \mid U_n\right]\right] \\ &\le \mathbb{E}\left[\sum_{n=1}^{N} r\left(\hat{s}_n(U_n); U_n\right) + \sum_{a\in\mathcal{A}} Q_{n,a}\hat{\imath}_{n,a}(U_n)\right] = J_N^D(\hat{s}, \hat{\imath}), \end{aligned}$$

where the first equality follows from the tower rule and because U_n is measurable w.r.t. the conditional expectation, and the inequality from applying Jensen's inequality to the concave revenue function. □

A.1.5 Proof of Theorem 2.2

The first bound follows from Proposition 2.3. We now prove the second bound.

Let $S_{n,a}^{\mu} = \sum_{i=1}^{n} (1 - X_i(s_i^{\mu}(U_i)))\, I_{i,a}^{\mu}(U_i)$ be the total number of impressions assigned to advertiser a by time n when following the stochastic policy μ^B. Additionally, we denote by $S_n^{\mu} = \{S_{n,a}^{\mu}\}_{a\in\mathcal{A}}$ the random vector of impressions assigned to advertisers. Then, $x_{n,a} = C_a - S_{n,a}^{\mu}$ is the total number of impressions left to assign to advertiser a to fulfill the contract, and $m = N - n$ is the total number of impressions remaining to arrive.

To simplify the proof, we let $C_0 = N - \sum_{a\in\mathcal{A}} C_a$ be the total number of impressions that are not assigned to any advertiser (accepted by AdX and discarded), and we refer to $S_{n,0}^{\mu} = n - \sum_{a\in\mathcal{A}} S_{n,a}$

as total number of impressions not assigned to any advertiser by time n when following the stochastic policy μ^B. Because C_0 is the total number of impressions we can dispense of, when the point is reached that $S_{n,0} = C_0$, then all remaining impressions need to be assigned to the advertisers.

Let the random time $N^* = \inf\left\{1 \leq n \leq N \ : \ \ x_{n,a} = 0 \text{ for some } a \in \mathcal{A} \text{ or } \sum_{a\in\mathcal{A}} x_{n,a} = m\right\}$ be the first time that any advertiser's contract is fulfilled or the point is reached where all arriving impressions need to be assigned to the advertisers. Clearly, N^* is a stopping time with respect to the stochastic process $\{S_n^\mu\}_{n=1,\dots,N}$.

In the following, let R_n^μ be the revenue from time n under policy μ^B. Similarly, we denote by R_n the revenue from time n when the deterministic control are used in an alternate system with no capacity constraints. Because the deterministic controls are time-homogeneous, and the underlying random variables are i.i.d., then the random variables $\{R_n\}_{n=1,\dots,N}$ are i.i.d. too. Moreover, it is the case that $\mathbb{E}R_n = J_1^D$. Notice that when $n < N^*$, the controls of stochastic policy μ^B behave exactly as the optimal deterministic controls. Thus, $R_n = R_n^\mu$ for $n < N^*$. Using this fact together with the fact that N^* is a stopping time we get that

$$J_N^B = \mathbb{E}\left[\sum_{n=1}^{N} R_n^\mu\right] = \mathbb{E}\left[\sum_{n=1}^{N^*} R_n + \sum_{n=N^*+1}^{N} R_n^\mu\right] \geq \mathbb{E}\left[\sum_{n=1}^{N^*} R_n\right] = \mathbb{E}N^* J_1^D, \tag{A.2}$$

where the inequality follows from the non-negativity of the revenues, and the last equality from Wald's equation. Then, we conclude that $J_N^B / J_N^D \geq \mathbb{E}N^*/N$.

Next, we turn to the problem of lower bounding $\mathbb{E}N^*$. Before proceeding we make some definitions. We define by $S_{n,a}$ the number of impressions assigned to advertiser a by time n when following the deterministic controls in the alternate system with no capacity constraints. As for the revenues, it is the case that $S_{n,a} = S_{n,a}^\mu$ for $n < N^*$. We define $S_{n,0}$ in a similar fashion.

Let $N_a = \inf\{n \geq 1 \ : \ S_{n,a} = C_a\}$ be the time when the contract of advertiser $a \in \mathcal{A}$ is fulfilled, and $N_0 = \inf\{n \geq 1 \ : \ S_{n,0} = C_0\}$ be the point in time where all arriving impressions need to be assigned to the advertisers. Even though these stopping times are defined with respect to the stochastic process that follows the deterministic controls, it is the case that $N^* = \min_{a\in\mathcal{A}_0}\{N_a\}$. In the remainder of the proof we study the mean and variance of each stopping time, and then conclude with a bound for $\mathbb{E}N^*$ based on those central moments.

For the case of $a \in \mathcal{A}$, the summands of $S_{n,a}$ are independent Bernoulli random variables with success probability ρ_a. The success probability follows from (2.3a). Hence, N_a is a negative binomial random variable with C_a successes and success probability ρ_a. The mean and variance are given by $\mathbb{E}N_a = N$, and $\mathrm{Var}[N_a] = N\frac{1-\rho_a}{\rho_a}$, where we used that $\rho_a = C_a/N$. Similarly, for the case of $a = 0$, now the summands of $S_{n,0}$ are Bernoulli random variables with success probability

ρ_0. Hence, N_0 is a negative binomial random variable with C_0 successes and success probability ρ_0.

Finally, using the lower bound on the mean of the minimum of a number of random variables of Aven (1985) we get that

$$\mathbb{E}N^* = \mathbb{E}\min_{a\in\mathcal{A}_0}\{N_a\} \geq \min_{a\in\mathcal{A}_0}\mathbb{E}N_a - \sqrt{\frac{A}{A+1}\sum_{a\in\mathcal{A}_0}\mathrm{Var}[N_a]}$$
$$= N - \sqrt{\frac{A}{A+1}}\sqrt{\sum_{a\in\mathcal{A}_0} N\frac{1-\rho_a}{\rho_a}} = N - \sqrt{N}K(\rho). \tag{A.3}$$

The result follows from combining (A.2) and (A.3). □

A.1.6 Proof of Corollary 2.1

We prove the complement, that is, the probability that $N^* \geq (1-\epsilon)N$ converges exponentially fast to one. Notice that $N^* \geq (1-\epsilon)N$ if and only if by time $(1-\epsilon)N$ the contract of each advertiser is not yet fulfilled ($S_{(1-\epsilon)N,a} < C_a$), and the point where all impressions need to be assigned to advertisers has not been reached ($S_{(1-\epsilon)N,0} < C_0$). Combining De Morgan's law and Boole's inequality we get that

$$\mathbb{P}\{N^* \geq (1-\epsilon)N\} = \mathbb{P}\{S_{(1-\epsilon)N,a} < C_a \ \forall a \in \mathcal{A}_0\} \geq 1 - \sum_{a\in\mathcal{A}_0}\mathbb{P}\{S_{(1-\epsilon)N,a} \geq C_a\}.$$

Recall that $S_{(1-\epsilon)N,0}$ is the sum of $(1-\epsilon)N$ independent Bernoulli random variables with success probability ρ_a. Hence, we conclude by applying Chernoff's bound to the each summand to obtain $\mathbb{P}\{S_{(1-\epsilon)N,a} \geq C_a\} \leq \exp(-2\epsilon^2\rho_a N)$. □

A.1.7 Proof of Proposition 2.4

The joint distribution of $B_{1:K}$ and $B_{2:K}$ has a density function (Laffont and Maskin, 1980)

$$f(b_1,b_2) = \begin{cases} K(K-1)F(b_2)^{K-2}f(b_1)f(b_2) & \text{if } b_1 \geq b_2 \\ 0 & \text{otherwise} \end{cases}.$$

Then, we have that

$$\begin{aligned} r(p) &= \mathbb{E}\left[\mathbf{1}\{B_{2:K} \geq p\}B_{2:K} + p\mathbf{1}\{B_{1:K} \geq p, B_{2:K} < p\}\right] \\ &= \int_p^\infty \int_p^{b_1} b_2 f(b_1,b_2)\ \mathrm{d}b_2\ \mathrm{d}b_1 + p\int_p^\infty \int_0^p f(b_1,b_2)\ \mathrm{d}b_2\ \mathrm{d}b_1 \\ &= K(K-1)\int_p^\infty b_2 F(b_2)^{K-2} f(b_2)(1-F(b_2))\ \mathrm{d}b_2 + KpF(p)^{K-1}(1-F(p)) \end{aligned}$$

Continuity of $r(s)$ follows because the p.d.f. is continuous, and $p(s)$ is continuous (if F not strictly monotone, the inverse may have jumps). Additionally, we may bound the revenue by

$$r(p) \leq \mathbb{E}\left[\mathbf{1}\{B_{1:K} \geq p\}B_{1:K}\right] \leq K\mathbb{E}\left[\mathbf{1}\{B \geq p\}B\right] \leq K\mathbb{E}B < \infty,$$

the first inequality follows because $B_{1:K}$ is the maximum, the second because any order statistic is upper bounded by the sum of the bids, and the fourth because bids are integrable. Moreover, integrability of B implies that $\lim_{p\to\infty} r(p) = 0$.

Next, we turn to the concavity of $r(s)$. Differentiating w.r.t to p we get

$$\frac{dr}{dp} = KF(p)^{K-1}(\bar{F}(p) - pf(p)).$$

Then, using the fact that $\frac{ds}{dp} = -KF(p)^{1-k}/f(p)$ we get from the composition rule that

$$\frac{dr}{ds} = \left.\frac{dr}{dp}\right|_{p(s)} \frac{dp}{ds} = p(s) - \frac{1}{h(p(s))},$$

where $h(p) = f(p)/\bar{F}(p)$ is the hazard rate of the bidder's valuation. Because $p(s)$ is non-increasing in s and the $h(p)$ is non-decreasing in p, we conclude that $\frac{dr}{ds}$ is non-increasing. Thus, the revenue function is concave.

Finally, notice that the that derivative of the objective w.r.t to s is

$$p(s) - \frac{1}{h(p(s))} - c, \tag{A.4}$$

which is non-increasing. When $c > p_\infty$ we have that (A.4) is negative, so $s^*(c) = 0$ and $p^*(c) = p_\infty$. Similarly, when $c < p_0 - 1/h(p_0)$ we that (A.4) is positive, so $s^*(c) = 1$ and $p^*(c) = p_0$. □

A.1.8 Directional Derivatives of the Dual Objective

Given a subset of the quality space $D \subseteq \Omega$, we define the measure $\mathbb{P}_R(D)$ as the probability that the quality vector belongs to that subset and the impression is rejected by the Ad Exchange when the optimal survival probability is used. More formally,

$$\mathbb{P}_R(D) = \mathbb{E}\left[\left(1 - s^*\left(\max_{a\in\mathcal{A}_0}\{Q_a - v_a\}; U\right)\right)\mathbf{1}\{Q \in D\}\right].$$

Notice that the latter is not a probability measure since $\mathbb{P}_R(\Omega) \leq 1$. Proposition A.1 characterizes the directional derivative of the objective function of the dual along some directions that, as we will show later, are of particular interest. Results are given in terms of the measure $\mathbb{P}_R$.

Proposition A.1. *Given a subset $\alpha \in \mathcal{A}$, the directional derivative of the objective function of the dual w.r.t. directions $\mathbf{1}_\alpha$ and $-\mathbf{1}_\alpha$ are respectively*

$$\nabla_{\mathbf{1}_\alpha}\psi(v) = -\mathbb{P}_R\left\{\max_{a\in\alpha}\{Q_a - v_a\} > \max_{a\in\mathcal{A}_0\setminus\alpha}\{Q_a - v_a\}\right\} + \sum_{a\in\alpha}\rho_a,$$
$$\nabla_{-\mathbf{1}_\alpha}\psi(v) = \mathbb{P}_R\left\{\max_{a\in\alpha}\{Q_a - v_a\} \geq \max_{a\in\mathcal{A}_0\setminus\alpha}\{Q_a - v_a\}\right\} - \sum_{a\in\alpha}\rho_a.$$

Proof. We consider first the direction $\mathbf{1}_\alpha$. Notice that the random function $R\left(\max_{a\in\mathcal{A}_0}\{Q_a - v_a\}; U\right)$ is convex, and thus directionally differentiable. We first show that $\psi(v)$ is finite. From Assumption 2.1 we have that the revenue function is bounded by $r(s;u) \leq M$, and thus $R(c;u) \leq M + \max(c, 0) \leq M + |c|$. Therefore, using the triangle inequality we obtain that

$$\psi(v) \leq M + \mathbb{E}|\max_{a\in\mathcal{A}_0}\{Q_a - v_a\}| + \sum_{a\in\mathcal{A}}|v_a| \leq M + \sum_{a\in\mathcal{A}}\mathbb{E}|Q_a| + 2|v_a| < \infty$$

We can now apply Theorem 7.46 in Shapiro et al. (2009) and obtain that $\psi(v)$ is directionally differentiable at v and that one can exchange expectation and directional derivative. Putting all together we get that

$$\begin{aligned}\nabla_{\mathbf{1}_\alpha}\psi(v) &= \mathbb{E}\left[\nabla_{\mathbf{1}_\alpha}R\left(\max_{a\in\mathcal{A}_0}\{Q_a - v_a\}; U\right)\right] + \sum_{a\in\alpha}\rho_a \\ &= \mathbb{E}\left[\frac{\mathrm{d}R}{\mathrm{d}c}\left(\max_{a\in\mathcal{A}_0}\{Q_a - v_a\}; U\right)\nabla_{\mathbf{1}_\alpha}\left\{\max_{a\in\mathcal{A}_0}\{Q_a - v_a\}\right\}\right] + \sum_{a\in\alpha}\rho_a,\end{aligned}$$

where the second equation follows from the chain rule. We conclude by the fact that $\frac{\mathrm{d}R}{\mathrm{d}c}(c;u) = 1 - s^*(c;u)$ and $\nabla_{\mathbf{1}_\alpha}\left\{\max_{a\in\mathcal{A}_0}\{Q_a - v_a\}\right\} = -\mathbf{1}\left\{\max_{a\in\alpha}\{Q_a - v_a\} > \max_{a\in\mathcal{A}_0\setminus\alpha}\{Q_a - v_a\}\right\}$. A similar result follows for the opposite direction $-\mathbf{1}_\alpha$ from the fact that $\nabla_{-\mathbf{1}_\alpha}\left\{\max_{a\in\mathcal{A}_0}\{Q_a - v_a\}\right\} = \mathbf{1}\left\{\max_{a\in\alpha}\{Q_a - v_a\} \geq \max_{a\in\mathcal{A}_0\setminus\alpha}\{Q_a - v_a\}\right\}$.

□

A.1.9 Proof of Proposition 2.5

The proof proceeds by contradiction, that is, we assume that there is no feasible flow. First, we cast the feasible flow problem as a maximum flow problem. Feasibility would imply the existence of a flow with value $1 - \mathbb{P}(\emptyset\text{-tie})$. But since we assume that no such feasible flow exists, by the max-flow min-cut theorem there should exists a cut with value strictly less than $1 - \mathbb{P}(\emptyset\text{-tie})$. The contradiction arises because the optimality conditions of v for the dual problem (2.4) imply that the every cut is lower bounded by $1 - \mathbb{P}(\emptyset\text{-tie})$.

