ABSTRACT

Title of Dissertation:	ESSAYS ON SPONSORED SEARCH AUCTIONS AND ONLINE PLATFORMS
	Panagiotis Dimitrellos Doctor of Philosophy, 2022
Dissertation Directed by:	Professor Andrew T. Sweeting Department of Economics

This dissertation covers two topics within the context of online platform design. In chapter 1, I develop a structural model to evaluate the effect of providing user traffic information to heterogeneous bidders in sponsored search auctions. Many internet platforms use sponsored search auctions as their primary source of revenue. In these auctions, advertisers bid for slots with different desirabilities. A standard assumption is that bidders know the click through rate (CTR) for each slot. I relax this assumption in two ways. First, I allow bidders to receive private signals about the CTR of the highest advertising slot, which they use to update their beliefs during an auction. Second, I allow for bidders to start each auction with different priors. I estimate my model using closed-form formulae and a new dataset from a large internet platform where the CTR for the highest slot varies significantly across auctions and over time. My estimates imply considerable variation in bidders' priors, and that this affects platform revenues. Specifically, I predict that the platform's revenues would increase by an average of 7% if the platform was able to credibly and accurately reveal the CTR of the highest slot. I show how this gain in revenues

relates to changes in revenue from bidders who, in the absence of knowledge of the CTR, have either optimistic or pessimistic priors about the CTR.

In chapter 2, I show how to calculate the theoretically optimal reserve prices in auctions for online advertisements with endogenous platform user behavior. In the case where the advertised content is useful to the platform's users, showing less advertisements due to increased reserve prices could imply less clicks on each advertisement from users because of a smaller choice set. Qualifying more bidders by lowering reserve prices creates a positive externality for all participating bidders. I present the results of a large field experiment in a sponsored search setting. Consistent with the theory, platform revenues increase substantially after the introduction of the optimal reserve prices, while users engage more with the website.

In chapter 3, I discuss the economic benefits of a Central Dispatcher for the New York City taxi industry. Drivers make dynamic spatial decisions without taking into account that their decisions impact their fellow drivers and consumer demand, increase traffic and affect matching efficiency. The Central Dispatcher internalizes that driver decisions affect the outcomes of other drivers and have an effect on congestion, demand, and matching efficiency. The Central Dispatcher makes decisions in an environment with search frictions, while taking into account the aforementioned externalities in order to maximize the market's social surplus. I solve for the Central Dispatcher allocation using a value function approximation algorithm based on neural networks. Results indicate that the competitive equilibrium leads to imperfect coordination between the drivers, excess supply and more traffic congestion than the optimum. The Central Dispatcher increases social surplus by 15%, or $798,000 per shift, reduces congestion by 5% on average and increases market thickness in Lower Manhattan and the Boroughs.

ESSAYS ON SPONSORED SEARCH AUCTIONS AND ONLINE PLATFORMS

by

Panagiotis Dimitrellos

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2022

Advisory Committee:
Professor Andrew T. Sweeting, Chair/Advisor
Professor Ginger Zhe Jin
Professor Liye Ma
Professor Daniel R. Vincent
Professor Chenyu Yang

Acknowledgments

I am profoundly grateful to my advisor, Professor Andrew Sweeting, for his constant guidance and support. His generous help and advice are the main reasons I consider myself an Economist today. I was very fortunate to have an advisor who is always there despite how packed his schedule may be. I am thankful to my dissertation committee for all their feedback and work. Professor Daniel Vincent provided invaluable comments on the theoretical aspects of this dissertation and his amazing teaching attracted me to Industrial Organization. Professor Chenyu Yang was always able to understand and give solutions to every research related problem of mine in a matter of seconds. Professor Ginger Jin heard my very first research ideas and was there to help me understand the empirical aspects of all my vague proposals during my first years in the program.

I would never forget to thank Peter, Tim, Dan, and Artashes from Tripadvisor. They gave me an opportunity to do research at a great company, and taught me how to think like a data scientist, and how to communicate my work effectively.

None of this would have been possible without the support of my family and friends; my parents, Giannis and Dimitra, who showed me how to strive and persevere. My brother, Makis, who taught me how to share. My friends, Nikos and Christos, who hold the fort back home.

Finally, I would like to thank Erika for making all this journey beside me, and for being there every day, rainy or bright.

Table of Contents

List of Tables

List of Figures

Chapter 1: Information, Bias, and Revenues in Sponsored Search Auctions

1.1 Introduction

Many internet platforms and networks generate a significant part of their revenue through the sale of advertising space. Most online platforms organize their space for advertisements in a list form, with different ads competing for user attention. Users engage with ads in the top of the list more often than with ads in lower slots. Therefore, it is potentially more valuable for an advertiser to place her ad in a high slot because it will receive more clicks from platform users. In platforms such as Google, Tripadvisor and Yahoo the advertisement slots are allocated with the help of an auction mechanism (the Sponsored Search Auction (SSA)), where advertisers become bidders and submit bids that reflect their valuation of advertisement.[1] A common form of these SSAs is that bidders submit a single bid which is then used to determine which bidders are allocated to which slots and the prices paid when an advertisement in a particular slot receives a click.

The literature studying SSAs has assumed that the click through rate (CTR) of each slot and the probability that a click will convert to a sale is known to all bidders.[2] However, this assumption may not be realistic. For example, on Tripadvisor, the platform that I study in

[1] The total valuation of an online advertisement slot is a function of the user attention it receives (clicks) and the unit value of attention (per click value).

[2] The click through rate describes the number of clicks that each slot offers.

this paper, the CTR and conversion rate for the top advertising slot on a hotel listing page can vary quite dramatically over time as a result of how many viewers are coming to the page by clicking on Tripadvisor's own paid search advertising, as these viewers tend to have quite different click and conversion patterns to Tripadvisor's own members who arrive at the page by performing searches on Tripadvisor. Anecdotal evidence suggests that different bidders, online travel agencies in my setting, may differ in their ability to predict the CTR for the top slot (slot 1) and may believe, systematically, that its average value for a given hotel is higher or lower than it really is.[3]

This naturally leads to the question of whether the platform could increase its revenues by providing bidders with accurate information about the CTR, and, if so, by how much. I investigate this question by estimating a structural model of bidding behavior in Tripadvisor SSAs. I depart from conventional assumptions in two ways. First, I allow bidders to have asymmetric information about the CTR of the top slot in the auction. For tractability, I assume that they know the CTR of other slots, which is plausible as, if they click at all, paid searchers are more likely to click on the top advertisement. The auction operates as a Generalized Second Price Auction (GSPA), which I model as a Generalized English auction.[4] Before the auction takes place, each bidder receives a signal about the top slot's CTR, which she updates, in Bayesian fashion, based on what happens as the auction plays out.

Second, I allow bidders to have different prior beliefs about the CTR to capture their biases, if any. A prior belief distribution centered close to the true number of clicks of the top slot reflects a bidder with the ability to predict the CTR well. Respectively, a prior belief distribution centered

[3]Not knowing the value of the auctioned items leads to lower bids due to the Winner's Curse. Even if bidders are oblivious to their biases, the winner of the auction is the bidder who overestimates valuation the most.

[4]See Edelman, Ostrovsky, and Schwarz (2007)

further away from the true number of clicks of the top slot shows a bidder that fails to predict the platform's traffic more often. I assume that bidders are oblivious about theirs and their opponent's biases because a rational agent would not be willingly biased.

By making appropriate distributional assumptions on prior beliefs and signals I am able to develop closed form estimators for the structural parameters that form bidder priors. I also utilize data on bids for hotels where there is never any paid search to control for differences in how bidders value listing slots which may reflect bidders' value of brand marketing in the city from which my data comes. I use my estimates to perform a counterfactual where the platform provides credible and accurate information about the click through rate. This counterfactual returns the model to the conventional SSA assumptions.[5] While the elimination of the winner's curse will tend to be good for revenues it is possible that, by eliminating the possibility that bidders have over-optimistic priors, the counterfactual could lower expected platform revenues. However, for the auctions that I study, I predict that the policy of revealing the CTR to bidders would raise platform revenues by an average of 7%.

Specifically, I, first, derive the equilibrium conditions for a generalized English auction where bidders are allowed to have asymmetric information and different priors about CTR.[6] I design a game theoretical model of the generalized English auction which allows for information asymmetry across bidders. The asymmetry is established by making the number of clicks in the highest slot unknown to bidders. Then, I solve for the Bayesian equilibrium of the game and derive the equilibrium strategies. I use the equilibrium conditions and available data on bids, click through rates, and bidder valuations to derive closed-form and consistent estimators for the

[5]In this context, a richer and more complicated bidding language could allow bidders to bid differently for different click qualities. Given that there is paid and organic traffic, this should be equivalent with revealing the ratio of paid to organic clicks.

[6]If a bidder's prior differs significantly from the true ratio, I refer to this bidder as "more biased".

mean and the variance of each bidder's prior belief for the CTR of each hotel. The possibility of constructing a reliable proxy of per click valuations using existing data allows me to identify bidder beliefs using observed bids.

The assumption that bidders return to their previous priors once the CTR changes, allows for some bidders to be consistently optimistic / pessimistic about the CTR, therefore bidding non-optimally. I also distinguish for reasons that lead to overbidding other than low quality information. I assume that there are two different reasons for advertisers to systematically bid higher: optimism, as discussed above, and marketing. Marketing refers to increased spending by the bidder to be more visible to consumers and reinforce their brand.

Second, I use the model's estimates to provide empirical evidence that bidders are biased in their beliefs about the CTR. I use the equilibrium conditions and available data on bids, click through rates, and bidder valuations in order to estimate the parameters connected to the information environment. The quality of a bidder's information is defined as the distance of their prior belief to the true ratio of clicks between slots the highest and second highest slots.[7] If a bidder's prior is not close to this ratio, then the bidder is expected to bid suboptimally, especially in the early stages of the game, when information depends heavily on the prior. The estimation results show that the variance in the bidders' quality of information is significant, suggesting that bidders can be biased.