In order to write the feasible flow problem as a maximum flow problem, we first add a source s and a sink t. Second, we add one arc from s to each node associated to a non-empty subset

$S \subseteq \mathcal{A}_0$ (left-hand side nodes) with capacity $\mathbb{P}(S\text{-tie})$. Third, we add one arc from each advertiser $a \in \mathcal{A}_0$ (right-hand side nodes) to t with capacity ρ_a. Lastly, we set the capacity of arcs from S to $a \in S$ to infinity.

Now, since no feasible flow exists, by the max-flow min-cut theorem there should be a cut with value strictly less than $1 - \mathbb{P}(\emptyset\text{-tie})$. Let $\alpha \subseteq \mathcal{A}_0$ be the advertiser nodes (right-hand) belonging to the t side of a *minimum* cut. Figure A.1 shows the minimum cut. Next we argue that subset nodes in the s side verify that $S \cap \alpha = \emptyset$, while those in the t side verify that $S \cap \alpha \neq \emptyset$. First, because the cut has minimum value, there is no arc from a subset node to an advertiser node crossing the cut (those arcs have infinity capacity). Equivalently, within the s side of the cut, all subsets nodes $S \subseteq \mathcal{A}_0$ should verify that $S \cap \alpha = \emptyset$. Second, observe that any subset node with $S \cap \alpha = \emptyset$ in the t side of the cut could be moved to the s side of the cut without increasing the value of the cut. Hence, with no loss of generality we can assume that all subset nodes in the t side of the cut verify that $S \cap \alpha \neq \emptyset$.

As a consequence, the only arcs crossing the cut are those from the source to the subsets $S \cap \alpha \neq \emptyset$, and those from advertisers $\mathcal{A}_0 \setminus \alpha$ to the sink. The value of this cut is

$$\sum_{S \subseteq \mathcal{A}_0 : S \cap \alpha \neq \emptyset} \mathbb{P}(S\text{-tie}) + \sum_{a \in \mathcal{A}_0 \setminus \alpha} \rho_a.$$

Because the value is strictly less than $1 - \mathbb{P}(\emptyset\text{-tie})$ we get that

$$\sum_{S \subseteq \mathcal{A}_0 : S \cap \alpha \neq \emptyset} \mathbb{P}(S\text{-tie}) < \sum_{a \in \alpha} \rho_a, \tag{A.5}$$

where we used that $\sum_{a \in \mathcal{A}} \rho_a + \rho_0^{\text{eff}} = 1 - \mathbb{P}(\emptyset\text{-tie})$.

Next, we look at the optimality conditions of v for the dual problem (2.4). We distinguish between the case that $0 \notin \alpha$ and $0 \in \alpha$. First suppose that $0 \notin \alpha$, and consider the direction $-\mathbf{1}_\alpha$ that has a -1 if $a \in \alpha$ and 0 elsewhere. According to proposition A.1 the directional derivative of the objective at v is

$$\begin{aligned} \nabla_{-\mathbf{1}_\alpha} \psi(v) &= \mathbb{P}_R \left\{ \max_{a \in \alpha} \{Q_a - v_a\} \geq \max_{a \in \mathcal{A}_0 \setminus \alpha} \{Q_a - v_a\} \right\} - \sum_{a \in \alpha} \rho_a \\ &= \sum_{S \subseteq \mathcal{A}_0 : S \cap \alpha \neq \emptyset} \mathbb{P}(S\text{-tie}) - \sum_{a \in \alpha} \rho_a, \end{aligned}$$

where we have written the event that the maximum is verified non-exclusively by some advertiser $a \in \alpha$ as all S-ties in which some advertiser $a \in \alpha$ is involved. The optimality of v implies that the directional derivative along that direction is greater or equal to zero, contradicting equation (A.5).

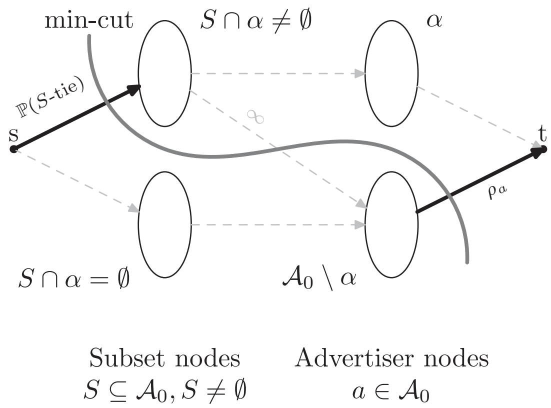

Figure A.1: The flow problem in the bipartite graph with a minimum cut. A source s connected to the subset nodes and a sink t connected to the advertisers nodes was included. $\alpha \subseteq \mathcal{A}_0$ is the subset of advertiser nodes (right-hand) belonging to the t side. Note that no there is no arc from a subset node to an advertiser node crossing the cut.

When $0 \in \alpha$ we consider the direction $\mathbf{1}_{\mathcal{A}\setminus\alpha}$ that has a 1 if $a \notin \alpha$ and 0 elsewhere. The direction derivative is now

$$\begin{aligned}\nabla_{\mathbf{1}_{\mathcal{A}\setminus\alpha}}\psi(v) &= -\mathbb{P}_R\left\{\max_{a\in\mathcal{A}\setminus\alpha}\{Q_a - v_a\} > \max_{a\in\alpha\cup\{0\}}\{Q_a - v_a\}\right\} + \sum_{a\in\mathcal{A}\setminus\alpha}\rho_a,\\ &= -\sum_{S\subseteq\mathcal{A}_0 : S\subseteq\mathcal{A}\setminus\alpha}\mathbb{P}(S\text{-tie}) + \sum_{a\in\mathcal{A}\setminus\alpha}\rho_a = \sum_{S\subseteq\mathcal{A}_0 : S\cap\alpha\neq\emptyset}\mathbb{P}(S\text{-tie}) - \sum_{a\in\alpha}\rho_a,\end{aligned}$$

where in the second equation we have written the event that the maximum is verified exclusively by some advertiser $a \in \alpha$ as all S-ties in which only advertisers in α are involved. Again, the optimality of v implies that the directional derivative along that direction is greater or equal to zero, contradicting equation (A.5). $\square$

A.2 Comparison with the Primal-Dual Method

Consider the allocation problem faced by a publisher in display advertising in which arriving impressions need to be assigned to advertisers, and there is no option of sending to an exchange. This problem is a particular case of our model where the winning bid random variable is identically zero, i.e. $B = 0$. The following proposition shows that the optimal controls admit simple analytical expressions.

Proposition A.2. *Suppose that ties have zero probability. Then, in the case without AdX the optimal controls are* $I_a(Q) = \mathbf{1}\{Q_a - v_a \geq Q_{a'} - v_{a'}\ \forall a' \in \mathcal{A}_0\}$ *where* $v = \{v_a\}_{a\in\mathcal{A}_0}$ *satisfies* $v_0 =$

0 *and*

$$\mathbb{P}\left\{Q_a - v_a \geq Q_{a'} - v_{a'}\ \forall a' \in \mathcal{A}_0\right\} = \rho_a \qquad \forall a \in \mathcal{A}.$$

Proof. Proof.Because $B = 0$, then it is not hard to show that $\bar{F}^{-1}(s) = 0$, and that the revenue function is $r(s) = 0$. Hence, the revenue function is regular and satisfies Assumption 2.1. Moreover, the optimal survival probability is $s^*(c) = 0$, and $R(c) = c$. The result follows from substituting these functions in Theorem 2.1. □ □

The resulting decision rule $\arg\max_a\{Q_a - v_a\}$ is identical to the rule studied in previous work (e.g., Devenur and Hayes (2009)), where v_a is an optimal dual variable resulting from solving an assignment problem on a sample of the data, where the distribution is unknown. Roughly speaking, in Devenur and Hayes (2009) (and similarly in other work Feldman et al. (2010); Vee et al. (2010); Agrawal et al. (2009)), it is shown that as long as the sample is of size $\approx\ \epsilon n$, the overall assignment will be $\approx \epsilon$ close to the optimal offline solution.

In our model, the parameters of the quality distribution are known, so we do not need to use a sample. Of course in practice, the parameters need to be learned, and so we would need to use a sample of the data in order to learn them; but in many settings (including online advertising) it is reasonable to assume that we at least know the *form* of the distribution (e.g., normal, exponential, Zipf), albeit not the specific parameters (mean, variance, covariance, etc.). The techniques in Devenur and Hayes (2009) are powerful *because* they don't need to assume anything about the distribution, but it is important to ask what can be gained from knowing the *form* of the distribution, which is what we do in the remainder of this section, both analytically and experimentally.

More formally, suppose the distribution of quality is not known with certainty, but we have at our disposal a sample of M quality vectors $\{q_m\}_{m=1}^M$ that may be used to pin-down the distribution. Additionally, it is known that the qualities are drawn independently from a population with continuous density $g(x|\theta)$, where θ is an unknown parameter to be estimated. Let $G(x|\theta)$ be the c.d.f. which we assume to be strictly monotonic. For simplicity, we deal with the case of one advertiser with capacity to impression ratio of ρ. From Proposition A.2 the optimal DAP control is the $(1-\rho)$-quantile of Q, that is, $v = \bar{G}^{-1}(\rho|\theta)$. We compare the asymptotic efficiency of a parametric and a non-parametric estimation of the model.

Parametric estimation method. Let $\hat{\theta}_M^{\text{mle}}$ be the maximum likelihood estimator (MLE) of the unknown parameter θ. That is, $\hat{\theta}_M^{\text{mle}}$ is an optimal solution of the program $\max_\theta \sum_{m=1}^M \log g(q_m|\theta)$. Once we have our estimator, we plug-in the estimated distribution in the dual problem, and solve for the optimal dual variable $\hat{v}_M^{\text{mle}}$. Again, from Proposition A.2 we have that the optimal dual variable, given our maximum likelihood estimation, is given by $\hat{v}_M^{\text{mle}} = \bar{G}^{-1}(\rho|\hat{\theta}_M^{\text{mle}})$.

In turn, by the invariance property of the MLE, it is the case that $\hat{v}_M^{\text{mle}}$ is the maximum likelihood estimator of the true optimally dual variable v (see, e.g., Casella and Berger (2002)). As a consequence, under some regularity conditions, we have that our new estimator is consistent, asymptotically efficient, and asymptotically normal

$$\sqrt{M}(\hat{v}_M^{\text{mle}} - v) \Rightarrow \mathcal{N}(0, u(\theta)),$$

where the $u(\theta)$ is the Cramér-Rao lower bound on the variance of any unbiased estimator. The Cramér-Rao lower bound is $u(\theta) = \left(\frac{\partial \bar{G}^{-1}}{\partial \theta}(\rho|\theta)\right)^2 I(\theta)^{-1}$, where $I(\theta) = \mathbb{E}\left[\left(\frac{\partial}{\partial \theta}\ln g(Q|\theta)\right)^2\right]$ is the Fisher information of parameter θ.

Non-parametric estimation method. Considering the Sample Average Approximation (SAA) of the dual problem (2.4) we can obtain a non-parametric estimator of the truly optimal dual variable (see Chapter 5 from Shapiro et al. (2009) for a review of the topic). In a SAA the expected value of the stochastic program is approximated by the sample average function over the observations $\{q_m\}_{m=1}^M$. In our case, we have that

$$\hat{v}_M = \arg\min_v \frac{1}{M}\sum_{m=1}^{M} \max\{q_m - v, 0\} + \rho v.$$

Equivalently, the previous problem can be stated as a linear program, and in this case one obtains the training-based Primal-Dual method as described in Devenur and Hayes (2009). It can be shown that, under some conditions, the non-parametric estimator $\hat{v}_M$ is consistent and asymptotically normal (Shapiro et al., 2009). As we shall see, this estimator is not necessarily efficient. It is not hard to prove that the sample $(1-\rho)$-quantile is an optimal solution to the SAA problem. Hence, from the asymptotic distribution of the $(1-\rho)$-quantile we have that

$$\sqrt{M}(\hat{v}_M - v) \Rightarrow \mathcal{N}(0, u'(\theta)),$$

where the variance is $u'(\theta) = \frac{\rho(1-\rho)}{g(v|\theta)^2}$.

Analysis. Both the parametric and the non-parametric estimators converge, as the number of samples increases, to the true optimal solution. However, the non-parametric estimator is not as efficient as the parametric counterpart. Indeed, this is expected since the maximum likelihood estimator is known to be asymptotically efficient. We measure the *relative efficiency* as the ratio of asymptotic variance of the non-parametric estimator to parametric one, i.e. $\varepsilon(\theta) = \frac{u'(\theta)}{u(\theta)}$.

Until now, our analysis has been in terms of the optimal dual solution. The rationale is that the closer the dual variable is to the true value v, the better the performance of the policy should be. Next, we quantify analytically how does a deviation from the optimal solution impacts the performance of the policy. To assess the performance of the policy we look at the fluid limit as described in §A.5. The next proposition shows that the relative efficiency in terms of the performance is exactly equal to $\varepsilon(\theta)$. Hence, there is no loss in looking at the relative efficiency of the estimators instead.

Proposition A.3. *The relative efficiency of the non-parametric estimator is*

$$\varepsilon(\theta) = \rho(1-\rho)I(\theta)\left(\frac{\partial \bar{G}}{\partial \theta}(v|\theta)\right)^{-2} \geq 1, \tag{A.6}$$

which is exactly equal to the relative efficiency in terms of the policies' performance.

Proof. Proof of Proposition A.3. For (A.6), we use that $\frac{\partial \bar{G}^{-1}}{\partial \theta}(\rho|\theta) = \frac{\partial \bar{G}}{\partial \theta}(v|\theta)/g(v|\theta)$, which follows from the implicit function theorem. In view of Cramér-Rao lower bound, we have that $\varepsilon(\theta) \geq 1$.

Next, we look at the average yield of the policy as the number of impressions grows to infinity when a bid price of u is employed, denoted by $\bar{J}(u)$. The limiting performance is given by

$$\bar{J}(u) = \begin{cases} \rho \mathbb{E}_\theta[Q|Q \geq u], & \text{if } u < v, \\ \mathbb{E}_\theta[Q] - (1-\rho)\mathbb{E}_\theta[Q|Q \leq u], & \text{if } u \geq v. \end{cases}$$

Under our assumptions, the performance function is continuous u. One would be tempted to apply the Delta Method to derive the asymptotic distribution of the performance. Unfortunately, $\bar{J}(\cdot)$ is not differentiable at v. However, it is the case that the performance function is semi-differentiable at v with finite right-derivative $\bar{J}'_+(v) \geq 0$ and left-derivative $\bar{J}'_-(v) \leq 0$. Thus, we can apply an extension of the Delta Method for directionally differentiable functions proved by Shapiro (1991), and obtain

$$\sqrt{m}(\bar{J}(\hat{v}_M) - \bar{J}(v)) \Rightarrow d\bar{J}\Big(v; \mathcal{N}(0, u'(\theta))\Big),$$
$$\sqrt{m}(\bar{J}(\hat{v}_M^{\text{mle}}) - \bar{J}(v)) \Rightarrow d\bar{J}\Big(v; \mathcal{N}(0, u(\theta))\Big),$$

where $d\bar{J}(v;\xi)$ is the Gâteaux derivative of $\bar{J}$ at the point v along the direction ξ, which is given by $d\bar{J}(v;\xi) = \bar{J}'_+(v)\xi$ when $\xi \geq 0$, and $d\bar{J}(v;\xi) = \bar{J}'_-(v)\xi$ when $\xi < 0$. Note that the asymptotic variance of the performances are $u'(\theta) \cdot K$, and $u(\theta) \cdot K$ respectively, where the performance scale factor is given by $K = \frac{1}{2}(\bar{J}'_+(v)^2 + \bar{J}'_-(v)^2) - \frac{1}{2\pi}(\bar{J}'_+(v) - \bar{J}'_-(v))^2$. Thus, the relative efficiency of the performance is identical to $\varepsilon(\theta)$. $\square$ $\square$

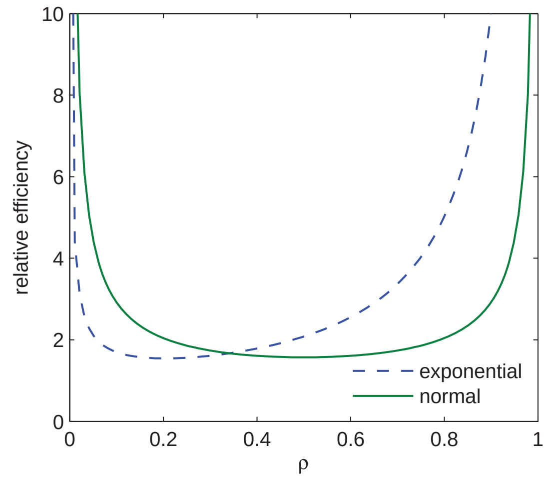

Figure A.2: Relative efficiency as a function of the capacity to impression ratio ρ for the exponential distribution, and the normal distribution with known variance.

Examples. To fix ideas we consider two simple examples. First, suppose that $Q \sim \exp(\theta)$. The maximum likelihood estimator is given by $\hat{\theta}_M^{\text{mle}} = \left(\frac{1}{M}\sum_{m=1}^{M} q_m\right)^{-1}$, and the Fisher information is $I(\theta) = \theta^{-2}$. The optimal dual variable is $v = -\theta^{-1}\ln\rho$. Hence, the relative efficiency is $\varepsilon(\theta) = (1-\rho)/(\rho\ln^2\rho)$. In this case, the relative efficiency is lower bounded by $\varepsilon(\theta) \geq 1.544$. The lower bound is tight, and attained at $\rho \approx 0.2032$. The relative efficiency as a function of the capacity to impression ratio is plotted in Figure A.2. As shown in the figure, the relative efficiency may be arbitrarily bad as the capacity to impression ratio gets close to zero or one.

For the next example we assume that qualities are normal with known variance σ^2 and unknown mean, that is, $Q \sim \mathcal{N}(\theta, \sigma^2)$. The maximum likelihood estimator is the sample mean, $\hat{\theta}_M^{\text{mle}} = \frac{1}{M}\sum_{m=M}^{m} q_m$, and the Fisher Information is $I(\theta) = \sigma^{-2}$. In this case the relative efficiency is given by $\varepsilon(\theta) = 2\pi\rho(1-\rho)\exp\left(\Phi^{-1}(1-\rho)^2\right)$, with Φ^{-1} being the inverse of the standard normal c.d.f. Here the relative efficiency is lower bounded by $\varepsilon(\theta) \geq \pi/2$, with the minimum attained at $\rho = 1/2$. Interestingly, the relative efficiency is invariant under monotonic transformations of any random variable. Hence, the previous result holds too for the log-normal distribution.

A.2.1 Numerical Experiments

In this experiment we study the performance the our algorithm, and contrast it with a Primal-Dual (PD) method, and show experimentally that the advantage of parametric estimation extends to multiple advertisers as well. Since no existing PD method is known yet for the AdX problem, we consider instead the case with no AdX. The Primal-Dual approach (Devenur and Hayes, 2009), uses a sample from data to estimate the dual variables and uses it in a bid-price control policy. In contrast, our algorithm, as stated, assumes the parameters of the quality distribution are known, and uses that to estimate the dual variables. So we do not need to use a sample. Of course in practice, the parameters need to be learned, and so we would need to use a sample of the data in order to learn them; but in many settings (including online advertising) it is reasonable to assume that we at least know the *form* of the distribution (e.g., normal, exponential, Zipf), albeit not the specific parameters (mean, variance, covariance, etc.). The techniques in Devenur and Hayes (2009) are powerful *because* they don't need to assume anything about the distribution, but it is important to ask what can be gained from knowing the *form* of the distribution, which is what we do in the remainder of this section.

In order to objectively assess the performance of our algorithm we adopt the user type model described in §2.3 as a generative model. The generative model is used to generate sample data on which both our algorithm and a PD method are tested. The advantages of adopting a generative model are twofold. First, it allows us to compute the truly optimal policy μ^{OPT}. Second, the true performance of any policy can be evaluated efficiently using a fluid limit (see §A.5).

The computational experiment is conducted as follows. First, a *training* data set of M impressions is generated. We denote the sampled quality vectors by $\{q_m\}_{m=1}^{M}$. Then, we estimate the parameters of the model on the training set as follows. For each type we estimate the type probabilities $\hat{\pi_T}$; and mean $\hat{\mu_T}$, and covariance matrix $\hat{\Sigma_T}$ of the logarithm of the qualities. Next, the dual problem (2.4) is solved on the estimated parametric model using a Gradient Descent Method as described in §A.4. Note that, since no AdX is considered, the maximum expected revenue function $R(\cdot)$ is the identity. Using the optimal solution v^{EST} we construct a policy, which be refer as μ^{EST}.

Simultaneously, we employ the PD method on the training data. The PD method amounts to solving a sample average approximation of problem (2.4), which results in the following linear program

$$
\begin{aligned}
\min_{v,\lambda} \; & \frac{1}{M} \sum_{m=1}^{M} \lambda_m + \sum_{a \in \mathcal{A}} \rho_a v_a \qquad\qquad \text{(A.7)} \\
\text{s.t. } \; & \lambda_m + v_a \geq q_{m,a}, \qquad \forall m, a \\
& \lambda_m \geq 0 \qquad \forall m.
\end{aligned}
$$

The linear program is solved using CPLEX 12. Again, using the dual optimal solution v^{PD} we construct a policy μ^{PD}.

Afterwards, we assess the performance of both policies using a fluid limit. These steps are replicated on 50 different training sets. Table A.1 reports the average results over the training

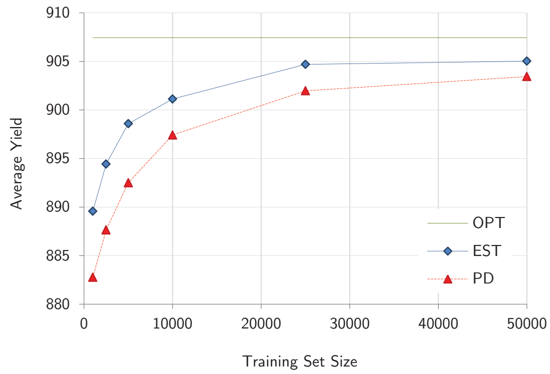

(a) Average Yield

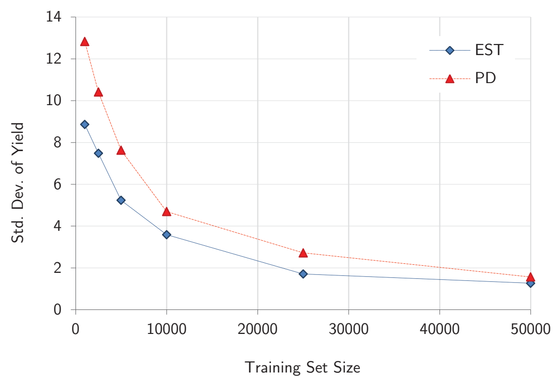

(b) Std. Dev. of Yield

Figure A.3: Average (a) and standard deviation (b) of yield as a function of training set sample size; results are shown for the parametric method (EST) based on our policy μ and the primal-dual method (PD) as in Devenur and Hayes (2009). Both policies converge to the optimal yield, but EST converges faster, and with less variance.

sets for different sizes of training sets, and instances. Plots of the results for a given instance are shown in Figure A.3.

Discussion. Results show that for both algorithms, as the size of the training set increases, the optimality gap decreases at a rate of $O(M^{\frac{1}{2}})$. However, the parametric method performs uniformly better that the non-parametric PD method. Additionally, the variability across different training sets diminishes as the size of the training set increases. Indeed, we observe that the standard deviation over training sets converges to zero for both methods, but the convergence is faster for the parametric one. In some sense this is expected, since the true data model follows exactly the distributional assumptions. However, the PD method is expected to be more robust to model misspecification.

Instance 1 (A = 3, T = 4, OPT = 2075.09)

Training	EST		PD	
Set Size	mean	std.dev.	mean.	std.dev
100	2004.16 (3.42%)	33.978	1990.32 (4.08%)	37.552
1000	2053.41 (1.04%)	10.008	2047.92 (1.31%)	12.365
2500	2065.12 (0.48%)	4.956	2062.76 (0.59%)	5.838
5000	2068.44 (0.32%)	3.681	2066.99 (0.39%)	4.224

Instance 2 (A = 6, T = 10, OPT = 907.44)

Training	EST		PD	
Set Size	mean	std.dev.	mean.	std.dev
1000	889.58 (1.97%)	8.861	882.77 (2.72%)	12.829
2500	894.43 (1.43%)	7.485	887.64 (2.18%)	10.418
5000	898.59 (0.98%)	5.231	892.51 (1.65%)	7.625
10000	901.13 (0.70%)	3.588	897.42 (1.10%)	4.692
25000	904.69 (0.30%)	1.712	901.97 (0.60%)	2.720
50000	905.03 (0.27%)	1.267	903.44 (0.44%)	1.567

Instance 3 (A = 17, T = 15, OPT = 894.82)

Training	EST		PD	
Set Size	mean	std.dev.	mean.	std.dev
2500	859.83 (3.91%)	9.937	849.44 (5.07%)	14.615
5000	868.61 (2.93%)	5.870	861.06 (3.77%)	7.954
10000	877.59 (1.92%)	5.226	873.46 (2.39%)	6.577
25000	884.04 (1.20%)	2.585	881.13 (1.53%)	3.747
50000	887.34 (0.84%)	1.926	885.11 (1.08%)	2.728

Instance 4 (A = 14, T = 10, OPT = 928.76)

Training	EST		PD	
Set Size	mean	std.dev.	mean.	std.dev
2500	892.55 (3.90%)	12.886	888.88 (4.29%)	13.427
5000	903.04 (2.77%)	8.537	901.79 (2.90%)	10.277
10000	911.25 (1.88%)	6.951	909.96 (2.02%)	6.935
25000	917.30 (1.23%)	3.353	915.81 (1.39%)	3.705
50000	921.36 (0.80%)	2.668	920.11 (0.93%)	2.716

Table A.1: Experimental results comparing the performance of our parametric method (EST) with the non-parametric primal-dual method (PD). No AdX present in this experiment.

Another experiment, though results are not reported, was conducted to test the strength of the parametric method on real data. We observed that, when the training set is small (around thousands), the parametric method performs better than the non-parametric one. However, as the sample size increases the non-parametric method outperforms the other. The rationale for this behavior is that, when data is scarce, the parametric method can exploit the distributional assumptions to reconstruct a fair representation of the data. However when the training set is larger, the fit of our model to real data is not perfect, and the non-parametric method can withstand deviations more robustly.

A.3 Incorrect Assignments in the User Type Model

In §2.3 we introduced a user-type model with good-will penalties to accommodate the fact that advertisers have specific targeting criteria. If the contracts are feasible, that is, there is enough inventory to satisfy the targeting criteria; one would expect our policy to assign only impressions within the criteria. In this section we formalize the concept of a feasible operation, and give sufficient conditions under which the stochastic control policy does not assign any impressions outside of the targeting criteria.

It is straightforward to state the problem of determining whether contracts can be satisfied or not, as a feasible flow problem on a bipartite graph. The problem can be formulated on an graph with one node for each user type T with a supply of $\pi(T)$, on the left side; and one node for each advertisers $a \in \mathcal{A}_0$ with a demand ρ_a, on the right side. Then, we say that the operation is *feasible* if the user type-advertiser graph admits a feasible flow.

The feasibility of the operation, albeit necessary, does not suffice to guarantee that no impressions outside the targeting criteria are assigned to the advertisers. When advertisers compete for the same type, and one of them obtains a potentially unbounded reward for that type; it may be optimal to allow the latter advertiser to cannibalize the user type, and force the others advertisers to take types outside of their criteria. This may occur, surprisingly, for all conceivable penalties. However, if qualities are bounded, and penalties are set high enough, then the optimal policy would not recommend the assignment of impressions outside the targeting criteria. Even in this case some impressions may be incorrectly assigned in the left-over regime, but the probability of this event decays exponentially fast. We formalize this discussion in the following proposition.

Proposition A.4. *Suppose that the user type-advertiser graph admits a feasible flow, and that qualities and bids from AdX are bounded by* $\frac{1}{A} \min_a \tau_a$*. Then, the stochastic control policy does not assign any impressions outside of the targeting criteria, except perhaps for the left-over regime.*

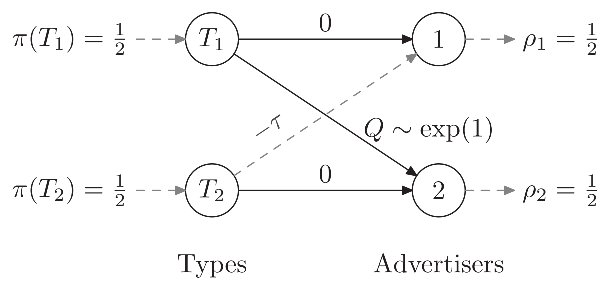

Figure A.4: Example with two user types, and two advertisers.

Proof. Proof Sketch of Proposition A.4. For simplicity we consider the case without AdX. Let i be an optimal solution to the DAP, and suppose that some advertisers are assigned some types outside of their targeting criteria. We will construct another solution with greater or equal yield in which no incorrect assignments are made.

An optimal solution to the DAP is a vector of functions $i_{T,a} : \Omega_T \to [0,1]$ for $a \in \mathcal{A}_0$ and $T \in \mathcal{T}$ such that

$$\sum_{T\in\mathcal{T}} \mathbb{E} i_{T,a}(Q) = \rho_a, \quad \forall a \in \mathcal{A}_0,$$
$$\sum_{a\in\mathcal{A}_0} i_{T,a}(Q) = \pi(T), \quad \text{(a.s.) } \forall T \in \mathcal{T},$$

We refer to each of the functions in a solution as components.

Next, we construct a feasible solution i^0 to the DAP from a feasible flow of the user type-advertiser graph. Take the difference $\Delta i = i^0 - i$, which is a circulation in the user type-advertiser graph. The circulation Δ may have components of mixed signs. Because i^0 has no incorrect assignments, if advertiser a is assigned a type $T \not\ni a$ not in her criteria, then the circulation verifies that $\Delta i_{T,a}(Q) = -i_{T,a}(Q)$. Hence, the components with incorrect assignments are negative.