Finally, I show that the seller's revenue would increase in the case where the platform decides to reveal CTR information and, therefore, erasing bidder bias. I perform a counterfactual where the platform reveals information about the click through rate, i.e. number of clicks in

[7]Assuming that clicks in the second highest slot are known to bidders, this is equivalent with a belief about the clicks in the highest slot.

the highest slot. This removes all information asymmetry and bias and the game transits to the unbiased (standard) case, introduced by Edelman, Ostrovsky and Schwarz (EOS) (2007). Results suggest that when bidders are biased, both optimistic and pessimistic bidders end up underbidding on average, however in different volumes. This underbidding allows less biased bidders to extract rents, facing less competition from both sides of the bias spectrum. More biased bidders generate higher revenue for the platform in the symmetric information case compared to the asymmetric information case with bias, while less biased bidders generate lower revenue for the platform when information becomes symmetric since they are not able to extract rents anymore.[8] The net effect of the policy is positive for the platform.

To my knowledge, this is the first study on asymmetric information and biased beliefs about click through rates in the SSA. It provides equilibrium conditions that can be used in a variety of applications for different platforms. Given the relevant data, the framework allows any platform that uses the GSPA to evaluate whether the transition to a known-CTR environment would be profitable.

This paper has two main contributions to the literature. First, it offers a game theoretic model that serves as a road map to estimate bidder beliefs given the adequate data. Most structural papers that aim to solve problems related to auctions use the equilibrium conditions to identify bidder valuations, while assuming the structure of bidder information as given and exogenous. However, in a realistic setting, online platforms have reliable data on the value of each click that they provide.[9] By constructing an accurate proxy for per click valuations, I allow for the identification of bidder beliefs from the equilibrium conditions and observed bids. Second, the

[8]Due to the environment's volatility, bidders cannot learn anything useful about the CTR over time. However, access to more human resources and better algorithms allow some bidders to have a better guess at any given moment.

[9]The value of the click is tightly related to the price of the advertised item and the probability of conversion.

paper shows that the claim that more information leads to higher seller revenue is not universally true. Despite that removing bidder bias is beneficial in Tripadvisor's case, the result is highly dependent on the bidder-bias composition. A platform where bidders have different level of biases could be worse off from this policy.[10]

This research is part of a broader study on the design of SSAs accounting for realistic features of the Tripadvisor environment. This study includes the current paper, and Dimitrellos (2020), (2021). The main result of Dimitrellos (2020) suggests that GSPA generates greater revenue for the platform compared to the generalized hybrid price auction, where the payment price is determined to be the maximum of the next highest bid and a fraction of bidder's own bid. Dimitrellos (2021) discusses the optimal reserve price of the auction given the results of the first two parts of the study. Adding to Edelman and Schwarz's (2010) environment, it allows CTR to change when the number of winning bidders changes, following the intuition that when users see less offers tend to engage less with the platform. In particular, the paper presents empirical evidence that platform users engage less when the ad supply is limited. This creates a trade-off between increasing revenue by increasing the reserve price, versus higher demand by allowing more bidders to appear in the search results. The optimal reserve prices are derived computationally and their performance is evaluated in a field experiment as in Ostrovsky and Schwarz (2016).

The rest of the paper is organized as follows: Section 2 discusses the existing literature and open questions on SSAs. Section 3 displays the auction framework used by Tripadvisor. In section 4, I describe the available data and show some aggregated statistics. Section 5 discusses

[10]For example, a platform with mostly less biased bidders might not want to reveal CTR because the existence of more biased bidders leads their less biased counterparts to overbid.

the information environment in the platform, the model, and its equilibrium. Section 6 summarizes identification, and the estimation strategy of the model's parameters. In section 7, I present the results of the estimation. Finally, section 8 discusses the counterfactual policy along with its effects, and section 9 concludes.

1.2 Literature

Sponsored Search Auctions is a subcategory of multi-unit auctions that has sparked substantial interest. Its multi-unit nature comes with several open questions, while its structure offers increased tractability compared to the general multi-unit auction. One aspect of the SSAs that has received considerable attention is optimal design. Questions of this type include the optimal pricing rule, optimal reserve prices, and the plausibility of each of the auction's equilibria. Many of the questions related to optimal auction design have been answered but mostly under restricting and often unrealistic assumptions. Another set of questions about SSAs that arise are related to the evaluation of theoretical results in a real setting. The majority of online platforms allocate their space to advertisers using a form of a SSA. Therefore, the impact of theoretical results on a real setting can significantly affect the platforms' revenues. However, the auction environment of an online platform is, at best, noisy. For this reason, it is important to model the environment as realistically as possible. Not many results exist in this direction because the data needed to capture additional frictions are difficult to find. This paper attempts to identify information frictions using a rich data set originated from a large online platform.

More specifically, this paper lies in the intersection of two research fields. First, it relates to the literature of optimal auction design. A major result in this field came in the 1980s with

Myerson (1981). Myerson solved the problem faced by a seller who has a single object to sell to one of several possible buyers, when the seller has imperfect information about how much the buyers might be willing to pay for the object. Edelman and Schwarz (2010) extend the single object results of Myerson (1981) to cover multiple items under complete information about the click through rates (CTR). More recently, Ulku (2013) considers an optimal mechanism design problem with several heterogeneous objects and interdependent values. The author characterizes ex post incentives using an appropriate monotonicity condition and reformulates the problem in such a way that the choice of an allocation rule can be separated from the choice of the payment rule. Ulku extends Myerson's results along two directions, the heterogeneity of the items for sale and the interdependent values which introduce the Winner's Curse.

Many results after Myerson refer to special cases of auctions. Iyengar and Kumar (2006) derive the optimal design for a dynamic game of incomplete information used to sell sponsored search advertisements. They also consider a corresponding static game of complete information. They analyze the underlying dynamic game of incomplete information, and they establish an upper bound on the revenue of any equilibrium of any dynamic game in this environment. Ostrovsky and Schwarz (2016) confirm the previous results by organizing a field experiment in a SSA platform.

Finally, optimization results have been derived while considering the dynamic nature of the auctions run by internet platforms. The dynamic side of the SSAs is based on the conjecture that the results that are showed to the users can affect their future behavior. Parkes and Singh (2004) and Said (2012) focus on cases where the intertemporal trade-offs arise through the dynamics of arrival, departure, or population. Rafieian (2020) provides structural evidence of interdependence in the sequence of auction outcomes and user engagement in the platform. Rafieian's approach is different from this paper's since it introduces a dynamic element in user behavior. Rafieian shows that there are significant revenue gains from using the dynamic revenue optimal auction compared with the static optimization since improved user engagement compensates for suboptimal revenue in the short term.

My paper adds to this literature by introducing an additional dimension to the problem, which is the symmetry of bidder information. In a realistic setting with fluctuating user demand, bidder knowledge of the CTR is at best limited as anecdotal evidence suggests. I provide empirical evidence that bidders' prior beliefs are different, and I examine whether removing the biases is optimal for seller revenue.

Second, the paper relates to the literature of game theory and Bayesian Nash equilibria in auctions with incomplete information. One of the early results in Bayesian equilibria for the SSA comes with Lahaie (2006). Lahaie analyzes the incentive, efficiency, and revenue properties of two slot auction designs: "rank by bid" (RBB) and "rank by revenue" (RBR). With incomplete information, neither RBB nor RBR are truthful, i.e. not incentive compatible to report type, with either first or second pricing. Lahaie finds that the informational requirements of RBB are much weaker than those of RBR, but that RBR is efficient whereas RBB is not. Finally, Lahaie shows that with complete information, no equilibrium exists with first pricing using either RBB or RBR.

Expanding the empirical literature on SSAs, Edelman and Ostrovsky (2006) provide empirical evidence on the inefficiency of the generalized first price auction, since it leads to strategic behavior, i.e. "Sawtooth" pattern, where bidders with low valuations may end up winning the highest slots. An important result comes from EOS (2007). They investigate the generalized second price auction (GSPA) and show that unlike the VCG mechanism, GSPA generally does not have an equilibrium in dominant strategies, and truth-telling is not an equilibrium of GSPA.

To analyze the properties of GSPA, they describe the generalized English auction that corresponds to GSPA and show that it has a unique equilibrium.

One of the few attempts to use a more realistic set of assumptions is Yan (2019), who introduces interdependent values into the auction. Yan finds that both the GSPA auction and the VCG-like auction with one-dimensional bidding language can be inefficient under interdependent values, where each bidder's per click value depends on its opponents' information as well as its own information. Furthermore, Yan shows this inefficiency problem can be fully resolved by adopting a multi-dimensional bidding language that allows bidders to bid differently across positions.[11]

My theoretical framework differs from Yan's work because the I study bidder uncertainty about the click through rate, not the value of each click separately. I offer a Bayesian perfect equilibrium in the case of asymmetric information and biased beliefs about the CTR and calculate the seller revenue increase when CTR becomes common knowledge. In the case for Tripadvisor, the bidders have perfect visibility on the conversion rates of organic clicks, and they set their own prices. However, bidders face difficulties in predicting the number of incoming clicks. These observations suggest that the value of a click should be known to the bidders, since it depends solely on the conversion rate, and the premium that the bidder receives from the hotelier when a conversion happens. Therefore, the assumption of interdependent values would not be plausible in this case.

[11]The adoption of this modification in a real setting can be problematic due to the complexity that bidders will face.