Let a be an advertiser that is assigned a type $T \not\ni a$ not in her criteria, that is, $\mathbb{E}[i_{T,a}(Q)] > 0$. We may find an augmenting cycle w containing the incorrect assignment, such that if we push some flow along this cycle, we construct another solution $i + w$ with fewer incorrect assignments. The cycle w has at most $A + 1$ positive components, and at most $A + 1$ negative components. The cost associated to the negative components is at most $A\frac{1}{A}\min_{a'} \tau_{a'} - \tau_a \leq 0$. All positive components arcs have a cost of at least zero, and the total cost of this cycle is positive. Thus, the new solution $i + w$ has greater or equal yield. Moreover, $\mathbb{E}[(i + w)_{T,a}(Q)] < \mathbb{E}[i_{T,a}(Q)]$, and no new incorrect assignments are introduced. Repeating this procedure, we may construct a solution with no incorrect assignments. □ □

Note that since the left-over regime is vanishingly small in proportion to the length of the horizon (Cor. 2.1) this implies that the number of unassigned impressions is small. Thus in practice, a publisher may set $C'_a = C_a + \epsilon$, discard any impressions assigned by the policy outside the targeting criteria, and ensure that contracts are filled properly.

Next, we prove by example that the requirement that qualities are bounded is necessary for the previous result to hold. Consider a publisher who contracts with two advertisers, and agrees to deliver one half of the arriving impressions to each one of them. Additionally, there are two impression types, denoted by T_1 and T_2, each occurring 50% of the time. The first advertiser only cares about the first type. She obtains a reward of zero for T_1, and the advertisers pays a positive penalty τ each time a T_2 impression is assigned to her. The second advertisers admits both types, but only obtains a positive reward $Q \sim \exp(1)$ for the first type. The setup is shown in Figure A.4.

A feasible policy could assign all T_1 impressions to the first advertiser, and T_2 impressions to the second advertiser. However, such policy is not optimal. Notice that both advertisers compete for the T_1 impressions, and the first advertiser could extract a potentially high quality from them. It is not hard to see that the optimal dual variables are $v_1 = -\tau$, and $v_2 = 0$; and the optimal objective value is $\frac{1}{2}\mathbb{E}[Q - \tau]^+ = \frac{1}{2}e^{-\tau}$. Hence, it is optimal to assign those T_1 impressions with quality greater than τ to the second advertiser. Thus, no matter the value of the penalty, a fraction $e^{-\tau}$ of the total impression assigned to the first advertiser are undesired.

A.4 Computation

In this section we describe show to compute the optimal policy for our data model. The main problem resides in the computation of the dual objective in (2.4) and its gradient given a vector of dual variables.

Objective. The first term of the objective can be written as

$$\begin{aligned}
\mathbb{E}R\left(\max_{a\in\mathcal{A}_0}\{Q_a - v_a\}\right) &= \sum_{\forall T}\pi(T)\mathbb{E}\left[R\left(\max_{a\in\mathcal{A}_0}\{Q_a - v_a\}\right) \mid T\right] \\
&= \sum_{\forall T}\pi(T)\sum_{a\in T\cup T^c}\mathbb{E}\left[R\left(Q_a - v_a\right)\mathbf{1}\{Q_a - v_a \geq Q_{a'} - v_{a'}\ \forall a' \neq a\} \mid T\right] \\
&= \sum_{\forall T}\pi(T)\left(I_{T,0}(v) + \sum_{a\in T} I_{T,a}(v)\right)
\end{aligned}$$

where the first equation follows by conditioning on the type, and the second because the events are a partition of the sample space. Next, we show to compute the expectations $I_{T,a}(v)$.

Let $M_T(v) = \max_{a \in \mathcal{A}_0 \setminus T}\{-\tau_a - v_a\}$ be the maximum contract adjusted quality of the advertisers (including the outside option) that are not in the type, and $\alpha_T(v)$ the set of advertisers that verify the maximum. Then, we have that

$$
\begin{aligned}
I_{T,0}(v) &= R\left(M_T(v)\right)\mathbb{P}\{Q_a - v_a \le M_T(v)\ \forall a' \in T\} \\
&= R\left(M_T(v)\right) G_T(M_T(v) + v_T),
\end{aligned}
$$

where $G_T(\cdot)$ is the c.d.f. of Q_T, and v_T is the vector of dual variables for the advertisers in the type.

For $a \in T$, we compute the expectation by conditioning on the continuous random variable Q_a. Further, suppose that we partition the mean vector and covariance matrix in a corresponding manner. That is, $\mu_T = \left(\begin{smallmatrix} \mu_a \\ \mu_{-a} \end{smallmatrix}\right)$, and $\Sigma_T = \left(\begin{smallmatrix} \Sigma_{a,a} & \Sigma_{a,-a} \\ \Sigma_{-a,a} & \Sigma_{-a,-a} \end{smallmatrix}\right)$. For instance, μ_{-a} gives the means for the variables in $T \setminus \{a\}$, and $\Sigma_{-a,-a}$ gives variances and covariances for the same variables. The matrix $\Sigma_{-a,a}$ gives covariances between variables in $T \setminus \{a\}$ and set a (as does matrix $\Sigma_{a,-a}$). Because the marginal distribution of a multivariate normal is an univariate normal, we have that $Q_a \sim \ln \mathcal{N}(\mu_a, \Sigma_{a,a})$. We denote by $g_{T,a}(\cdot)$ the p.d.f. of Q_a. Similarly, let Q_{-a} be the vector of qualities for advertisers in $T \setminus \{a\}$. Conditioning on $Q_a = q_a$, the distribution of Q_{-a} is log-normal with mean vector $\mu_{-a} - \Sigma_{-a,a}(q_a - \mu_a)/(\Sigma_{a,a})$, and covariance matrix $\Sigma_{-a,-a} - (\Sigma_{-a,a}\Sigma_{a,-a})/(\Sigma_{a,a})$. We denote its c.d.f. by $G_{T,-a}(\cdot)$. Putting all together, we have that

$$
\begin{aligned}
I_{T,a}(v) &= \mathbb{E}\left[R\left(Q_a - v_a\right)\mathbb{P}\{Q_{a'} - v_{a'} \le Q_a - v_a\ \forall a' \neq a \mid Q_a\} \mid T\right] \\
&= \int_{v_a + M_T(v)}^{\infty} R(q_a - v_a) G_{T,-a}(q_a - v_a + v_{-a}) g_{T,a}(q_a)\, \mathrm{d}q_a,
\end{aligned}
$$

where v_{-a} is the vector of dual variables for advertisers in $T \setminus \{a\}$.

Gradient. The forward derivative of the dual objective can be written as

$$
\begin{aligned}
\nabla_a \psi(v) &= -\mathbb{P}_R\left\{Q_a - v_a > \max_{a \in \mathcal{A}_0 \setminus a}\{Q_{a'} - v'_a\}\right\} + \rho_a \\
&= -\sum_{\forall T} \pi(T)\mathbb{E}\left[\left(1 - s^*(Q_a - v_a)\right)\mathbf{1}\left\{Q_a - v_a > \max_{a \in \mathcal{A}_0 \setminus a}\{Q_{a'} - v'_a\}\right\} \mid T\right] + \rho_a \\
&= -\sum_{T: a \in T} \pi(T) P_{T,a}(v) - \sum_{\substack{T: a \notin T \\ a \in \alpha_T(v), |\alpha_T(v)| = 1}} \pi(T) P_{T,a}(v) + \rho_a,
\end{aligned}
$$

where the contributing types for the forward derivative are those where a is in, and those where a is not in but verifies exclusively the maximum of the types not in ($M_T(v)$). If two or more advertisers verify the maximum $M_T(v)$, then increasing v_a does not have an impact of the type's contribution to the objective. When $a \notin T$, the expectation is given by

$$P_{T,a}(v) = \big(1 - s^*(M_T(v))\big) G_T(M_T(v) + v_T).$$

Similarly to the objective, when $a \in T$ we have that

$$P_{T,a}(v) = \int_{v_a + M_T(v)}^{\infty} \big(1 - s^*(q_a - v_a)\big) G_{T,-a}(q_a - v_a + v_{-a}) g_{T,a}(q_a) \, \mathrm{d}q_a.$$

The backward derivative is computed in a similar fashion. The only exception is that, when $a \notin T$, and a verifies the maximum $M_T(v)$,the advertiser always contributes to the derivative regardless of the number of advertisers that attain the maximum. Hence,

$$\nabla_{-a}\psi(v) = \sum_{T: a \in T} \pi(T) P_{T,a}(v) + \sum_{\substack{T: a \notin T \\ a \in \alpha_T(v)}} \pi(T) P_{T,a}(v) - \rho_a.$$

Optimization. We solve the dual problem (2.4) using a Gradient Descent Method. At each step the objective and its objective are computed as described previously. Notice that, when multiple advertisers verify a tie, the objective is not differentiable. In this case a descent direction is constructed using the forward and backward derivatives (if possible).

Ties. For the following, we assume that the instance is not *degenerate*, that is, the variances within the types are positive, and no two advertisers are perfectly correlated. Then, within each type, non-trivial ties can only occur between the advertisers that are not in the type (we refer to the non-trivial ties as those in which multiple advertisers attain the same contract adjusted quality). Moreover, there can be at most one tie within each type, and this happens when the maximum $M_T(v)$ is verified by many advertisers, that is $|\alpha_T(v)| > 1$. With some abuse of notation, the probability of such a tie is given by $\pi(T) P_{T,\alpha_T(v)}$ and it should be split among the advertisers $\alpha_T(v)$. Note that the number of non-trivial ties is $O(T)$, and the tie-breaking rule can be computed efficiently by solving a feasible flow problem.

A.5 Fluid Limit

Exploiting our generative model we can construct a fluid model, and obtain the limiting performance of an arbitrary bid-price policy as the number of impressions grows to infinity. We first describe the fluid equations governing the dynamics of the fluid model, then we construct a solution to such system, and then prove that the stochastic algorithm satisfies the fluid equation in the limit.For simplicity, we focus our analysis on the case with no AdX, and no ties; though, a similar analysis applies to the more general case.

In the following we analyze the performance of the stochastic control policy when implementing some (sub-optimal) bid-prices v. Let ω be a sample path, $J_n(\omega)$ be the cumulative yield collected up to impression n, and $S_{n,a}(\omega)$ the number of impressions assigned to advertiser a up to time n. We extended the previous definitions for an arbitrary time, by taking their linear interpolations, so that they are continuous. The previous functions are random elements on $\mathbb{C}[0,\infty)$. We shall construct the fluid limit by scaling capacity and time proportionally to infinity, and considering a continuous flow of impressions arriving during an horizon of length 1. More formally, we define $\bar{S}_a(t) = \lim_{N\to\infty} N^{-1} S_{tN,a}(\omega)$, which can be interpreted as the fraction of impressions assigned to advertiser a by time t. Similarly, we define $\bar{J}(t) = \lim_{N\to\infty} N^{-1} J_{tN}(\omega)$ as the cumulative yield up to time t. We are interested in computing $\bar{J}(1)$, the total limiting yield of the algorithm under bid-prices v.

When capacity is scaled, each advertiser has a capacity of ρ_a, and the fluid model should satisfy the following differential equations

$$J'(t) = \sum_{a\in\mathcal{A}(t)} \mathbb{E}\Big[Q_a \mathbf{1}\big\{a = \arg\max_{a'\in\mathcal{A}(t)}\{Q_{a'} - v_{a'}\}\big\}\Big], \tag{A.8a}$$

$$S'_a(t) = \mathbb{P}\Big\{a = \arg\max_{a'\in\mathcal{A}(t)}\{Q_{a'} - v_{a'}\}\Big\}, \quad \forall a \in \mathcal{A}_0 \tag{A.8b}$$

$$\mathcal{A}(t) = \big\{a \in \mathcal{A}_0 : \bar{S}_a(t) < \rho_a\big\}, \tag{A.8c}$$

with the initial conditions $S_a(0) = 0$, and $J(0) = 0$. In (A.8c), $\mathcal{A}(t)$ is the set of advertisers that are yet to be fulfilled (including advertiser 0, which is fulfilled when the time comes all impressions should be assigned directly to the advertisers), and (A.8b) determines the rate at which impressions are assigned to each advertiser. When one advertiser is fulfilled and the fraction of impressions $\bar{S}'_a(t)$ reaches its capacity ρ_a, it is excluded from $\mathcal{A}(t)$, and its rate is driven to zero. Finally, (A.8a) determines the rate at which yield is generated.

It is not hard to see that the solution to the fluid equations (A.8) is piecewise linear, and continuous. We construct a solution as follows. Let an *epoch*, denoted by t^k, be the time in which the contract of any advertiser is fulfilled (including advertiser 0). The horizon $[0,1]$ is partitioned in consecutive *pieces*, each culminating with an epoch. The k^{th} piece spans the interval $[t^k, t^{k+1})$, and has a length of $\Delta^k = t^{k+1} - t^k$. Since one advertiser is fulfilled at each epoch, there are at most $A + 1$ pieces.

Let $\mathcal{A}^k$ be set of advertisers yet to be satisfied at the beginning of stage k as given by (A.8c), r_a^k be the service rate for advertiser a during stage k as given by (A.8b), and y^k be the yield rate during stage k as given by (A.8a). The length of a stage is determined by the advertiser that is first fulfilled. Once determined, the solution is constructed recursively as follows

$$\Delta^k = \min_{a\in\mathcal{A}^k}\left\{\frac{\rho_a - S_a(t_k)}{r_a^k}\right\}, \tag{A.9a}$$

$$t^{k+1} = \Delta^k + t^k,$$

$$S_a(t^{k+1}) = S_a(t^k) + r_a^k\Delta^k,$$

$$J(t^{k+1}) = J(t^k) + y^k\Delta^k,$$

$$\mathcal{A}^{k+1} = \mathcal{A}^k \setminus a^k,$$

where a^k is the advertiser that verifies the minimum in (A.9a), and initially $t^0 = 0$, and $\mathcal{A}^0 = \mathcal{A}_0$. The functions S_a and J are obtained as the linear interpolation of the values at the endpoints of the interval. Fortunately, the rates can be easily obtained by evaluating the dual objective and its gradient. Let $\psi^k(v)$ be the objective in (2.4) when restricting to the set of advertisers $\mathcal{A}^k$. Then, we have that $r^k = \rho - \nabla\psi^k(v)$, and $y^k = \psi^k(v) - v\cdot\nabla\psi^k(v)$.

We conclude by showing that functions obtained are actually the fluid limit of the stochastic process induced by the algorithm.

Proposition A.5. *The fluid limits $\bar{S}_a(t)$, and $\bar{J}(t)$ are a solution to the fluid equations* (A.8).

Proof. Proof. Consider the sequence of random elements $\bar{S}_{N,a}(t,\omega) = N^{-1}S_{tN,a}(\omega)$, and $\bar{J}_N(t,\omega) = N^{-1}J_{tN}(\omega)$. We would like to show that the previous sequences are tight. By Theorem 8.3 in Billingsley (1968), a sequence of random elements $\{X_N\}$ in $\mathbb{C}[0,\infty)$ is tight iff (i) $\{X_N(0)\}$ is tight, and (ii) for all $\epsilon > 0$ and $\eta > 0$, there is a $\delta > 0$ and an integer N_0 such that $\mathbb{P}\{\sup_{t\le t'\le t+\delta}|X_N(t') - X_N(t)| \ge \epsilon\} \le \eta$ for $N \ge N_0$.

The first condition is trivially satisfied for both sequences. Disregarding integrality issues, which are not important in the analysis, and using the fact that the processes are non-decreasing, we have that

$$\sup_{t\le t'\le t+\delta}|\bar{S}_{N,a}(t') - \bar{S}_{N,a}(t)| \le \frac{1}{N}\left(S_{(t+\delta)N,a} - S_{tN,a}\right) \le \delta.$$

Thus, by picking $\delta < \epsilon$ the second condition is satisfied for the number of impressions assigned. For the yield processes, employing Markov's inequality, and the bound $\mathbb{E}R_n^2 \le A^2\max_a\{\mathbb{E}Q_a^2\}$ for the yield in any single period, we obtain

$$\frac{1}{\delta}\mathbb{P}\left\{J_{(t+\delta)N} - J_{tN} \ge N\epsilon\right\} \le \frac{1}{\delta N^2\epsilon^2}\mathbb{E}\Big[\sum_{n=tN}^{(t+\delta)N} R_n\Big]^2 \le \delta\frac{A^2\max_a\{\mathbb{E}Q_a^2\}}{\epsilon^2},$$

which can be bounded from above by η by picking a small enough δ, and $N > 1/\delta$.

Next, we show that the fluid limit converges to the solution of the equation (A.8b). Before proceeding we state some definitions. Let the stopping time $n_N^k = \inf\{n : S_{n,a} \ge C_a$ for some $a \in$

$\mathcal{A}_N^{k-1}\}$ be the time in which the contract of the k^{th} advertisers is fulfilled. In our previous terminology, n_N^k is the k^{th} epoch and the beginning of the k^{th} piece. Similarly, we let $\mathcal{A}_N^k = \{a \in \mathcal{A}_0 : S_{n_N^k,a} < C_a\}$ to be the set of advertisers that are active during the k^{th} piece. The initial values are given by $n_N^0 = 1$ and $\mathcal{A}_N^0 = \mathcal{A}_0$.

We intend to show that $N^{-1}n_N^k \to t^k$, $\mathcal{A}_N^k \to \mathcal{A}^k$, and $N^{-1}S_{n_N^k,a} \to S_a(t^k)$ as $N \to \infty$ in an almost sure sense. We proceed by induction in k. The base case follows trivially. Suppose that our claim holds for k, we intend to show that it holds for $k+1$. By the definition of the $(k+1)^{\text{th}}$ epoch, it should be the case that $S_{n_N^{k+1}-1,a} < C_a$ for all $a \in \mathcal{A}_N^k$ and $S_{n_N^{k+1},a} \geq C_a$ for exactly one advertiser $a \in \mathcal{A}_N^k$. Notice that the number of impressions assigned to ad a up to the stopping time can be written as

$$\frac{S_{n_N^{k+1},a}}{N} = \frac{S_{n_N^k,a}}{N} + \frac{n_N^{k+1} - n_N^k}{N} \frac{\sum_{n=n_N^k+1}^{n_N^{k+1}} I_{n,a}}{n_N^{k+1} - n_N^k}, \tag{A.10}$$

where the summands on the right-hand size are i.i.d. bernoulli random variables with success probability $\mathbb{P}\Big\{a = \arg\max_{a' \in \mathcal{A}_N^k} \{Q_{a'} - v_{a'}\}\Big\}$. From the induction hypothesis, together with the Triangular Strong Law of Large Numbers, we get that (A.10) goes to $S_a(t^k) + (\lim N^{-1}n_N^{k+1} - t^k)r_a^k$ since the success probability converges to r_a^k. It is not hard to see that the same limit holds for $N^{-1}S_{n_N^{k+1}-1,a}$. Thus, we have that for all advertisers yet to be satisfied it should be the case that $S_a(t^k) + (\lim N^{-1}n_N^{k+1} - t^k)r_a^k \leq \rho_a$, and for at least one advertiser $S_a(t^k) + (\lim N^{-1}n_N^{k+1} - t^k)r_a^k \geq \rho_a$. Combining this expressions we get that in the limit the stopping time should satisfy

$$\lim \frac{n_N^{k+1}}{N} = t^k + \min_{a \in \mathcal{A}^k} \left\{ \frac{\rho_a - S_a(t_k)}{r_a^k} \right\} = t^{k+1}.$$

In turn this implies that $N^{-1}S_{n_N^{k+1},a} \to S_a(t^{k+1})$ and $\mathcal{A}_N^{k+1} \to \mathcal{A}^{k+1}$, thus concluding the inductive step. $\square$

Appendix B

Appendix to Chapter 3

B.1 Proof of Statements

B.1.1 Proof of Proposition 3.1

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

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

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

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

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

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

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

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

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

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

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

B.1.2 Proof of Theorem 3.1

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

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

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

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

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

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

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

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

B.1.3 Proof of Theorem 3.2

Exploiting the fact that the dual objective is convex and differentiable, one may write the equilibrium condition (3.4) as a Nonlinear Complementarity Problem (NCP). From the optimality conditions of the dual, it should be the case that for each type $\theta \in \Theta$, one of the following alternatives holds

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

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

$$\mu_\theta^* \geq 0 \quad \perp \quad \alpha_\theta \eta s_\theta G_\theta(\boldsymbol{\mu}^*, r) \leq b_\theta, \qquad \forall \theta \in \Theta, \tag{B.1}$$

where $\perp$ indicates a complementarity condition between the multiplier and the expenditure, that is, at least one condition should be met with equality. From item (ii) of Lemma B.3 we have that the mapping $\mathbf{G}$ is differentiable. The latter, together with the P-matrix assumption, allows one to invoke (Facchinei and Pang, 2003*a*, Proposition 3.5.10) and conclude that the NCP (B.1) has at most one solution.

B.1.4 Proof of Proposition 3.2

Fix $r \geq 0$. The existence of the equilibrium follows from Theorem 3.1. The uniqueness follows from the fact that Assumption 3.1 is automatically satisfied in the present case from item *iii.*) of Lemma B.3. We next derive the characterization of the FMFE.

Suppose first that $\alpha\eta s G_0(r) < b$. By Lemma B.3 *iii.*), increasing the multiplier cannot increase the expenditure, and no solution to the NCP with $\mu > 0$ exists. Thus $\mu^* = 0$ is the unique equilibrium multiplier. Suppose now that $\alpha\eta s G_0(r) \geq b$, then advertisers need to shade their bids by picking a non-negative equilibrium multiplier μ^* that solves for $\alpha\eta s G(\mu^*, r) = b$ (and such a solution exists by the proof of Theorem 3.1). Noting that $(1+\mu^*)G(\mu^*, r) = G_0((1+\mu^*)r)$ concludes the proof.

B.1.5 Proof of Theorem 3.3

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

Using (3.6), the publisher's profit as a function of the reserve price and the multiplier can be written as $\Pi(\mu, r) = \alpha\lambda\eta s G(\mu, r) - \eta c I(\mu, r)$, with $I(\mu, r)$ the probability that the impression is won by some advertiser in the exchange when advertisers employ a multiplier μ and the publisher sets a reserve price r. Note that $I(\mu, r) = I_0((1+\mu)r)$, where by Lemma B.2(i), $I_0(r) = 1 - e^{-\alpha\lambda s \bar{F}_v(r)}$ is the probability that the impression is won in the exchange by truthful advertisers . The publisher's problem amounts to solving $\max_{r \geq 0} \Pi(\mu(r), r)$, where $\mu(r)$ is the unique equilibrium multiplier for price r.

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

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

$$\frac{\partial \Pi}{\partial r}(0, r) = \alpha\lambda\eta s \left(G_0'(r) + ce^{-\alpha\lambda s \bar{F}_v(r)} f_v(r) \right) = \alpha\lambda\eta s \bar{F}_v(r) e^{-\alpha\lambda s \bar{F}_v(r)} \left(1 - \frac{r-c}{r}\xi(r) \right), \tag{B.2}$$

where the second equation follows by Lemma B.2(ii). The previous expression vanishes at r_c^*. Notice that the leading terms in the derivative are non-negative, and by the IGFR assumption, it follows that the derivative is non-negative for $r < r_c^*$ and non-positive for $r > r_c^*$. Thus, $\Pi(0, r)$ is strictly quasi-concave on $[\underline{V},\ \bar{V}]$.

(ii) *Then the set* $\mathcal{R}^*$ *is a closed bounded interval.* The proof follows by noticing that setting $c = 0$ in (B.2) implies that $G_0(r)$ is strictly quasi-concave in r (in the interval $[\underline{V},\ \bar{V}]$). Since $\mathcal{R}^*$ is an upper-level set of G_0, and G_0 is continuous we get that $\mathcal{R}^*$ is a closed interval. The boundedness of $\mathcal{R}^*$ follows from Lemma B.3 *iv.*).

(iii) *The equilibrium multiplier verifies* $\mu > 0$ *for* r *in the interior of* $\mathcal{R}^*$*, and zero otherwise.* That $\mu = 0$ outside the interior of $\mathcal{R}^*$ follows directly from the statement of Proposition 3.2. By the strict quasi-concavity of $G_0(r)$ in r, $\alpha\eta s G_0(r) > b$ for r in the interior of $\mathcal{R}^*$, so by Proposition 3.2, $\mu > 0$ for r in this set.

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

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

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

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

where the second equation follows from the fact that $G'_r(\mu, r) = G'_0((1+\mu)r)$ and $G'_\mu(\mu, r) = G'_0((1+\mu)r)r/(1+\mu) - G(\mu, r)/(1+\mu)$, and the inequality from the fact that $G(\mu, r) = g \geq 0$, and that $G(\mu, r)$ is decreasing in μ for fixed r by Lemma B.3 *iii.*).

Now, we study the two cases.

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

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

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

B.1.6 Proof of Theorem 3.4

We use the following lemma to prove the theorem.

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

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

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

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

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

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

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

□

Proof of Theorem 3.4. Fix $r \ge 0$ and let $\Pi(\mu,\eta)$ be the publisher's profit as a function of the rate of impressions, and the equilibrium multiplier, respectively. The publisher's problem amounts to solving $\max_{0\le\eta\le\bar\eta} \Pi(\mu(\eta),\eta)$. We use Proposition 3.2 to analyze the dependence of the FMFE multiplier on the rate of impressions, $\mu(\eta)$. When $\eta < \eta^0$ advertisers bid truthfully and the equilibrium multiplier is $\mu(\eta) = 0$. When $\eta \ge \eta^0$ advertisers shade their bids so as to deplete their budgets in expectation and the multiplier is the unique solution of the equation $\alpha\eta s G_0\Big((1+\mu)r\Big) = (1+\mu)b$. We deduce that

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

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

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

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

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

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

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

where the second equation follows from writing $G(\mu, r) = G_((1+\mu)r)/(1+\mu)$, and using the fact that $\alpha\eta s G(\mu, r) = b$. Note that from Lemma B.3 point *iii.*) one gets that $G_\mu'(\mu, r) < 0$, which allows one to conclude that the multiplier is increasing with the rate of impressions. In the remainder of the proof we show that the direct effect dominates the indirect effect.

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

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

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

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

$$\begin{aligned}
\lambda b - \eta r I_0 &= \eta \mathbb{E}\left[\mathbf{1}\{\hat{W}_{1:\hat{M}} \geq r\}\Big(\max\{\hat{W}_{2:\hat{M}}, r\}\Big)\right] - \eta r \mathbb{P}\{\hat{W}_{1:\hat{M}} \geq r\} \\
&= \eta \mathbb{E}\left[\mathbf{1}\{\hat{W}_{1:\hat{M}} \geq r\}\Big(\hat{W}_{2:\hat{M}} - r\Big)^+\right] = \eta \mathbb{E}\left[\mathbf{1}\{\hat{W}_{2:\hat{M}} \geq r\}\Big(\hat{W}_{2:\hat{M}} - r\Big)\right] \qquad \text{(B.3)} \\
&= \eta \mathbb{E}\left[\mathbf{1}\{\hat{W}_{2:\hat{M}} \geq r\}\hat{W}_{2:\hat{M}}\right] - \eta r \mathbb{P}\{\hat{W}_{2:\hat{M}} \geq r\},
\end{aligned}$$

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

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

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

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

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

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

For the second term in parenthesis, we proceed in a similar fashion. Using the joint distribution of the highest and second-highest bid (see, e.g., David and Nagaraja (2003)) we have that the probability that the impression is won and the reserve price is paid is given by $\mathbb{P}\{\hat{W}_{1:\hat{M}} \geq r, W_{2:\hat{M}} < r\} = (\alpha\lambda s)\bar{F}_v(1 - I_0)$. Thus, we obtain that

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

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

Putting it all together, one obtains

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

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

B.1.7 Proof of Corollary 3.2

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

Let $\Pi(\eta) = \max_{r\geq 0} \Pi(\mu(r,\eta), r, \eta)$ be the objective of the inner optimization. By Theorem 3.3 we have that

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

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

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

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

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

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

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

B.1.8 Proof of Theorem 3.5

Fix α in $(0, 1]$. In view of Theorem 3.3, advertisers bid truthfully at the optimal reserve price. Note that the generalized failure rate of the value distribution (3.7) is $\xi_{v(\alpha)}(x) = \xi_v(x/\sigma(\alpha))$, and the failure rate is $h_{v(\alpha)}(x) = h_v(x/\sigma(\alpha))/\sigma(\alpha)$. Let $\Pi_0(r,\alpha)$ denote the publisher's profit when advertisers bid truthfully, which after integrating by parts is given by

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

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

For any given α, by Theorem 3.3, the optimal reserve price is unique, bidders bid truthfully at the optimal reserve, and the optimal profit is given by $\Pi_0(\max\{r_c^*(\alpha),\bar{r}(\alpha)\},\alpha)$ (with some abuse of notation, we make the dependence on α explicit). The result follows by separately analyzing the two possible cases: (1) $r_c^*(\alpha)$ is the optimal reserve price; and (2) $\bar{r}(\alpha)$ is the optimal reserve price. With some abuse of notation, let $G_0(r,\alpha)$ denote the expected expenditure-per-auction in the absence of budget constraints when advertisers bid truthfully.

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

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

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

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

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

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

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

$$b = \alpha\eta s G_0(\bar{r}(\alpha), \alpha) = \eta s \int_{\alpha\bar{r}(\alpha)}^{\infty} \bar{F}_v(x)\Big(\xi_v(x) - 1\Big)e^{-\alpha\lambda s \bar{F}_v(x)}\,\mathrm{d}x. \tag{B.4}$$

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

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

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

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

To conclude that the profit is non-increasing we shall show that both terms are non-positive. Indeed, the partial derivative w.r.t. the matching probability is $\partial\Pi_y/\partial\alpha = -c\lambda s\eta\bar{F}_v(\bar{y}(\alpha))e^{-\alpha\lambda s\bar{F}_v(\bar{y}(\alpha))} \leq 0$. Similarly, the partial derivative w.r.t. the scaled reserve price is $\partial\Pi_y/\partial y = c\eta\alpha\lambda s f_v(\bar{y}(\alpha))e^{-\alpha\lambda s\bar{F}_v(\bar{y}(\alpha))} \geq 0$. Finally, invoking the Implicit Function Theorem we get from equation (B.4) that the total derivative of the scaled optimal reserve price is

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

For the last inequality recall that, by assumption, $\bar{r}(\alpha) > r^*_c(\alpha)$, which implies that $\bar{y}(\alpha) > r^*_{\alpha c} \geq r^*_0$. Using the IGFR assumption we obtain that $\xi_v(y(\alpha)) > \xi_v(r^*_0) \geq 1$, and then both the numerator and the denominator are non-negative. Hence, the optimal reserve price is non-increasing with the matching probability.