1.3 Auction Overview

This paper focuses on the case of Tripadvisor. Tripadvisor is the world's largest online travel company that operates a website and a mobile app with a comparison shopping website. It offers online hotel reservations and other travel services. Its website reached 463 million average monthly unique visitors in 2019. In 2019, Tripadvisor earned 33% of its revenues from Expedia Group and Booking Holdings and their subsidiaries, primarily for pay-per-click advertising. Online travel agents compete for advertisement slots in the platform which are allocated through the generalized second price auction. Tripadvisor holds repeatedly simultaneous generalized second price auctions for different user searches. A user search is defined as combination of search variables such as hotel id (location id), length of stay, days to arrival and number of guests. First, the user sees a list of hotels that match her search, as displayed in figure 2.1. The order of appearance for hotels after step 1 is not related to the auction but is personalized for each different user. OTAs only compete for user attention within hotels. When a user clicks on a hotel, a limited number of travel agents (OTA) is displayed in a list form, as displayed in figure 2.2.

Different positions have different desirability for OTAs. When a user clicks on an OTA's listing, she is redirected to the OTA's website, as displayed in figure 2.3. Then, the OTA pays the platform for sending the user to its website - "pay-per-click" rule. All positions are auctioned at once. However, bidders can submit only one bid. The highest bid wins the highest slot and so on, until the auction is repeated.

The bidders who bid more than a preset reserve price, r , are ranked in terms of decreasing bids. Rarely, this is not the final order of appearance since further adjustments take place to the rules. When the order is determined, a bidder pays when it receives a click, regardless of whether

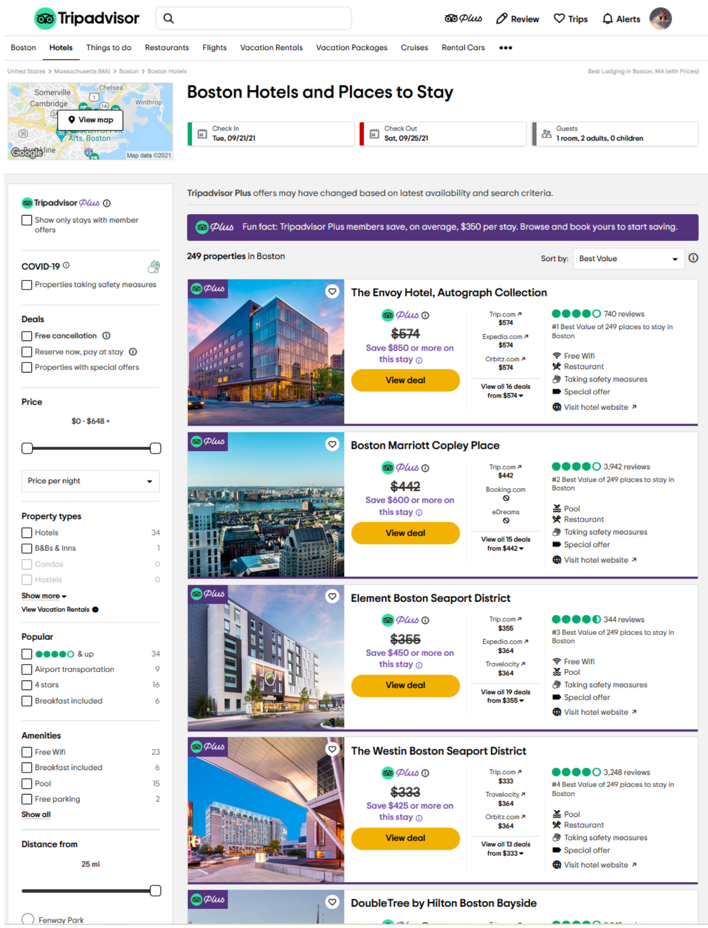

Figure 1.1: Hotel list

the click leads to a booking.

The pricing scheme used by Tripadvisor in 2021 is the second price payment (next highest bid). Bidders are allowed to update their bids twice a day. Each bidder is notified about the position she achieved in case of winning, their cost per click, and if she bid less than the reserve price. Theoretically, it is possible to infer the position of other winning bidders since each auction's ranking is available to platform users.

I discuss a 2-bidder example with private values and perfect information about the click through rate. Note that the rest of the paper does not assume perfect information about the click through rate, but this example aims to display the auction rules in a simple way. Suppose that

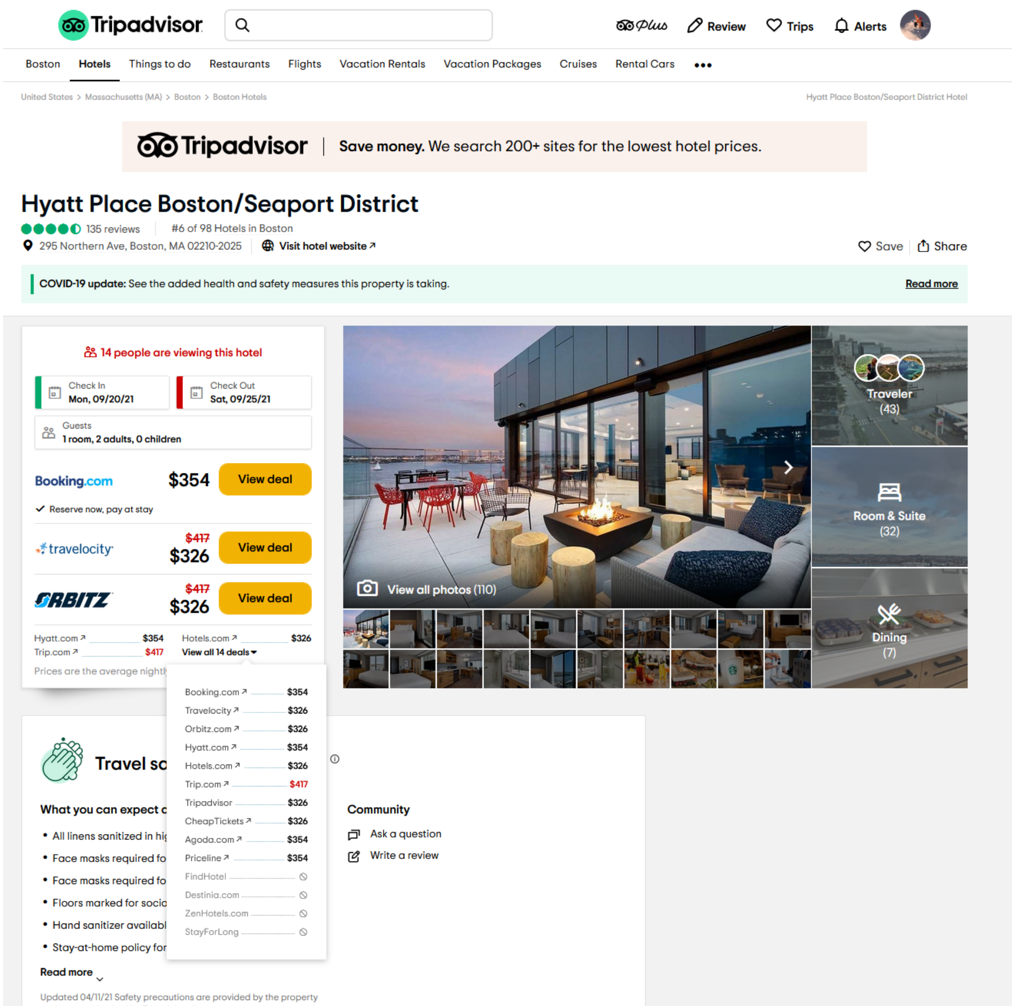

Figure 1.2: OTA display

hotel H has an available room for a given date D. Assume that this hotel is not a partner of Tripadvisor, therefore its only option to advertise the room in Tripadvisor is to assign it to travel agents affiliated with Tripadvisor. Suppose that 2 travel agents are available, B_1 and B_2. The hotelier announces that she will require e.g. \$100 from the agent that will sell the room. The two travel agents compete for getting slot 1 in Tripadvisor's page about this room. The platform runs the auction for this room and this date and asks from B_1 and B_2 to submit their bids.

The click through rate is 20 clicks for slot 1 and 5 clicks slot 2. In addition, suppose that both bidders offer the room for \$150, therefore the bidder who sells the room will have a profit

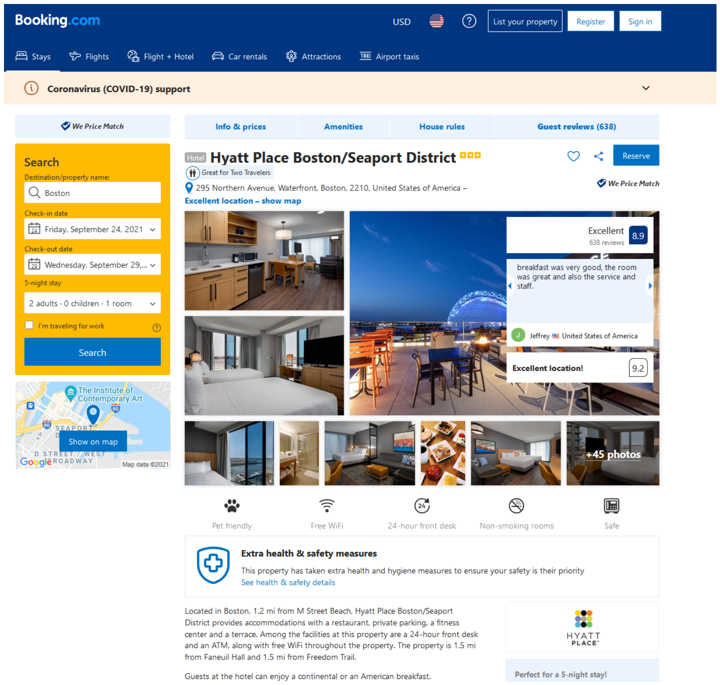

Figure 1.3: Redirection

of \$50 minus the cost of advertisement, that goes to Tripadvisor.