Putting it all together. The optimal profit is given by

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

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

B.1.9 General properties of the expenditure function

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

Lemma B.2. *i.) The distribution of the maximum competing bid when $x \geq r$ can be characterized by*

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

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

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

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

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

$$F_d(x;\boldsymbol{\mu}) = \mathbb{E}\left[F_w(x;\boldsymbol{\mu})^{\hat{M}}\right] = \exp\left\{-\mathbb{E}[\alpha_\Theta\lambda s_\Theta]\bar{F}_w(x;\boldsymbol{\mu})\right\}, \tag{B.5}$$

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

ii.) The expenditure function can be written as

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

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

□

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

Lemma B.3. *i.) For any $\boldsymbol{\mu}$, the maximum bid $D \sim F_d(\boldsymbol{\mu})$ is integrable, that is, $\mathbb{E}[D] < \infty$.*

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

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

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

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

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

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

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

ii.) By Lemma B.2(i), the distribution of the maximum competing bid when $x \geq r$ is given by $F_d(x; \boldsymbol{\mu}) = \exp\left\{-\mathbb{E}[\alpha_\Theta \lambda s_\Theta] \sum_\theta \hat{p}_\theta \bar{F}_{v_\theta}((1+\mu_\theta)x)\right\}$, where $\hat{p}_\theta = \mathbb{P}_{\hat{\Theta}}\{\theta\}$. Since the cumulative distribution of values is differentiable, the distribution of the maximum bid is differentiable w.r.t. x and $\boldsymbol{\mu}$. Indeed, its partial derivatives are given by $\partial F_d/\partial \mu_\theta = F_d(x; \boldsymbol{\mu})\mathbb{E}[\alpha_\Theta \lambda s_\Theta]\hat{p}_\theta x f_{v_\theta}((1+\mu_\theta)x)$, and $\partial F_d/\partial x = F_d(x; \boldsymbol{\mu})\mathbb{E}[\alpha_\Theta \lambda s_\Theta] \sum_\theta \hat{p}_\theta (1+\mu_\theta) f_{v_\theta}((1+\mu_\theta)x)$. Moreover, the second derivatives of the distribution of the maximum bid are continuous because densities $f_{v_\theta}(\cdot)$ are continuously differentiable.

By Lemma B.2(ii), the expenditure function can be written as $G_\theta(\boldsymbol{\mu}, r) = r\bar{F}_{v_\theta}((1+\mu_\theta)r)F_d(r; \boldsymbol{\mu}) + \int_r^{\bar{V}} x\bar{F}_{v_\theta}((1+\mu_\theta)x)\,\mathrm{d}F_d(x; \boldsymbol{\mu})$, which is clearly differentiable in r. Moreover, for any $\theta' \in \boldsymbol{\Theta}$ the first term is differentiable w.r.t. $\mu_{\theta'}$, while the integrand is continuously differentiable. We conclude by an application of Leibniz's integral rule, which holds because $[\underline{V}, \bar{V}] \times U$ is bounded.

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

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

where

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

and

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

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

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

$$\frac{\partial G_\theta}{\partial \mu_\theta} = -\sum_{\theta' \neq \theta}(1+\mu_{\theta'})\hat{p}_{\theta'}\langle f_\theta, f_{\theta'}\rangle - \hat{p}_\theta\langle f_\theta, \bar{F}_\theta\rangle - r f_\theta(r)F_d(r;\boldsymbol{\mu}), \tag{B.6}$$

which is strictly negative.

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

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

When $r = 0$ we have that

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

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

□

B.2 Sufficient Conditions for P-matrix Assumption to Hold

We establish here sufficient conditions for Assumption 3.1, that was required for uniqueness of a FMFE, to hold.

Proposition B.1. *The P-matrix condition (Assumption 3.1) holds in either of the following cases.*

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

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

The *positively homogeneous* condition in *ii.)* imposes that there is some $n \geq 0$ such that $h_v(ax) = a^n h_v(x)$ for all $x \in \text{dom}(V)$ and $a > 0$. This property is satisfied by distributions whose failure rates are power functions; such as the exponential, Weibull, and Rayleigh distributions. Additionally, it is not difficult to show from first principles that, for the case of two types with common value distribution, Assumption 3.1 holds when values are uniformly distributed with support $[0, \bar{V}]$.

B.2.1 Proof of Proposition B.1

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

i.) In this case $J_{\mathbf{G}} = \partial G(\mu, r)/\partial\mu$, and the result follows directly from item *iii.)* of Lemma B.3.

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

Step 1. In the proof of item *iii.)* from Lemma B.3 we characterized the diagonal entries of the Jacobian, that is, $\partial G_\theta(\boldsymbol{\mu}, r)/\partial\mu_\theta$. Using a similar notation, we characterize the off-diagonal entries as follows.

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

$$\begin{aligned}\frac{\partial G_\theta}{\partial \mu_{\theta'}} &= r\bar{F}_v((1+\mu_\theta)r)\frac{\partial F_d}{\partial \mu_{\theta'}} + \frac{\partial}{\partial \mu_{\theta'}}\int_r^{\bar{V}} x\bar{F}_v((1+\mu_\theta)x)\frac{\partial F_d}{\partial x}\,\mathrm{d}x \\ &= r\bar{F}_v((1+\mu_\theta)r)\frac{\partial F_d}{\partial \mu_{\theta'}} + \int_r^{\bar{V}} x\bar{F}_v((1+\mu_\theta)x)\frac{\partial^2 F_d}{\partial \mu_{\theta'}\partial x}\,\mathrm{d}x \\ &= -\int_r^{\bar{V}} \frac{\partial}{\partial x}\left(x\bar{F}_v((1+\mu_\theta)x)\right)\frac{\partial F_d}{\partial \mu_{\theta'}}\,\mathrm{d}x \\ &= (1+\mu_\theta)\hat{p}_{\theta'}\langle f_\theta f_{\theta'}\rangle - \hat{p}_{\theta'}\langle f_{\theta'}\bar{F}_\theta\rangle, \end{aligned} \tag{B.7}$$

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

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

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

From item *iii.)* of Lemma B.3 one concludes that the principal minors $J|_{\{1\}}$ and $J|_{\{2\}}$ are negative (they are, in fact, negative scalars), so the corresponding principal minors of $-\mathbf{G}$ are positive. The determinant of the remaining minor $J|_{\{1,2\}}$ is that of the whole Jacobian, which is given by

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

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

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

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

Appendix C

Appendix to Chapter 4

C.1 Proof of Statements for Synchronous Campaigns

C.1.1 Proof of Theorem 4.1

Proof. The proof follows directly from combining the lower bound in Proposition 4.1 and the upper bound in Proposition 4.2. □

C.1.2 Proof of Proposition 4.1

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

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

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

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

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

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

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

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

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

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

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

Term 2. Define $\tilde{I}_{k,i}$ as the number of auctions that advertiser k^{th} participates until agent i^{th} runs out of budget, that is, $\tilde{I}_{k,i} = \sum_{n=1}^{\tilde{N}_i} M_{n,k}$. Using this notation we obtain that the number of auctions the k^{th} advertiser participates until someone runs out of budget can be alternatively written as $\tilde{I}_k = \sum_{n=1}^{\min_i \tilde{N}_i} M_{n,k} = \min_i \sum_{n=1}^{\tilde{N}_i} M_{n,k} = \min_i \tilde{I}_{k,i}$. Using this identity we obtain that

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

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

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

where the second inequality follows from the fact that the expected expenditure in the FMFE never exceeds the budget, that is, $\alpha_i \eta s G_i(\boldsymbol{\mu}) \leq b_i$, and by setting $\bar{b} = \max_i b_i$; and the last because $\alpha_k \eta s = O(\bar{b})$ from Assumption 4.1. □

C.1.3 Proof of Proposition 4.2

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

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

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

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

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

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

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

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

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

$$x_n \le m_{n,k}. \tag{C.1c}$$

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

Consider the continuous relaxation of the hindsight program (C.1) in which we replace the integrality constraints by $0 \le x_n \le m_{n,k}$. Let μ_k be the equilibrium multiplier of the FMFE for k^{th} advertiser. Introducing dual variables $\mu \ge 0$ for the budget constraint and $z_n \ge 0$ for the constraints $x_n \le m_{n,k}$, we get by weak duality that

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

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

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

Term 2. Using the same argument that in the proof of Proposition 4.1 we obtain that

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

Term 3. In order to bound the term $\mathbb{E}[\alpha_k \eta s - \tilde{I}_k]^+$ we shall consider a second alternate system (II) in which all advertisers (including k^{th}) implement the FMFE strategies and the initial budgets for every advertiser is discounted are discounted in advance to take into account the potential impact that the k^{th} advertiser may have on its competitors. Since a competitor can spend at most

$\bar{V}$ in an auction, this potential impact can be upper bounded by $\bar{V}$ times the number of auctions in which she competes against the k^{th} advertiser. That is, we set the budgets to $b_i^{\text{II}} = \big(b_i - \bar{V}T_{k,i}\big)^+$ for all $i \neq k$ where $T_{k,i} = \sum_{n=1}^{N(s)} M_{n,k}M_{n,i}$ is the number of auctions in which k and i compete together. Defining by $\tilde{N}^{\text{II}}$ be the first auction in which some advertiser runs out of budget in the alternate system (II), we obtain using a coupling argument that

$$\tilde{N}^{\text{II}} \leq \tilde{N} \quad \text{(a.s.).} \tag{C.2}$$

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

Let $\tilde{I}_{k,i}^{\text{II}}$ be the number of auctions that advertiser k^{th} participates until agent i^{th} runs out of budget, and $\tilde{I}_k^{\text{II}}$ be the number of auctions the k^{th} advertiser participates until someone runs out of budget. Equation (C.2) implies that $\tilde{I}_k^{\text{II}} \leq \tilde{I}_k$ almost surely, which implies using the steps in the proof of Proposition 4.1 that

$$\begin{aligned}\mathbb{E}\left[\alpha_k \eta s - \tilde{I}_k\right]^+ &\leq \mathbb{E}\left[\alpha_k \eta s - \tilde{I}_k^{\text{II}}\right]^+ \leq \mathbb{E}\left[\max_i \{\alpha_k \eta s - \tilde{I}_{k,i}^{\text{II}}\}^+\right] \\ &\leq \max_i \left\{\alpha_k \eta s - \mathbb{E}\tilde{I}_{k,i}^{\text{II}}\right\}^+ + \sqrt{\sum_i \mathrm{Var}[\tilde{I}_{k,i}^{\text{II}}]},\end{aligned} \tag{C.3}$$

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

For the mean we obtain that

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

where the inequality follows from property (v) of Lemma C.1, and the last equality from the fact that $T_{k,i}$ is Poisson with mean $\alpha_k\alpha_i\eta s$.

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

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

For the first term we use that from property (vi) of Lemma C.1 there exists non-negative constants C_0, C_1 such that $\mathrm{Var}[\tilde{I}_{k,i}^{\text{II}} \mid b_i^{\text{II}}] \leq C_0 + C_1 b_i^{\text{II}}$, together with the fact that $b_i^{\text{II}} \leq b_i$ to obtain that

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

For the second term we combine the upper and lower bounds on property (v) of Lemma C.1 to obtain that there exists some different non-negative constants C_0, C_1 such that $\left|\mathbb{E}[\tilde{I}_{k,i}^{\text{II}} \mid b_i^{\text{II}}] - C_1 b_i^{\text{II}}\right| \le C_0$. Together with Lemma C.2 we obtain that

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

where the third inequality follows from the fact that truncation reduces variance, that is, for any random variable X and constant x we have that $\text{Var}(x - X)^+ \le \text{Var} X$ (see, e.g., Liu and Li (2009)); and last inequality follows from the fact that $\alpha_k \le 1$ and $\alpha_i \eta s = O(b_i)$ from Assumption 4.1. Combining the bounds for the first and second terms we get that $\text{Var}[\tilde{I}_{k,i}^{\text{II}}] = O(b_i)$.

We put everything together by plugging in our bounds for the mean and variance of the stopping times in the main bound (C.3) and dividing by the expected number of impressions in the horizon to obtain that

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

where the second inequality follows from the fact that the maximum of a sum is dominated by the sum of the maximums and by setting $\bar{b} = \max_i b_i$, the third inequality because expected expenditure in the FMFE never exceeds the budget, that is, $\alpha_i \eta s G_i(\boldsymbol{\mu}) \le b_i$, and because the second term is $O(\alpha_k)$ because the expected expenditure is bounded from below from Assumption 4.1; and the last because $\alpha_k \eta s = O(\bar{b})$ from Assumption 4.1 too. $\square$

C.1.4 Additional Results

Lemma C.1 (Identities and Bounds for Stopping Times)**.** *Suppose that Assumption 4.1 holds. We have that*

(i) $\frac{b_k}{G_k(\boldsymbol{\mu})} \le \mathbb{E}[\tilde{I}_{k,k}] \le \frac{b_k + \bar{V}}{G_k(\boldsymbol{\mu})}$ *for all advertiser* k,

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

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

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

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

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

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

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

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

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

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

Item (ii). The variance is bounded from above by $\mathrm{Var}(\tilde{I}_{k,k}) \leq (b_k + \bar{V})\mathrm{Var}(Z'_k)/\mathbb{E}[Z'_k]^3 + \bar{V}^2/\mathbb{E}[Z'_k]^2$ (use Wald's second identity to get $\mathbb{E}[C_{\tilde{I}_{k,k},k} - \tilde{I}_{k,k}\mathbb{E}Z'_k]^2 = \mathrm{Var}(Z'_k)\mathbb{E}\tilde{I}_{k,k}$). The result follows because expenditures are bounded from above by V and because expected expenditures are bounded from below by Assumption 4.1.

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

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

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

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

where the first inequality follows from Minkowski's inequality, the third equality follows from Wald's second identity ($\mathbb{E}[\tilde{I}_{k,i} - \alpha_k \tilde{N}_i]^2 = \alpha_k(1-\alpha_k)\mathbb{E}\tilde{N}_i$) and the definition of variance; and the last bounds from items (iii) and (iv) from this lemma and Assumption 4.1's restriction of matching probabilities.

Item (vii). Note that when $\mu_k > 0$ we have by the FMFE characterization that the advertiser is budget constrained and thus $G_k(\boldsymbol{\mu}) = b_k/(\alpha_k \eta s) \geq \underline{g}$ by Assumption 4.1. Next, we show that the expenditure is lower bounded when the advertiser is not shading her bids. Recall from the proof of Theorem 3.1 that FMFE multipliers are upper bounded by $\mu_k \leq \alpha_k \eta s \overline{V}/b_k \leq \underline{g}\overline{V}$ with the second inequality by Assumption 4.1. Let $D_{-k}(\boldsymbol{\nu}_{-k}) = \max_{i \neq k, M_i = 1}(V_i/(1+\nu_i)) \vee r$ be the maximum competing bid observed by the k^{th} advertiser when competitors shade their bids according to $\boldsymbol{\nu}_{-k}$. The expected expenditure of the k^{th} advertiser is bounded from below by

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

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

Lemma C.2. *Let X and Y be two random variables, then $\|X - \mathbb{E}X\|_2 \leq \|X-Y\|_2 + \|Y - \mathbb{E}Y\|_2$.*

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

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

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

C.2 Proof of Statements for Asynchronous Campaigns

C.2.1 Proof of Theorem 4.2

Write the relative expected payoff of the deviation as

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

Proposition 4.7 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 4.6 implies that the limsup of the second term of the numerator is bounded from above by $\bar{J}_\theta^{\mathrm{F}}(F_d)$ and Proposition 4.5 implies that the liminf of the second term of the denominator is bounded from below by $\bar{J}_\theta^{\mathrm{F}}(F_d)$. The result follows.