In this example, B_1 submits a bid equal to 1.9 and B_2 submits a bid equal to 1.5, while the reserve price is equal to 1. Both clicks are above the reserve price and therefore qualify for a slot. B_1 submitted the highest bid and gets slot 1. B_2 gets slot 2. Users of the platform searching for hotels for date D, happen to encounter hotel H which offers the room. When users click to hotel H, they access a list of 2 offers, containing B_1 in slot 1 and B_2 in slot 2. As the click through rate expresses, 20 users click on B_1's offer and 5 users click on B_2's offer. All realizations of the parameters related to CTR, and conversion are not necessarily equal to their expected values according to bidders. Bidder beliefs on CTR are irrelevant in this example. What matters is that their bids reflect their equilibrium strategies given their beliefs and auction rules.

In this example only one user that clicked on B_1's offer decides to book the room, while no user that clicked on B_2's offer decides convert. Therefore, the revenue of B_1 is \$150 and the revenue of B_2 is zero. However, B_1 must pay \$100 to the hotelier. Furthermore, both bidders must pay the advertising cost to the platform. Given that the platform runs a generalized second price auction, each bidder pays the maximum of the next highest bid and the reserve price when clicked. Thus, the advertising cost for B_1 would be $20{\cdot}1.5$ while B_2 would pay $5{\cdot}1$. Summarizing:

$\pi_1^{GSPA} = 150 - 100 - 20 \cdot 1.5 = \$20$

$\pi_2^{GSPA} = -5 \cdot 1 = -\$5$

In addition, the CTR is not known to the bidders since the fraction of valuable to non valuable clicks in slot 1 changes over time. This makes the number of high value clicks in slot 1 for a given auction unknown to bidders. In the case of Tripadvisor, the reason for low value clicks is traffic coming through ads in third party websites (paid traffic). These users are less engaged with the platform and they are more often exploring than trying to book a specific room. The vast majority of these users click in slot 1 for convenience reasons, thus I assume that the only slot with variance in click value is slot 1. All these reasons make Tripadvisor an appropriate environment for studying the effect of asymmetric information and biased beliefs about the CTR in a generalized second price auction environment.

1.4 Data

The data sets used in this study are provided by Tripadvisor and include the entire data collection of the platform since 2015, therefore it covers all auctions for all realized searches. The data set covers both sides of the market; it contains all user activity, including clicks and conversions. Although I have visibility on user activity along with certain demographics, I do not connect each activity with a unique user id because of confidentiality reasons and the difficulty to track user activity on a platform that does not requires membership to use. The respective location, bidder, auction position and conversion outcome are available for each click, therefore I can infer the click through rates for each location across time and the conversion frequency. The data about the supply side describes bidder activity. First, it displays in which category each bidder belongs across locations and time.[12] Furthermore, it contains the submitted bids for all realized auctions and the auction outcomes. Finally, the data set contains partial information about the bidder margins for each location, i.e hotel, when a conversion is realized. Given that conversion frequencies and bidder margins are known, the data allows the calculation of the value of a user click to a bidder. Denote the unknown per click valuation of a bidder in a specific auction with v. I calculate a proxy for v: $\tilde{v}$ as follows:

$$\tilde{v} = margin \cdot p \cdot \mathbb{P}(conversion) \cdot (1 - \mathbb{P}(cancellation))$$

where

- $margin$: The percentage of the transaction value that goes to the bidder from the hotelier.
- p: The booking price that the user sees on the platform.[13]
- $\mathbb{P}(conversion)$: The probability that a user books, given that the user has clicked on the bidder's ad. Bidders share conversion data with the platform, hence the calculation of the rate is feasible since it does not change significantly across time.

[12] There are two main categories of bidders, which differ in how much they derive their bids with the platform's assistance or by themselves.

[13] The price consists of the per night price times the number of nights.

- $\mathbb{P}(cancellation)$: The probability that a user cancels her reservation, given a booking.[14]

The main data set includes the outcome of every user search along with the information that the user was provided, and information provided to the platform by OTAs, e.g. their conversions. The provided information is extensive; therefore I present the most important variables. The data set is sizable, e.g. a subset of the data set contains more than 5 million searches in the US for the 200 most popular hotels for April 2021. More specifically, each observation contains search characteristics (days to arrival, length of stay, day of the week, if date is specified, hotels sorted by number of guests, day, and hour of search), hotel characteristics (name, location, stars, reviews, amenities, hotel parent brand, clicks in last month), user characteristics (device-browser, country, member, targeted, booked before), offer characteristics for every bidder in this search (bidder id, bidder category, bid, auction position, price, is refundable, pay at stay, all inclusive), and outcome characteristics (offer clicked, offer converted).

The data set is sparse with respect to conversions, since the overall conversion rate is low. The second data set contains cancellation data. When a user is converted the partners provide Tripadvisor with the outcome of the booking. The two possible outcomes are that the stay happened as planned or that it has been canceled by the user. The third data set contains the commissions that partners receive for each conversion. The source of revenue for the partners is a commission paid by the hoteliers when a conversion happens as a percentage of the total amount received by the hotel, since OTAs do not ”own” rooms. More specifically, each observation contains the bidder ID, the hotel ID, and the commission rate that the OTA extracts from the hotel as a percentage rate.

I use auction data from the downtown area of a big US city, customers from the US and

[14] This means that the user (partially) receives her payment back and the bidder is not paid by the hotelier.

desktop device. The subset of the data set of hotels in the platform that receive no paid traffic contains auctions that took place in April 2021. This dataset serves for the identification of spending for marketing reasons. Bidders have knowledge of the no paid traffic property for these hotels since the origin of a click is ex post visible to the bidder. The data set of hotels in the platform that receive no paid traffic contains 130,000 auctions in 10 popular hotels in this city.

The paper assumes that all OTAs would face the same CTR and conversion rates in counterfactual auctions where results are different, given a search.[15] The data on user clicks allow the reconstruction of the click through rate for every auction across time. A small subset of hotels in the platform receives no paid traffic for exogenous reasons. The paid traffic volume is not determined by the auction designer but by other segments of the platform on which the auction designer has no control. The existence of hotels with zero paid traffic implies that in these cases there is no variance in the click values for slot 1, hence there is no asymmetry in information. This subset of the data is useful for the estimation of over(under)bidding because of the information asymmetry. There is an additional reason that can lead to over(under)bidding that the researcher should control: Bidders can focus on specific markets for marketing reasons, e.g. increasing their market share. If a bidder wants to build their brand name in a geographical area, then they are expected to bid more than the optimal value and win slot 1 more often than they should in order to be seen more by the platform's users. When there is no paid traffic, any deviation from the theoretical optimal bid can be attributed to marketing, which allows its identification.

The main question that the paper aims to answer in whether it is in the platform benefit to change its information policy regarding the value of each click.[16] While the clicks in slots 2 and

[15]For example, in a given auction slot 1 receives 20 clicks regardless of which bidder wins. In addition, a certain bidder converts a certain percentage of clicks in an auction, e.g. 3%, regardless of her slot.

[16]The value of a click depends on the probability of conversion. A click from a disengaged user who is looking around is less valuable than a click from a user that has already decide to spend her vacation in the place that the

lower are of the same kind, organic clicks, the clicks generated by slot 1 is a mix of organic and paid traffic clicks. The difference between organic and paid traffic clicks lies in their probability of conversion to a booking. Paid traffic is less interested in booking a listing, because the user is rarely a platform member and visits the platform through an ad in a third website. A screenshot of a google search that can lead someone to Tripadvisor's website through an ad can be found in figure 1.4. A user reaching Tripadvisor through a third-party ad is expected to be less engaged with the platform than a signed up member that searches directly at Tripadvisor, despite the fact that they see the same screen. The reason for the this is that an organic user that visits Tripadvisor directly is more likely to be familiar with the platform and book a room than a person looking around search engines.

Therefore, when the number of clicks that come from slot 1 in an auction is 5 times higher than the number of clicks that come from slot 2, the total value generated for the bidder is slot 1 is less than five times higher than the total value generated for the bidder is slot 2. This would not create a complication if the ratio of paid traffic clicks to organic clicks in an auction was predictable for the bidders. However, the way that paid traffic clicks are generated in each auction is not related to the auction mechanics, auction rules, and varies across time. Anecdotal evidence from bidders suggests that their ability in predicting the ratio of paid traffic clicks to organic clicks in an auction is limited at best.

Figure 1.5 displays the share of paid traffic clicks received by slot 1 in a hotel in a range of 30 days.

Paid traffic intensity appears to change over time in a volatile and unpredictable way. The

hotel is located. The reason for this is that the former user pays to book a room fairly rarely, while the probability of conversion is much higher for the latter user.

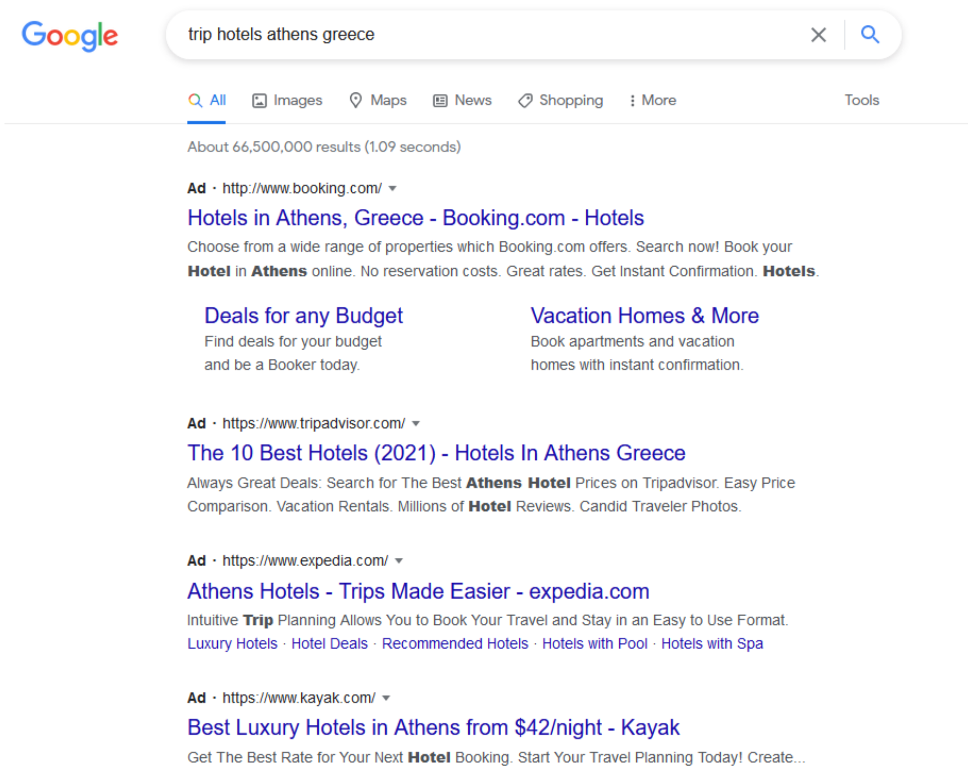

Figure 1.4: Screenshot of a Tripadvisor ad in a third party website

respective paid traffic plots for other hotel show similar levels of variance across time. Table A.1 shows an attempt to predict the share of paid traffic in slot 1 for different hotels. I use panel data containing the paid traffic intensity for 100 hotels in the studied geographical area over 30 days. The independent variables of the dataset are previous observations for paid traffic (1,2 and 3 auctions lag), the total number of clicks in slot 1 for the auction, the ratio of clicks between slots 1 and 2 for the auction, the average submitted bid, each individual bid submitted by bidders, and hotel and time fixed effects. The prediction power of the data seems discouraging at best, showing an $R^2 = 0.03$. The fact that each individual bidder has access to only a small subset of this data supports the claim that predicting paid traffic intensity is a difficult problem for bidders.[17]

[17] A bidder has access to past values of paid traffic if and only if the won slot 1 for the corresponding auction. In addition, the ratio of clicks between slots 1 and 2 is unobservable for all bidders, because a bidder can win mostly one slot. Finally, a bidder cannot observe the bids of her competitors since the platform does not reveal them.

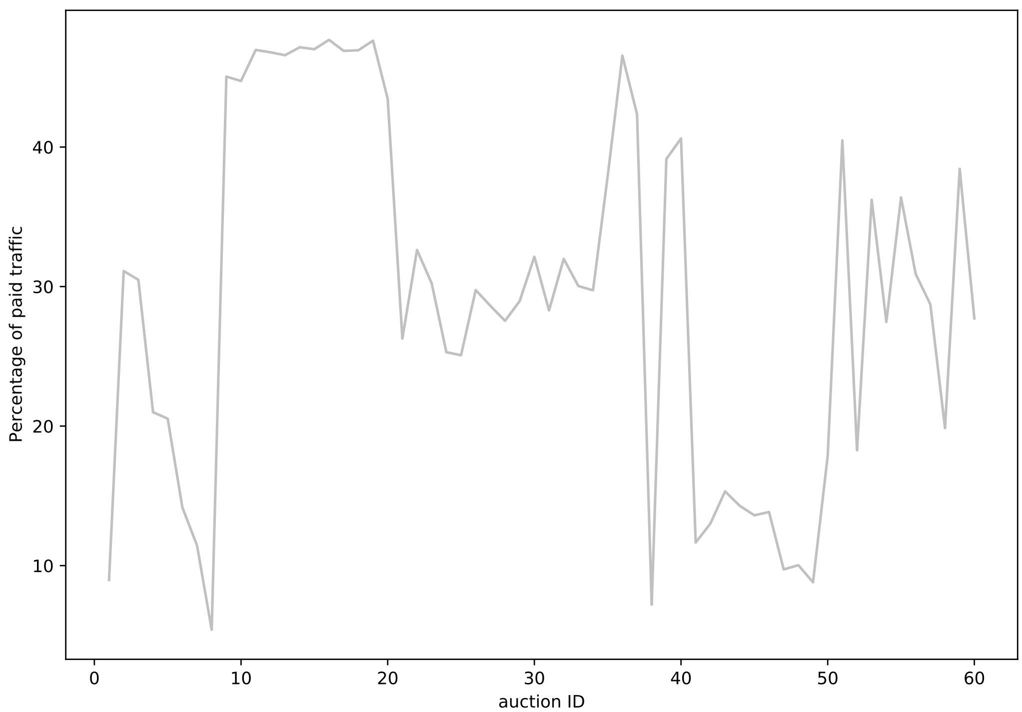

Figure 1.5: Paid traffic intensity within hotel across time

1.5 Model

1.5.1 Environment

The peculiarity of the platform that introduces uncertainty about the slot value is the exogenous generation of paid traffic whose intensity changes over time. Paid traffic focuses on slot 1 and converts rarely, reducing its per-click value. However, bidders pay full price for it. There are two ways to model unpredictable paid traffic. First, one could represent it as uncertainty about the per-click value in slot 1, in the sense that the mix of organic and paid traffic averages to a per-click value in between of the organic and paid click valuations. Another way to model paid traffic would be to represent it as uncertainty in the number of ”equivalent” clicks offered by slot 1, assuming that all clicks have the organic per-click valuation. The two approaches are equivalent and differ only in their calibration. Assume that in the original setting there are n organic clicks of value $\bar{v}$ each and m paid clicks of value $\underline{v}$. The bidder pays p for each click. Therefore, bidder return is $ret = (n \cdot \bar{v} + m \cdot \underline{v}) - (n + m) \cdot p$. The first approach attempts to find a value in between that produces the same revenue for the bidder: $ret_1 = (n + m) \cdot y - (n + m) \cdot p$. Setting $ret = ret_1$ and solving for y: $y = \frac{n \cdot \bar{v} + m \cdot \underline{v}}{n+m}$, which depends on the unpredictable volume of paid traffic, m. The second approach aims to find a number of clicks, x, that return the same revenue and are of value $\bar{v}$. Hence, $ret_2 = x \cdot \bar{v} - x \cdot p$. Setting $ret = ret_2$, and solving for x: $x = n + m\frac{\underline{v}-p}{\bar{v}-p}$. In addition, the ”equivalent” clicks approach can capture the effects of bidder bias in platforms where the source of uncertainty is different. All that matters is that the click volume cannot be predicted. For these reasons, I model the uncertainty about the slot valuation coming from CTR.

The fact that some bidders have access to better software infrastructure and information than others creates the hazard of overpaying when winning the auction, since the total value offered by slot 1 is unknown. In addition, the facts that the ratios change significantly across time and that the bidders' feedback mechanism is restricted, creates room for consistently optimistic or pessimistic bidders.[18] The latter implies that a bidder may be optimistic in an auction and turn to pessimistic when the auction is repeated in the future with different ratios. However, the problem for the bidder reduces to guessing the ratio given the available information, while past ratios do not matter.

An optimistic bidder can be perceived as a bidder who often fails to obtain information that suggests a lower ratio. Changes in information policy can alleviate this complication and affect platform and bidder revenue. The first counterfactual that I evaluate is the case where the platform reveals the ratios to the bidders. This is feasible since the platform decides the amount of paid traffic for each auction. However, it would be impossible for the platform to exactly predict the amount of organic traffic ex ante. Despite the uncertainty, data suggests that organic traffic is highly seasonal and can be predicted by the platform accurately. I also consider the policy of allowing bidders to bid differently for different kinds of clicks, i.e. different bids for organic and paid traffic clicks.

I approach the question by developing a theoretical framework for bidder behavior in this environment. I derive a closed form expression of the Bayes Nash Equilibrium under certain assumptions: Bidders are modeled to have different priors on the ratio of clicks and update their priors based on the bidding behavior of their opponents. The ability of bidders to update

[18] By optimistic (pessimistic) bidder, I define a bidder who consistently over(under)estimates the total value offered by slot 1 in an auction.

come from the modeling of the auction as a generalized English auction. This is a common technique in the literature, e.g. EOS (2007), which captures the procedure that bidders learn about their opponents as time progresses. Different auctions across time for the same hotel are treated as independent and bidders maintain the same priors. Inference of opponent information is introduced in the model by the English auction environment. In Tripadvisor's case the auction rankings are posted in online and are visible by everybody, while bidders can often change their bids in order to measure any changes. Hence, the use of the generalized English auction seems justified. I solve in an identical private value environment, which allows for a closed form solution, and is plausible at the same time, since bidder click valuations do not vary significantly across bidders.[19]

1.5.1.1 Assumptions

There are two key assumptions in this model. These assumptions form the setup of the model and set the guideline of approaching the research question. The key assumptions embedded in the model follow.

Assumption 1. The variation in the click through rate of different slots is only coming from variation in paid traffic. Given an auction, the clicks for slots 2 and lower are assumed to be the same across time, while clicks for slot 1 change according to paid traffic shock.

The ratios of organic to paid traffic are assumed to be determined exogenously. The data provide visibility on paid traffic within hotels and confirms its volatility over time. However, there are exceptions. Certain hotels receive no paid traffic over time, making the ratio of organic to total clicks equal to 1. This group of hotels represents less than 5% of the population and

[19] I discuss the identical private value assumption later in the paper and provide evidence in favor of it.

is assumed to be determined exogenously. I later show how this group of hotels helps with the identification of overbidding due to marketing reasons.

Assumption 2. There is no heterogeneity in per click valuations within auctions, i.e., bidders have identical private values.

This assumption is based on the fact that per click valuations of bidders do not differ significantly for a certain auction. The estimation method provides evidence that variance in information and marketing spending describes the submitted bids sufficiently well. Evidence in favor of the identical private value assumption can be found in the Appendix, where I perform a statistical test of identical private values using data on bidder margins, booking prices, conversion, and cancellation rates.