C.2.2 Proof of Proposition 4.3

The second part of the statement is direct. We prove the first one. Let $\mathbf{f} : [0,1]^{|\boldsymbol{\Theta}|} \to [0,1]^{|\boldsymbol{\Theta}|}$ 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 (4.3) we have that

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

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 $\theta \in \boldsymbol{\Theta}$, and let $\mathbf{q}$ and $\mathbf{q}'$ be two distinct vectors of active probabilities. Let X_θ^{MF} and $X_\theta'^{\mathrm{MF}}$ 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 $a_{n,k} = \mathbf{1}\{U_{n,k} \le q_{\theta_{n,k}}\}$, while for Y_θ' is $q_{n,k}' = \mathbf{1}\{U_{n,k} \le q_{\theta_{n,k}}'\}$, 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., $\mathcal{L}(X_\theta^{\mathrm{MF}}) = \mathcal{L}(Y_\theta)$ and $\mathcal{L}(X_\theta'^{\mathrm{MF}}) = \mathcal{L}(Y_\theta')$. Let $A = \left\{\cup_{n=1}^{N_\theta(s_\theta)} \cup_{k=1}^{M_n} a_{n,k} \neq a_{n,k}'\right\}$ 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

$$\left\|\mathcal{L}(X_\theta^{\mathrm{MF}}) - \mathcal{L}(X_\theta^{\mathrm{MF}\prime})\right\|_{TV} \le \mathbb{P}\{Y_\theta \neq Y_\theta'\} \le \mathbb{P}\{A\}.$$

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 $\Big(\min\{q_{\theta_{n,k}}, q_{\theta_{n,k}}'\}, \max\{q_{\theta_{n,k}}, q_{\theta_{n,k}}'\}\Big]$, an event occurring with probability $|q_{\theta_{n,k}} - q_{\theta_{n,k}}'|$. Taking expectations over the types, we have that $\mathbb{P}\{a_{n,k} \neq a_{n,k}'\} = \sum_{\theta\in\Theta} \mathbb{P}\{\hat{\Theta}_{n,k} = \theta\}|q_\theta - q_\theta'| \le \|\mathbf{q} - \mathbf{q}'\|_\infty$. Using a union bound together with the independence assumption on the primitives, one may bound the probability of the event A by

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

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

C.2.3 Proof of Proposition 4.4

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 $\mathbf{f}^\infty$. Second, we show that the unique fixed point of the function $\mathbf{f}^\infty$ 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 $\mathbf{f}^\infty$. 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 $X_\theta^{\mathrm{MF}(\kappa)}(\mathbf{q})$, 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 $D_{\theta,n}(\mathbf{q}) = \max\left(\left\{\{\beta_\theta(V_{n,k})\}_{k=1}^{M_{\theta,n}^a(\mathbf{q})}\right\}_{\theta\in\mathbf{\Theta}}, r\right)$, 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 $\lim_{\kappa\to\infty} \tilde{N}_\theta^\kappa = \infty$ almost surely since $b_\theta^\kappa \leq C_{\theta,N_\theta^\kappa} \leq \tilde{N}_\theta^\kappa \bar{V}$, and $\lim_{\kappa\to\infty} b_\theta^\kappa = \infty$. Hence, taking the limit as $\kappa \to \infty$ and using the Strong Law of Large Numbers (SLLN), one obtains that

$(1/\tilde{N}^\kappa_\theta)\sum_{j=1}^{\tilde{N}^\kappa_\theta} Z_{\theta,j}$ converges to $\mathbb{E}Z_\theta$ a.s. In turn, we obtain that $\tilde{N}^\kappa_\theta/(\alpha^\kappa_\theta\eta^\kappa s_\theta) \to g_\theta/\mathbb{E}Z_\theta$ a.s.

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

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

The Dominated Convergence Theorem enables to conclude that

$$f^\infty_\theta = \lim_{\kappa\to\infty} f^\kappa_\theta = \lim_{\kappa\to\infty}\mathbb{E}\left[\min\{\tilde{S}^\kappa_\theta/s_\theta, 1\}\right] = \mathbb{E}\left[\min\{\lim_{\kappa\to\infty}\tilde{S}^\kappa_\theta/s_\theta, 1\}\right] = \min\left\{\frac{g_\theta}{\mathbb{E}Z_\theta}, 1\right\},$$

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

Step 2 (Fixed-points of $\mathbf{f}^\infty$). In this section we study the fixed-points of the limit function $\mathbf{f}^\infty$, 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 $f^\infty_\theta(\mathbf{q}) = \min\{g_\theta/\mathbb{E}[Z_\theta(\mathbf{q})], 1\}$ coincide with the solutions of the NCP

$$H_\theta(\mathbf{q}) \le g_\theta \quad \perp \quad 0 \le q_\theta \le 1, \quad \forall\theta\in\boldsymbol{\Theta}, \tag{C.4}$$

where the complementary condition is with the inequality $q_\theta \le 1$. 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 (C.4). Additionally, it is not hard to show that the functions

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

are differentiable. Also, by Lemma C.3 (stated and proved in Appendix C.5), the Jacobian of $\mathbf{H}$ is a P-matrix. Then, by Facchinei and Pang (2003*a*, Proposition 3.5.10) we conclude that 1 is the unique vector of active probabilities that solves (C.4), and thus the unique fixed-point of $\mathbf{f}^\infty$.

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 $\mathbf{f}^\infty$ is *uniform*, together with the continuity of the mapping, one would be able to invoke Lemma C.4 (stated and proved in Appendix C.5) to conclude that these fixed-points converge to the unique fixed-point of $\mathbf{f}^\infty$, i.e., $\lim_{\kappa\to\infty}\|1 - \mathbf{q}^\kappa\| = 0$. 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

$$
\begin{aligned}
|f_\theta^\kappa(\mathbf{q}) - f_\theta^{\kappa'}(\mathbf{q})| &= \left|\mathbb{E}\left[\min\{\tilde{S}_\theta^\kappa(\mathbf{q})/s_\theta, 1\}\right] - \mathbb{E}\left[\min\{\tilde{S}_\theta^{\kappa'}(\mathbf{q})/s_\theta, 1\}\right]\right| \\
&\leq \mathbb{E}\left|\min\{\tilde{S}_\theta^\kappa(\mathbf{q})/s_\theta, 1\} - \min\{\tilde{S}_\theta^{\kappa'}(\mathbf{q})/s_\theta, 1\}\right| \\
&\leq \mathbb{E}\left|\tilde{S}_\theta^\kappa(\mathbf{q})/s_\theta - \tilde{S}_\theta^{\kappa'}(\mathbf{q})/s_\theta\right| \\
&\leq \frac{1}{s_\theta}\mathbb{E}\left|\tilde{S}_\theta^\kappa(\mathbf{q}) - \mathbb{E}\left[\tilde{S}_\theta^\kappa(\mathbf{q})\right]\right| + \frac{1}{s_\theta}\mathbb{E}\left|\tilde{S}_\theta^{\kappa'}(\mathbf{q}) - \mathbb{E}\left[\tilde{S}_\theta^{\kappa'}(\mathbf{q})\right]\right| + \frac{1}{s_\theta}\left|\mathbb{E}\left[\tilde{S}_\theta^\kappa(\mathbf{q})\right] - \mathbb{E}\left[\tilde{S}_\theta^{\kappa'}(\mathbf{q})\right]\right| \\
&\leq \frac{1}{s_\theta}\sqrt{\mathrm{Var}[\tilde{S}_\theta^\kappa(\mathbf{q})]} + \frac{1}{s_\theta}\sqrt{\mathrm{Var}[\tilde{S}_\theta^{\kappa'}(\mathbf{q})]} + \frac{1}{s_\theta}\left|\mathbb{E}\left[\tilde{S}_\theta^\kappa(\mathbf{q})\right] - \mathbb{E}\left[\tilde{S}_\theta^{\kappa'}(\mathbf{q})\right]\right|, \qquad \text{(C.5)}
\end{aligned}
$$

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 Lyaponov's inequality. We next turn to the problem of bounding the mean and variance of the hitting time $\tilde{S}_\theta^\kappa$.

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 $\tilde{N}_\theta$. 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 $\mathrm{Var}(\tilde{N}_\theta) \leq (b_\theta + \bar{V})\mathrm{Var}(Z_\theta)/\mathbb{E}[Z_\theta]^3 + \bar{V}^2/\mathbb{E}[Z_\theta]^2$ (use Wald's second identity to get $\mathbb{E}[C_{\theta,\tilde{N}_\theta} - \tilde{N}_\theta \mathbb{E}Z_\theta]^2 = \mathrm{Var}(Z_\theta)\mathbb{E}\tilde{N}_\theta$). Next, recall that $\tilde{S}_\theta$ is a sum of a random number $\tilde{N}_\theta$ 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 $\mathrm{Var}[\tilde{S}_\theta] = (\alpha_\theta\eta)^{-2}(\mathbb{E}[\tilde{N}_\theta] + \mathrm{Var}[\tilde{N}_\theta])$ (see, e.g., Ross (1996, pp.22)).

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

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

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. Similarly, the variance can be upper bounded by

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

for some $K > 0$ independent of the scaling, the vector of active probabilities, and the type. The last follows from the facts that (i) $\mathrm{Var}[Z_\theta(\mathbf{q})] \leq \bar{V}^2/4$ 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 is (C.5) 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 $\kappa, \kappa' \to \infty$ uniformly in $\mathbf{q}$, we get that the sequence of functions is uniformly convergent.

Step 4 (Convergence in distribution of the maximum bid.) Let $\mathbf{q}^\kappa$ 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

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

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

C.2.4 Proof of Proposition 4.5

Fix a type θ and the scaling κ. Let $\mathbf{q}^\kappa$ be a consistent vector of active probabilities for the κ-th scaling (which exists according to Proposition 4.3). As in the proof of Proposition 4.4 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 $\beta_\theta^{\mathrm{F}}(x) = x/(1+\mu_\theta)$ 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

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

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 $\bar{J}_\theta^{\mathrm{F}}(F_d)$, 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 $\bar{\Psi}_\theta(\mu_\theta; F_d^\kappa) \to \bar{\Psi}_\theta(\mu_\theta; F_d)$ as $\kappa \to \infty$, since $F_d^\kappa \Rightarrow F_d$ from Proposition 4.4, and $\bar{\Psi}_\theta$ is continuous w.r.t. the distribution of the maximum bid from the proof of Theorem 3.1 of the main paper. Furthermore, since μ_θ is an optimal dual variable we get that $\bar{\Psi}_\theta(\mu_\theta; F_d) = \bar{J}_\theta^{\mathrm{F}}(F_d)$, in view of Proposition 3.1 of the main paper. Additionally, from Proposition 4.4 we have $\mathbb{E}[Z_\theta(\mathbf{q})]$ is continuous in $\mathbf{q}$, and thus $\mathbb{E}[Z_\theta(\mathbf{q}^\kappa)] \to \mathbb{E}[Z_\theta(1)]$ as $\kappa \to \infty$. 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 $\mathbb{E}(X-x)^+ \leq (\mathbb{E}X - x)^+ + \sqrt{\mathrm{Var}(X)/2}$, 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 4.4 to get

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

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

C.2.5 Proof of Proposition 4.6

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 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 $J_\theta^{\mathrm{MF}(\kappa)}(\beta^{\mathrm{H}}, \boldsymbol{\beta}^{\mathrm{F}})$ 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

$$
J_\theta^{\mathrm{MF}(\kappa)}(\beta, \boldsymbol{\beta}^{\mathrm{F}}) \le J_\theta^{\mathrm{MF}(\kappa)}(\beta^{\mathrm{H}}; \boldsymbol{\beta}^{\mathrm{F}}).
$$

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 advertisers decides to wins the auction or not. In hindsight, the zeroth advertiser needs to solve, for each realization ω, the following knapsack problem

$$J_\theta^{\mathrm{H}(\kappa)}(\omega) = \max_{x_{n,0}\in\{0,1\}} \sum_{n=1}^{n_\theta} x_{n,0}(v_{n,0} - d_{n,0}) \tag{C.6a}$$

$$\text{s.t.} \sum_{n=1}^{n_\theta} x_{n,0} d_{n,0} \le b_\theta. \tag{C.6b}$$

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 $J_\theta^{\mathrm{MF}(\kappa)}(\beta^{\mathrm{H}}, \boldsymbol{\beta}^{\mathrm{F}}) = \mathbb{E}_\omega\left[J_\theta^{\mathrm{H}(\kappa)}(\omega)\right]$.

Consider the continuous relaxation of the hindsight program (C.6) in which we replace the integrality constraints by $0 \le x_{n,0} \le 1$. Let μ_θ be the equilibrium multiplier of the FMFE for type θ. Introducing dual variables $\mu \ge 0$ for the budget constraint and $z_n \ge 0$ for the constraints $x_{n,0} \le 1$, we get by weak duality that

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

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

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

We conclude by noting that $\lim_{\kappa\to\infty} \bar{\Psi}_\theta(\mu_\theta; F_d^\kappa) = \bar{J}_\theta^{\mathrm{F}}(F_d)$ as in the proof of Proposition 4.5.

C.2.6 Proof of Proposition 4.7

Section C.4 shows that the AdX market may be modeled as a closed system with a random number of agents. Furthermore, Proposition C.2 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 C.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 $g'(\eta, F_k, \alpha, T)$, as defined in Corollary C.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 $\mathbb{E}_{Y^\kappa}\left[g^\kappa(Y^\kappa)\right]$,

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

with $A^\kappa(y) = 2\bar{\alpha}^\kappa + \sqrt{\mathrm{Var}Y^\kappa} + |y - \mathbb{E}[Y^\kappa]|$, $B(y) = y$, 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

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

The first term in parenthesis can be bounded as

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

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

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

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

for $y, \xi \geq 0$ to obtain

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

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 C.3.5, 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

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

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 C.1 for the extended states.

C.3 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.

C.3.1 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 $v : \mathbb{X} \times \mathbb{R} \to \mathbb{X}$, 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 $\vec{x} = \langle x_1, x_2, \ldots, x_{|\vec{x}|} \rangle$ 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 $\vec{x} \cdot \vec{y} = \langle x_1, \ldots, x_{|\vec{x}|}, y_1, \ldots, y_{|\vec{y}|} \rangle$. 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.[1] The indices of the participating agents is given by the set $\mathcal{M}_n = \{k = 1, \ldots, K : A_{k,n} = 1\}$, 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.

[1] As 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.

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* $f : \mathbb{X}^{\mathbb{N}} \times \mathbb{E}^{\mathbb{N}} \times \mathbb{R} \to \mathbb{X}^{\mathbb{N}}$, such that $f(\vec{x}, \vec{\xi}, t) = \vec{y}$ gives an additive change in the 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 $\{1, \ldots, |\vec{x}|\}$ 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)

$$\mathrm{d}X_k(t) = v(X_k(t), t)\mathrm{d}t + f_1\left(X_k(t) \cdot \vec{X}_{-k}(t), \xi_k(t) \cdot \vec{\xi}_{-k}(t), t\right) A_k(t)\mathrm{d}N(t), \tag{C.7}$$

for all agent $k = 1, \ldots, K$. 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 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.