Assumption 3. Bidder spending on marketing happens on the geographical level. A bidder exhibit the same marketing spending in every hotel in the dataset since all are located in the same geographical area.

This assumption originates from the intuitive claim that a bidder would organize a marketing campaign on the country or the city level rather than the hotel level. The complexity of running different marketing campaigns for every hotel is significant given the number of hotels worldwide. In addition, the volume of click data that a bidder receives from each hotel is not enough to evaluate the campaign's performance.

1.5.2 Modeling framework - Examples

In this game, I examine the generalized English auction, an analogue of the standard English auction corresponding to GSPA. The notion that described the SSA in the previous sections was the generalized English auction, where advertisers have converged to a long-run steady state, have learned each other's values, and the equilibrium is stable, as in EOS (2007). However, in this application the bidders are assumed to be unable to recognize the biases in their beliefs about the click through rate. At the same time, bidders know that their opponents have private information about the click through rate and use their bids to infer this information. More specifically, given possibly biased beliefs, bidders try to calculate the true click through rate given information from opponent bids in a rational way.

The generalized English Auction, a mechanism introduced by EOS (2007), helps to simulate the aforementioned environment. In the generalized English auction, there is a clock showing the current price, which continuously increases over time. Initially, the price on the clock is zero, and all advertisers are in the auction. An advertiser can drop out at any time, and his bid is the price on the clock at the time when he drops out. The auction is over when the next-to-last advertiser drops out. The ad of the last remaining advertiser is placed in the best position on the screen, and this advertiser's payment per click is equal to the price at which the next-to-last advertiser dropped out. The ad of the next-to-last advertiser is placed second, and his payment per click is equal to the third-highest advertiser's bid, and so on. In other words, the vector of bids obtained in the generalized English auction is used to allocate the objects and compute the prices according to the rules of GSP. With one object, the generalized English auction becomes a simple English auction. The strategy of an advertiser assigns the choice of dropping out or not for any history of the game, given that the advertiser has not previously dropped out. Therefore:

- The strategy of each bidder is a single number (the drop-out price) for every history of the game. A history of the game at time t includes the signal that the bidder gets and the prices that other bidders dropped out before time t.

- A bidder can win at most one slot. If this bidder is the i^{th} bidder who dropped out, she will get the i^{th} lowest slot.

In reality, such a clock does not exist, and bidders submit their bids simultaneously. However, the constant repetition of the auction across time can allow bidders to learn from their opponents' behavior, making the environment very similar to the generalized English auction. In the following subsections I present two simple versions of the game along with their equilibrium conditions. These simplified versions serve as examples to display how bidders derive their bids and infer information about the click through rate by their opponents' bids. The game follows the assumption of identical private per click valuation. In addition, bidders are assumed to have the same prior for simplicity reasons. This assumption is removed in later sections.

1.5.2.1 Example - Two bidders

- 2 bidders
- Same valuation per click, $v_1 = v_2 = v$
- Two slots, with click rates c_1, c_2.
- c_2 is known to both bidders, while c_1 is not known.
- Each bidder i gets a signal s_i from a Weibull distribution, $W[1, c_1]$.
- The signals that bidders get are conditionally independent, given c_1.
- Reserve price: r

Both bidders' prior distribution is an Inverse Gamma $(2, \alpha \cdot c_2)$, with mean $\alpha \cdot c_2$. At this stage, α is common for both bidders. However, the analysis stays the same in the case where it is not, the algebraic part is just a bit longer. A large α, implies that the bidder is optimistic regarding the number of valuable clicks in slot a, while when α is closer to 1, it implies a more pessimistic bidder. Inverse Gamma is a proper prior for the Weibull distribution, given that its shape is greater than 1. Inverse Gamma was selected because of its conjugate prior property when paired with Weibull. The timing of the game is as follows:

1. Each bidder receives an independent private signal from $W[1, c_1]$.
2. All bidders update their beliefs on c_1 based on their private signal.
3. A clock showing the current price, which continuously increases over time, starts at the reserve price r.
4. As the current price continuously increases a bidder drops out - denote with p_2. This bidder gets the last slot (slot 2) and pays r for each click in this slot.
5. The clicks in slot 1 are realized and the remaining bidder gets slot 1 and pays p_2 for each click in slot 1.

I assume that the equilibrium strategy profile is symmetric, and the equilibrium strategy is strictly increasing in s_i.

Proposition 1 *The dropout price in the two bidder Generalized English Auction with asymmetric information is:*

$$p_i^* = v - \frac{4c_2}{\alpha \cdot c_2 + 2s_i} \cdot (v - r)$$

Proof: See Appendix.

It is visible in the result that the bidder with the highest signal wins slot 1, a property of pure common value games. In addition, the drop out price is inversely correlated with the clicks of slot 2 because the more clicks slot 2 offers, the more attractive it becomes. Finally, it is straightforward that the more optimistic the bidders are, the higher the drop out price is; the latter is reflected by α. In the appendix, I discuss an example with 3 bidders, and the implications of bidder bias for the seller's revenue in the case of 2 bidders.

1.5.3 Model with N bidders

I generalize the drop out conditions over n bidders and I describe the Bayesian Nash equilibrium of the game. A clock shows a price which continuously increases over time, and it starts at the reserve price r at the beginning of the game. As in the previous sections, each bidder receives a private signal on the ratio of clicks in slots 1 and 2. As the price increases, bidders drop out one by one and slots are allocated. When a bidder drops out at a certain price, the remaining bidders infer her signal and update their beliefs about the ratio. I split the game in $n-1$ stages: $0, 1, ..., n-2$, each stage representing the time between two consecutive bidder i drop outs. At any stage of the game, t, a participating bidder has the following information set:

1. Their private signal, s_i.

2. The history of the game $h^t = \{b_1, b_2, ..., b_t\}$, which contains all the prices that previous bidders dropped out.

3. The current price on the clock, p.

Each bidder infers the signals of the bidders that have already quit and drops out when the clock reaches p_i^t, unless another bidder quits before p_i^t, and the game progresses to stage $t+1$. Expanding the result of the previous section:[20]

$$p_i^t = v - \frac{(n+2)\cdot c_t}{\alpha\cdot c_2+(n-t)s_i+\sum_{j=1}^{t} s_j(b_j)} \cdot (v - b_t)$$

As shown before, when the game starts $(t = 0)$, bidder i would quit if the price reaches p_i^0 and no other bidder has quit until then, where:

$$p_i^0 = v - \frac{(n+2)\cdot c_0}{\alpha\cdot c_2+n\cdot s_i} \cdot (v - r)$$

In summary, the equilibrium strategy profile of each bidder in the game is described by the following matrix of dimension n x $(n - 2)$:

$$\begin{cases} p_1^t(h^t), & \text{where } t \in \{0,1,...,n-2\}, h^t \in H^t \\ p_2^t(h^t), & \text{where } t \in \{0,1,...,n-2\}, h^t \in H^t \\ ... \\ p_n^t(h^t), & \text{where } t \in \{0,1,...,n-2\}, h^t \in H^t \end{cases} \tag{1.1}$$

where H^t is the space of histories for the game at stage t:

[20]Abusing notation, I denote the clicks in slot t from the last slot as c_t.

$$
H^t = \begin{cases} \varnothing, & \text{if } t = 0 \\ \{x_1\}, & \text{where } x_1 > r, \text{ if } t = 1 \\ \{x_1, x_2\}, & \text{where } x_2 > x_1 > r, \text{ if } t = 2 \\ \ldots & \\ \{x_1, x_2, ..., x_{n-2}\}, & \text{where } x_{n-2} > ... > x_2 > x_1 > r, \text{ if } t = n-2 \end{cases} \tag{1.2}
$$

1.5.3.1 Bidder Bias

One last step is required to allow for optimism and pessimism for bidders. The difference between optimism and getting a high signal is its persistence across auctions because of the distribution that generates them, given that signals across auctions are assumed to be independent. Data suggest that some bidders consistently overbid, while others bid consistently lower than the bids suggested by EOS (2007). This cannot be explained by signal draws since signals are random while the deviating behavior is consistent towards one direction. An optimistic bidder typically believes that the ratio of organic to paid traffic is higher than its real value. I model this by allowing bidders to have different priors in the dimension of expected value. The mean of an optimistic bidder's prior is higher than the real ratio, while the mean of a pessimist bidder's prior is lower.

More specifically, the prior distribution for bidder i is an Inverse Gamma $(2, \alpha_i \cdot c_2)$, with mean $\alpha_i \cdot c_2$. I do not allow bidders to know the priors of their opponents, since rational bidders would not have different priors. I also do not allow bidders to know that their opponents have different priors, since the signal identification becomes impossible. This is because when a bidder sees an opponent dropping out of the auction, it is not possible to identify how this bid was affected by the prior of their opponent and how it was affected by the signal of their opponent. In the appendix, figure A.1 depicts bids over time and justifies the form of bias in the model.

The game in this case unfolds as follows:

1. Each bidder gets a ”prior parameter” a_i from nature that applies to every auction in a hotel. Their prior parameter reflects the mean of their prior for every auction in this hotel.

2. When the game starts each bidder gets a private signal from $W[1, c_1]$.

3. When a bidder quits, the bidder updates her prior, while ignoring that step 1 happens (naivety).

I assume that the marketing effect is incorporated in the bidder’s prior. I model the marketing effect to be a part of the prior because a bidder’s brand is reinforced when it is being seen more by the platform’s user, i.e. being in slot 1 more often. Increasing a bidder’s prior makes slot 1 more attractive to the bidder and subsequently increases their equilibrium bid. More specifically, consider H hotels in the same geographical area. It is plausible to assume that bidder marketing happens at the geographical area level, not the auction level. For example, a bidder might decide to spend additional resources for auctions of hotels in Long Beach, CA for US users for marketing reasons. I consider implausible that bidder marketing campaigns can be more specific than the latter example because of the large number of auctions and the increased level of complexity. Hence, for hotel h in the geographical area l for bidder i, the prior parameter of a bidder is formulated as follows.

$$\alpha_{i,h} = ratio_h \cdot opt_{i,h} \cdot mar_i^l \tag{1.3}$$

There is a prior parameter, $\alpha_{i,h}$ for every hotel-bidder combination, the real ratio of clicks between slot 1 and 2, $ratioh$, is one per hotel, same for all participating bidders. The "optimism" effect, $opt_{i,h}$, reflects each bidder's prior belief on the real ratio of clicks between slot 1 and 2, one for every hotel-bidder combination. Finally, the marketing effect, mar_i^l, describes the additional effect of spending of bidder i in geographical location l for marketing reasons. Note that increased (decreased) bidder spending on marketing is modeled to apply to an increase (decrease) of the prior parameter, α. This implies that marketing effect can translate to a behavior where the bidder acts as the ratio of clicks between slot 1 and slot 2 is differs from $ratio_h$. While in practice, bidder may multiply all their bids with a constant greater than one when they want to market themselves in the platform, I assume that bidders choose a number mar_i^l, that is multiplied with the mean of their prior. More specifically, I assume that when bidders change their bids for marketing reasons, they do not multiply their bid presented in equation 1.4 by a constant, but they incorporate it into the prior term α_i. In addition, assumption 3 justifies why the marketing effect variable mar_i^l is indexed solely by the geographical area l bidder id n.

1.5.3.2 Equilibrium

The main difference of this game with the game in the previous section is that bidders are allowed to have different priors among themselves. However, to maintain the rationality of bidders, bidders ignore that nature gives them different prior parameters and assume that every bidder has the same prior parameter with them.[21] This leads to wrong inference of opponents' signals, since they assume that their beliefs are the same. In addition, a higher drop out price does

[21] As mentioned before, the prior incorporates both optimism and marketing effects.

not always imply a higher signal. A bidder with a lower signal can drop out later than a bidder with a higher signal if the prior parameter of the later is high enough. At any stage of the game, t, a participating bidder has the following information set:

1. Their private signal, s_i.

2. Their prior parameter, α_i.

3. The history of the game $h^t = \{b_1, b_2, ..., b_t\}$, which contains all the prices that previous bidders dropped out.

4. The current price on the clock, p.

Each bidder infers the signals of the bidders that have already quit and drops out when the clock reaches p_i^t, unless another bidder quits before p_i^t, and the game progresses to stage $t+1$. Expanding the result of the previous section:

$$p_i^t = v - \frac{(n+2) \cdot c_t}{\alpha_i \cdot c_2 + (n-t)s_i + \sum_{j=1}^{t} s_j^i(b_j)} \cdot (v - b_t) \tag{1.4}$$

As shown before, when the game starts $(t = 0)$, bidder i would quit if the price reaches p_i^t and no other bidder has quit until then, where:

$$p_i^0 = v - \frac{(n+2) \cdot c_0}{\alpha_i \cdot c_2 + n \cdot s_i} \cdot (v - r) \tag{1.5}$$

Equation (1.5) suggests that the bidder who drops out first is based solely on her prior. However, equation (1.4) suggests that as the auction progresses and bidders drop out, the remaining bidders' decisions are affected increasingly by their inferences of opponents' signals. Bidder's bias affects inference heavily. The inferred signal is different than the original signal in case bidders have different prior parameters. For instance, if a bidder has a high prior parameter, she infers the signal of a bidder with low prior parameter as their prior parameters are the same. Therefore, if a bidder drops out because of a low prior parameter, another bidder with a high parameter will interpret it as a very low signal. Hence, as the auction progresses, even initially optimistic bidders can bid lower than optimal given that bidders have different priors.

In summary, the equilibrium strategy profile of each bidder in the game is described by the following matrix of dimension n x $(n-2)$:

$$\begin{cases} p_1^t(h^t), & \text{where } t \in \{0,1,...,n-2\}, h^t \in H^t \\ p_2^t(h^t), & \text{where } t \in \{0,1,...,n-2\}, h^t \in H^t \\ ... \\ p_n^t(h^t), & \text{where } t \in \{0,1,...,n-2\}, h^t \in H^t \end{cases} \tag{1.6}$$

where H^t is the space of histories for the game at stage t:

$$H^t = \begin{cases} \varnothing, & \text{if } t = 0 \\ \{x_1\}, & \text{where } x_1 > r, \text{ if } t = 1 \\ \{x_1, x_2\}, & \text{where } x_2 > x_1 > r, \text{ if } t = 2 \\ \dots \\ \{x_1, x_2, ..., x_{n-2}\}, & \text{where } x_{n-2} > ... > x_2 > x_1 > r, \text{ if } t = n-2 \end{cases} \tag{1.7}$$

Note that the inference of opponents' signals is a different function for every bidder i: $s_j^i(b_j)$.

In this game, N bidders compete for slots in M auctions in a property. Each property (hotel) is associated with multiple auctions, since differences in variables such as number of guests, length of stay and days to arrival create different auctions that allow for different bids and different winners. Before the auctions begin, bidders are assigned prior parameters, α, by nature that reflect their priors. Vector α is of dimension $(N \text{ x } 1)$. A bidder with prior parameter α_n has a prior distribution on the ratio of clicks for slots 1 and 2 equal to α_n for all auctions held in this property. When each of the M auctions start, each bidder receives a private and conditionally independent signal s_{nm} , $n \in \{1, 2, ..., N\}$, $m \in \{1, 2, ...M\}$, from a Weibull distribution, $W[1, c_a^m]$.[22] Each bidder updates her prior based on the history of each game m at stage t, H_m^t. At any stage, the policy of each bidder is the price that she would drop out of each auction if no opponent bidder drops out from this auction when the clock reaches this price.

Definition *Equilibrium is a sequence of bidder policy vectors* $\{D_t\}$ *over each stage* $t = \{1, 2, ...N-$

[22] c_a^m denotes the real number of clicks in slot 1 in auction m.

$2\}$, *bidder beliefs* $\{\tilde{B}_t^m\}$ *and an initial state* $\{D_0^m\}_{\forall m}$ *such that:*

(a) *Each bidder* i *believes that the bidder prior parameters are described by the vector* $\alpha^n = [\alpha_n, ..., \alpha_n]_{Nx1}$.

(b) *The initial state* $\{D_0^m\}_{\forall m}$ *contains the prices that each bidder* i *would drop out from auction* m, *if no opponent bidder drops out when the clock reaches the prices in vector* $\{D_0^m\}_{\forall m}$. *The initial state depends on the bidder prior parameters* α, *the known click curves* $\{C\}_{(N-1)xM}$, *the private signals* S_{NxM} *and the known reserve prices* $\{r\}_{Mx1}$. $\{D_0^m\}_{\forall m}$ *is described by equation 1.5.*

(c) *Bidder policy for stage* $t \in \{1, 2, ...N-2\}$ *of auction* $m \in \{1, 2, ...M\}$, $D_t^m(\alpha, H_m^t, S_m^t, C_m) \in [b_t, v]^N$ *is derived by equation 1.4.*

(d) *Transition to the next stage* $D_{t+1}^m(\alpha, H_m^{t+1}, S_m^t, C_m)$ *from* $D_t^m(\alpha, H_m^t, S_m^t, C_m)$ *is described by Bayesian update of bidder beliefs given the most recent part of history* $H_m^{t+1} = H_m^t \cup \{b_m^t\}$, *where* $b_m^t = min(D_t^m(\alpha, H_m^t, S_m^t, C_m))$.

Proposition 3 *The outcome of the equilibrium is a unique vector of bids that formed after* $N-1$ *stages, is increasing and is equal with* $\{b_0^m = min(D_0^m), b_1^m = min(D_1^m \setminus \mathbb{I}D_1^m(b_0^m)), ..., b_{N-2}^m = min(D_{N-2}^m \setminus \mathbb{I}D_{N-2}^m(b_0^m, ..., b_{N-3}^m))\}$.

Proposition 3 ensures the existence and uniqueness of the outcome of the game. Additionally, the history only matters through beliefs and beliefs are payoff relevant. This displays the Markovian nature of the game, since bidder policies at stage t depend on the inference made over all previous stages through their posterior at stage t.

Corollary 1 *Bidder inference of the signals of their competitors is incorrect in general and correct if and only if all bidders are given the same prior parameter by nature. Hence, the naivety of bidders about the nature's action and their rational decisions for the rest of the game may lead to suboptimal allocation of slots. Bidders do not reach contradictions with their mistaken inference since the system of equations for the opponent's signal and the opponent's prior is underidentified.*

1.6 Estimation

This section aims to provide a method to estimate bidder prior parameters. Despite the richness of the data, a straightforward calculation of bidder prior parameters is impossible. The researcher can observe bids, and the click curve from the data. In addition, data provide enough information to calculate a reliable proxy for the bidders' per click valuations, as described in the data section, outside the marketing component which is estimated separately.

However, there are two components in each bid which cannot be known to the researcher: the bidder's signal in an auction and the bidder's prior parameter for the hotel. However, not all is lost: A bidder is assumed to have a unique prior parameter for all the auctions for a hotel, but a different signal in each auction.[23] I observe a large set of auctions for each property, therefore one can expect the set of bidder signals to closely be described by their distribution due to their sheer numbers. Hence, the knowledge on bids can allow the researcher to pin down the bidder prior parameters. The algorithm I use to perform the estimation is based on maximum likelihood and

[23]This comes from the nature of the problem: The platform decides the click quality in the hotel level, not the individual auction level.

exploits the intuition. For given bidder prior parameters, it is possible to calculate the signal that is associated with each bid.[24] Then, knowing the signal distribution for each auction, I calculate the likelihood that the estimated signals are derived from this distribution. By expanding equations (1.5) and (1.4), I get an expression for the signal s_i that corresponds to each bid b_i.