C.3.2 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 $\mathbb{P}_c : \mathcal{B}(\mathbb{X}) \times \mathbb{R} \to [0,1]$, 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 differential equation

$$\mathrm{d}\tilde{X}_0(t) = v(\tilde{X}_0(t), t)\mathrm{d}t + f_1\left(\tilde{X}_0(t) \cdot \vec{\tilde{X}}_{-0}(t), \xi_0(t) \cdot \vec{\tilde{\xi}}_{-0}(t), t\right) A_0(t)\mathrm{d}N(t). \tag{C.8}$$

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 $\vec{\tilde{X}}_{-0,n} = \left\langle \tilde{X}_{1,n}, \ldots, \tilde{X}_{\tilde{M}_n,n} \right\rangle$, 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 $\vec{\tilde{\xi}}_{-0,n} = \left\langle \xi_{1,n}, \ldots, \xi_{\tilde{M}_n,n} \right\rangle$. 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 3. *A distribution $\mathbb{P}_c : \mathcal{B}(\mathbb{X}) \times \mathbb{R} \to [0,1]$ is* T-consistent *if for any Borel-measurable set of states $\mathcal{X}$ and time t*

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

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)).

C.3.3 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 $\vec{k} \in \mathbb{Z}_+^{\mathbb{N}}$, denoted as $\vec{k} = \langle k_1, k_2, \ldots \rangle$.

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 area 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 $\vec{k} = \langle 0, \ldots, 0, k_{n+1}, k_{n+2} \ldots, k_{|\vec{k}|} \rangle$ 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 $\langle k_{n+2}, \ldots, k_{|\vec{k}|} \rangle$ 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 $\langle 0, 0, \ldots, 0 \rangle$, $\langle 0, \ldots, 0 \rangle$, 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 Boltzman tree is T-consistent for the mean-field model (Definition 3). This is proved formally in, for example, Chauvin and Giroux (1990).

C.3.4 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 C.1. *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 $A = \alpha + \alpha\sqrt{VarK} + \alpha|K - \mathbb{E}[K]|$, $B = \alpha(K-1)$, 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 $\{\widetilde{K}_n\}_{n\in\mathbb{N}}$ 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 $\widetilde{K}_n$ 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 $\widetilde{M}_{n,k} = \sum_{i=1}^{\widetilde{K}_n+1} m_{n,i} - 1$, 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 $\widetilde{M}_{n,k}$ 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 $\Delta(\Gamma_k(t)) = \left| M_{\mathcal{N}_k(t),k} - \widetilde{M}_{\mathcal{N}_k(t),k} \right|$ 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 a 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 C.3.3. 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

$$\begin{aligned} U(t;K) \leq \int_0^t & \underbrace{\mathbb{P}\left\{\Delta(\Gamma_k(x)) \neq 0 \mid K, k \in \mathcal{M}(x)\right\}}_{(I)} \\ & + \underbrace{\mathbb{P}\left\{\bigcup_{k' \in \mathcal{M}(x)} \Gamma_{k'}(x^-) \text{ not uncorrelated} \mid K, k \in \mathcal{M}(x)\right\}}_{(II)} \mathrm{d}\bar{F}_{\alpha\eta}(t-x) \end{aligned} \tag{C.9}$$

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

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

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 (C.9), 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 (C.9), one obtains the functional equation $U(t;K) \leq AF_{\alpha\eta}(t) + (1+B)\Big(U(\cdot;K) * F_{\alpha\eta}\Big)(t)$; where we denoted the Stieltjes convolution by $(F * G)(x) = \int F(x-u)\,\mathrm{d}G(u)$. Iterating the functional equation we obtain the following exponential bound on the probability that the tree is not uncorrelated

$$
\begin{aligned}
U(t;K) &\le A\sum_{i=0}^{\infty}(1+B)^i F_{\alpha\eta}^{(i+1)}(t) = A\sum_{i=0}^{\infty}(1+B)^i \int_0^t \frac{(\alpha\eta)^{i+1}x^i}{i!}e^{-\alpha\eta x}\,\mathrm{d}x \\
&= A\int_0^t \alpha\eta e^{-\alpha\eta x}\sum_{i=0}^{\infty}\frac{(\alpha\eta x(1+B))^i}{i!}\,\mathrm{d}x = A\int_0^t \alpha\eta e^{\alpha\eta x B}\,\mathrm{d}x \\
&= A\frac{e^{\alpha\eta t B}-1}{B} \le A\frac{e^{2\alpha\eta t B}-1}{2B}
\end{aligned} \tag{C.10}
$$

where $F_{\alpha\eta}^{(i)}(t) = \int_0^t \frac{(\alpha\eta)^i x^{i-1}}{(i-1)!}e^{-\alpha\eta x}\,\mathrm{d}x$ 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 $(e^x-1)/x \le (e^{2x}-1)/(2x)$ for $x \ge 0$.

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 $Q(t;K)=\sum_{d\ge 0}Q_d(t;K)$.

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,

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

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

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

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

$$\begin{aligned}
Q(t;K) = \sum_{d=0}^{\infty} Q_d(t;K) &\le \sum_{d=0}^{\infty} B^d\Big(Q_0(\cdot;K) * F_{2\alpha\eta}^{(d)}\Big)(t) \\
&\le Q_0(t;K) + B\sum_{d=0}^{\infty}(1+B)^d\Big(Q_0(\cdot;K) * F_{2\alpha\eta}^{(d+1)}\Big)(t) \\
&= Q_0(t;K) + B\int_0^t Q_0(t-x;K)2\alpha\eta e^{2\alpha\eta x B}\,\mathrm{d}x = \frac{\alpha}{2B}\left(e^{2\alpha\eta t B} - 1\right), \qquad \text{(C.11)}
\end{aligned}$$

where the second inequality follows from partitioning the sum, the non-negativity of the terms and using that $B \le B+1$; 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,

$$
\begin{aligned}
L(t;K) &\le \int_0^t (CQ(x;K) + (1+B)L(x;K))\ \mathrm{d}\bar{F}_{\alpha\eta}(t-x) \\
&= C\Big(Q(\cdot;K) * F_{\alpha\eta}\Big)(t) + (1+B)\Big(L(\cdot;K) * F_{\alpha\eta}\Big)(t) \\
&\le C\sum_{i=0}^{\infty}(1+B)^i\Big(Q(\cdot;K) * F_{\alpha\eta}^{(i+1)}\Big)(t) \\
&= C\int_0^t Q(t-x;K)\alpha\eta e^{\alpha\eta x B}\ \mathrm{d}x \\
&\le \frac{\alpha C}{2}\frac{e^{2\alpha\eta t B} + 1 - 2e^{\alpha\eta t B}}{B^2} \le \frac{\alpha^2\eta t C}{2}\frac{e^{2\alpha\eta t B} - 1}{2B},
\end{aligned}
\tag{C.12}
$$

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 (C.11), integrating and discarding negative terms; and the last inequality from the fact that $(e^{2x} - e^x - 1)/x^2 \le (e^{2x} - 1)/(2x)$ for $x \ge 0$.

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

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

□

C.3.5 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 $H'_0(t) = \left\{X'_0(t_n^-), \vec{X}'_{-0}(t_n^-), \xi_0(t_n), \vec{\xi}_{-0}(t_n)\right\}_{t_n \le t}$. 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

$$\mathrm{d}\tilde{X}_0'(t) = v(\tilde{X}_0'(t), t)\mathrm{d}t + f_1'\left(\tilde{H}_0'(t)\right) A_0(t)\mathrm{d}N(t), \tag{C.13}$$

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 C.1. *We have that*

$$\left\|\mathcal{L}(X_0') - \mathbb{P}_c'\right\|_{[0,T]} \leq g'(\eta, F_k, \alpha, T)$$

where

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

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

Proof. The proof follows as in Proposition C.1, 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 C.1. We conclude by taking expectation w.r.t. the number of agents in the system, which is now equal to $K + 1$. □

C.3.6 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 $\alpha : \mathbb{X} \times \mathbb{R} \to [0, 1]$, 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

$$\tilde{M}_n = \sum_{k=1}^{\tilde{K}_n} \tilde{A}_{k,n}, \tag{C.14}$$

where $\tilde{A}_{k,n}$ is Bernoulli with success probability $\tilde{\alpha}(t_n)$, and $\tilde{\alpha}(t) = \int_{\mathbb{X}} \alpha(x,t)\ \mathrm{d}\mathbb{P}_c(x,t)$ 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

$$\tilde{\mathbb{P}}_c(\mathcal{X}, t) \triangleq \frac{1}{\tilde{\alpha}(t)} \int_{\mathcal{X}} \alpha(x,t)\ \mathrm{d}\mathbb{P}_c(x,t). \tag{C.15}$$

Note that in this case Definition 3 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 (C.15).

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 $\vec{x}$ the states and $\vec{u}$ the uniform noise terms of the agents that pass the first round, we denote the set of agents that pass the second round by $\bar{\mathcal{M}}(\vec{x}, \vec{u}, t) = \{i = 1, \ldots, |\vec{x}| : u_i \leq \alpha(x_i, t)/\bar{\alpha}\}$. The previous construction guarantees that agents interact with their correct state and time-dependent probability. Then, the new interaction function is defined as $\bar{f}_i\left(\vec{x}, (\vec{\xi}, \vec{u}), t\right) = 0$ for $i \notin \bar{\mathcal{M}}(\vec{x}, \vec{u}, t)$ and

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

where $\vec{x}|_{\mathcal{I}} = \langle x_i \rangle_{i \in \mathcal{I}}$ is a slice of $\vec{x}$ 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.4 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 $X_k^{\mathrm{C}}(t) \in \mathbb{R}^3 \times \mathbf{\Theta}$; where given a state $x_k^{\mathrm{C}}(t) = (b_k(t), s_k(t), \tau_k, \theta_k)$, 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 $\tau_k \le t \le \tau_k + s_{\theta_k}$. 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 $\alpha\big((b, s, \tau, \theta), t\big) = \alpha_\theta \mathbf{1}\{\tau \le t \le \tau + s\}$, that is, advertisers match with their type-dependent probability only during their campaign.
- The deterministic drift is given by $v\big((b, s, \tau, \theta), t\big) = -e_s \mathbf{1}\{\tau \le t \le \tau + s\}$, 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 unchange 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 4.2 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

$$\hat{p}_\theta \int_0^{s_\theta} \mathbb{P}_e^{\text{BMFM}}\Big(\mathcal{B}, s_\theta - u\Big|\theta, u\Big)\mathbf{1}\{s_\theta - u \in \mathcal{S}\}\frac{1}{s_\theta}\,\mathrm{d}u = \mathbb{P}_e(\mathcal{B}, \mathcal{S}, \theta). \tag{C.16}$$

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 4. *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 $\mathbb{P}_c^{\text{BMFM}}(B \in \mathcal{B}, S \in \mathcal{S}, \mathrm{T} = \tau, \Theta = \theta, t)$, is given as follows. First, the probability that an advertiser is of type $\theta \in \mathbf{\Theta}$ 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,

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

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

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

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

When we specify the initial conditions X_0 as drawn from $\mathbb{P}_c^{\mathrm{BMFM}}$ 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 form the consistent distribution of the BMFM. We have the following result.

Proposition C.2. *Let* $\mathbb{P}_e^{\mathrm{BMFM}}$ *be a time dependent consistent distribution for the BMFM, and* $\mathbb{P}_c^{\mathrm{BMFM}}$ *the induced distribution for the closed market, as given by Definition 4. Then,* $\mathbb{P}_c^{\mathrm{BMFM}}$ *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 $\mathbb{P}_c^{\mathrm{BMFM}}$, which is determined from a consistent distribution of the BMFM through Definition 4. 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

$$\begin{aligned}\tilde{\alpha}(t) &= \int_{\mathbb{X}} \alpha(x,t)\ \mathrm{d}\mathbb{P}_c^{\mathrm{BMFM}}(x,t) = \int_{\mathbb{X}} \alpha_{\theta_x} \mathbf{1}\{\tau_x \in [t - s_{\theta_x}, t]\}\ \mathrm{d}\mathbb{P}_c^{\mathrm{BMFM}}(x,t) \\ &= \int_{\mathbb{X}} \alpha_{\theta_x} \frac{s_{\theta_x}}{T + s_{\theta_x}}\ \mathrm{d}\mathbb{P}_c^{\mathrm{BMFM}}(x,t) = \sum_\theta p_\theta \alpha_\theta \frac{s_\theta}{\mathbb{E}[s_\Theta] + T} = \frac{\mathbb{E}[\alpha_\Theta s_\Theta]}{\mathbb{E}[s_\Theta] + T},\end{aligned} \tag{C.17}$$

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 (C.14) 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 (C.15) we obtain that an agent's state conditional on interacting at time t is distributed as:

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

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 4, 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 (C.17); 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 (C.16). 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. $\square$

C.5 Auxiliary Results

Lemma C.3. *Fix a vector of multipliers $\boldsymbol{\mu}$ and consider the differentiable vector function $\mathbf{H} : [0,1]^{|\boldsymbol{\Theta}|} \to \mathbb{R}_+^{|\boldsymbol{\Theta}|}$ 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 $\partial F_d/\partial q_\theta = -F_d(x;\boldsymbol{q})\mathbb{E}[\alpha_\Theta \lambda s_\Theta]\hat{p}_\theta \bar{F}_{v_\theta}((1+\mu_\theta)x)$, 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

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

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 B.3 of the main paper and canceling terms we can write

$$\frac{\partial H_\theta}{\partial q_\theta} = \sum_{\theta' \neq \theta} \hat{p}_{\theta'}(1+\mu_{\theta'})q_{\theta'}\langle f_{\theta'}, \bar{F}_\theta\rangle + \hat{p}_\theta q_\theta \langle \bar{F}_\theta, \bar{F}_\theta\rangle + r\bar{F}_{v_\theta}((1+\mu_\theta)r)F_d(x;\boldsymbol{q}). \tag{C.18}$$

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

$$\begin{aligned}\frac{\partial H_\theta}{\partial q'_\theta} &= q_\theta \frac{\partial}{\partial q'_\theta}\mathbb{E}[Z_\theta(\mathbf{q})] = -q_\theta \int_r^{\bar{V}} \frac{\partial}{\partial x}\left(x\bar{F}_{v_\theta}((1+\mu_\theta)x)\right)\frac{\partial F_d}{\partial q'_\theta}\,\mathrm{d}x \\ &= -\hat{p}_{\theta'}(1+\mu_\theta)q_\theta\langle f_\theta, \bar{F}_{\theta'}\rangle + \hat{p}_{\theta'}q_\theta\langle \bar{F}_\theta, \bar{F}_{\theta'}\rangle. \end{aligned} \tag{C.19}$$

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 (C.18) 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

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

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. □

Lemma C.4. *Consider a sequence of continuous mappings* $\{\mathbf{g}^n\}_{n\geq 1}$ *with* $\mathbf{g}^n : [0,1]^d \to [0,1]^d$ *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* $\lim_{n\to} \mathbb{D}_\infty(X^n, X) = 0$[2].

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

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 $\mathrm{dist}(\mathbf{x}^n, X) \geq \epsilon$ for some $\epsilon > 0$ and that $\mathbf{x}^n$ converges to a point $\mathbf{x}^* \in [0,1]^d$. It follows that $\mathbf{x}^* \notin X$. 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. $\square$