Proposition 4 *The signal that each bidder receives in each auction can be calculated given the history of the game within an auction and has the following form:*

$$s_i = c(c_2, n, t) \cdot \alpha_i + d(c_2, c_t, b_i, b_t, t, n, r) \tag{1.8}$$

where:

$$c(c_2, n, t) = -\frac{c_2}{n-t+1} + \frac{c_2}{n^2} + \sum_{j=2}^{t-1} \frac{c_2}{(n-j)\cdot(n-j+1)}$$

$$d(c_2, c_t, b_i, b_t, t, n, r) = \frac{n+2}{n-t+1} \cdot \frac{v-b_{t-1}}{v-b_t} \cdot c_t - \frac{1}{n} \cdot \frac{n+2}{n} \cdot \frac{v-r}{v-b_0} \cdot c_0 - \sum_{j=2}^{t-1} \frac{1}{n-j} \cdot \frac{n+2}{n-j+1} \cdot \frac{v-b_{j-1}}{v-b_j} \cdot c_j$$

and:

- s_i : The signal of bidder i in the auction.
- c_2 : Clicks provided by the second highest slot.
- n : Number of participating bidders in the auction.
- c_t : Clicks provided by the t^{th} highest slot.

[24] As mentioned, bids, per click valuations and click curves for each auction are known by the data.

- b_i : The bid of bidder i in the auction.
- b_t : The bid of bidder in slot t in the auction, when lower than bidder i.
- t : The position of bidder i in the auction.
- r : The reserve price in the auction.
- v : The per click value in the auction.

Proof: See Appendix.

The model's assumption that bidders do not recognize that their opponents may have different priors allows for a closed form solution of their signal by using the aforementioned variables from the game's history, and the separate estimation of each bidder's prior separately. The latter property is a result of bidders' naivety believing that their opponents have the same prior with them. Equation (1.8) allows for the estimation of α_i for each bidder i, by using the data from all auctions that bidder i participates. All the needed variables to calculate c and d can be found in the data, and the distribution of the signals is assumed by the model. The first part of deriving an estimator for α_i is to solve equation (1.8) with respect to $d(c_2, c_t, b_i, b_t, t, n, r)$ and get for every auction m that bidder i has participated in hotel h:

$$d^m_{i,h}(c_2, c_t, b_i, b_t, t, n, r) = -c^m_{i,h}(c_2, n, t) \cdot \alpha_{i,h} + s^m_{i,h}$$

Note that the random part is the received signal of the bidder which follows the Weibull distribution $W[1, c_1]$. Variable c is treated as non-random since it does not depend on the bidder's signal but on the auction rules and the equilibrium behavior. The expected value of $s^m_{i,h}$ equals c_1. I subtract c_1 from both sides since having a random part with zero mean makes the derivation of the estimator easier. Hence:

$$\underbrace{\tilde{d_{i,h}^{m}}}_{d_{i,h}^{m}-c_1} = -c_{i,h}^{m} \cdot \alpha_{i,h} + \underbrace{\tilde{s_{i,h}^{m}}}_{s_{i,h}^{m}-c_1} \tag{1.9}$$

Proposition 5 *The estimator for* $\alpha_{i,h}$

$$\hat{\alpha_{i,h}} = -\frac{\sum_{m=1}^{n} \tilde{d_{i,h}^{m}} \cdot c_{i,h}^{m}}{\sum_{m=1}^{n} (c_{i,h}^{m})^2}$$

is consistent, i.e.

$$\lim_{n \to \infty} \hat{\alpha_{i,h}} = \alpha_{i,h}$$

with MSE for bidder i *over auctions indexed by* m*:*

$$\hat{\epsilon}_i^m = \frac{1}{M} \sum_{m=1}^{M} (\tilde{d_{i,h}^{m}} + \hat{\alpha_{i,h}} \cdot c_{i,h}^{m})$$

Proof: See Appendix.

Another property of the estimator is that it equals $\alpha_{i,h}$ on average. This can increase the confidence on the estimator's value even when the sample is not large. However, the availability of data alleviates all concerns of this kind.[25]

[25]Each bidder I examine appears around 15,000 times in auctions for each location in the data. The number of

Proposition 6 *The estimator* $\hat{\alpha_{i,h}}$ is unbiased, i.e.

$$\mathbb{E}[\hat{\alpha_{i,h}}|C_{i,h}] = \alpha_{i,h}$$

Proof: See Appendix.

Finally, the variance of the estimator has a closed form expression. In addition, it does not include any variables unknown to the econometrician, making the standard errors' calculation straightforward.

Proposition 7 $var(\hat{\alpha_{i,h}}|C_{i,h}) = \frac{c_1^2}{\sum_{m=1}^{n}(c_{i,h}^m)^2}$

Proof: See Appendix.

The answer to which bidders are optimistic and which bidders bid lower than optimal should be reflected in their priors. If a bidder has a prior that is higher than the actual value of the ratio of slot 1 to slot 2 total value, then this bidder submits higher bids than the optimal. I estimate bidder priors for the subset of hotels with click through rates without paid traffic clicks. This allows for marketing effect identification. Then I estimate bidder priors for hotels with click through rates without paid traffic clicks and subtract the marketing effect, which should be constant in a certain

instances for each bidder is large due to the number of auctions. Note that when a variable changes value, e.g. days to arrival, it is considered a different auction.

geographical area, e.g. Lower Manhattan, French users, desktop.[26]

Following equation (1.8) for the prior parameter, for bidder i, hotel h in the geographical area l and auction m, equation (1.3) becomes:

$$-\frac{d_{i,h}^{\tilde{m}}}{c_{i,h}^{m}} + \frac{s_{i,h}^{\tilde{m}}}{c_{i,h}^{m}} = ratio_h \cdot opt_{i,h} \cdot mar_i^l \tag{1.10}$$

Equation 1.10 has two unknowns: the marketing effect and optimism effect. Hence, the next challenge is to estimate the marketing effect. This becomes possible by observing that the marketing effect for a bidder is the same across all auctions, in contrast with the optimism effect. Therefore, a subset of auctions where the optimism effect is known can allow for the calculation of the marketing effect. I find this subset of auctions by including hotel in the geographical area for which the ratio of clicks between slots 1 and 2 are known to all bidders. The platform exogenously chooses not to redirect any paid traffic to certain hotels over time. This becomes clear to the bidders, since they do not receive any paid traffic clicks from this hotel and click type (paid/organic) is visible to the bidder. Therefore, the optimism effect for the auctions in hotels with no paid traffic is equal to 1 for all bidders since the ratios are known for this auction. By denoting the set of hotels with known ratios as $\mathbb{H}_0$, I derive the estimator for the marketing effect for every bidder i in geographical location l.

[26]There is no straightforward way to see this in the data because marketing spending has an entangled relationship with the bid. More specifically, the observed bid is shaped by per-click valuation, bias and marketing spending. Hence, the part of marketing spending in bid calculation is not directly visible. Given this, I rely on anecdotal evidence to assume that bidders target their marketing campaigns at the geographical area level, and not by specific searches which would be tedious and impractical.

Proposition 8 *The estimator for mar_i^l*

$$\hat{mar}_i^l = -\frac{\sum_{h\in\mathbb{H}_0}\sum_{m=1}^{n_h} ratio_h \cdot \tilde{d_{i,h}^m} \cdot c_{i,h}^m}{\sum_{h\in\mathbb{H}_0}\sum_{m=1}^{n_h} (ratio_h \cdot c_{i,h}^m)^2}$$

is consistent as the number of auctions per hotel increases to infinity for all hotels, i.e.

$$\lim_{\substack{n_h\to\infty,\\ \forall h\in\mathbb{H}_0}} \hat{mar}_i^l = mar_i^l$$

Proof: See Appendix.

Finally, the estimator for the marketing effect can be used to derive an estimator for the optimism/pessimism effect for every bidder i in hotel h, where $h \notin \mathbb{H}_0$.

Proposition 9 *The estimator for $opt_{i,h}$*

$$\hat{opt}_{i,h} = -\frac{\sum_{m=1}^{n} \hat{mar}_i^l \cdot ratio_h \cdot \tilde{d_{i,h}^m} \cdot c_{i,h}^m}{\sum_{m=1}^{n} (\hat{mar}_i^l \cdot ratio_h \cdot c_{i,h}^m)^2}$$

is consistent, i.e.

$$\lim_{n\to\infty} \hat{opt}_{i,h} = opt_{i,h}$$

Proof: See Appendix.

1.6.1 Identification

While the estimation of the model is based on parametric assumptions, I provide an informal description of the identification strategy here. There are three factors that create variance in bids conditional on per click valuations: prior beliefs, signals and marketing spending. Bidder valuations are identified directly by the data, which is unusual in auction papers and allows me to quantify factors that other papers have not captured before. An important assumption is that I can estimate a common bidder valuation for a click at a given hotel without using the bidding data. This is possible because I observe margins for each bidder and the frequency of conversion and the frequency of cancellation.

The identification of marketing spending is more difficult. It enters the model in a multiplicative way with the bias parameter (see equation (1.3)). This would make their identification impossible unless there was an idiosyncratic characteristic in the data to allow it. Indeed, the data provides an opportunity in the hotels that receive no paid traffic. These hotels are known to bidders while they do not display any different characteristics than other hotels; rather the platform chooses them exogenously. In those no-paid-traffic hotels, by definition, the bias parameter is nonexistent (equal to 1). Therefore, the marketing effect can be identified by the observed bids in the no-paid-traffic hotels. Specifically, the only way to explain any discrepancies between the observed bids and the theoretically optimal bids in auctions in no-paid-traffic hotels is the marketing effect. In no-paid-traffic hotels, any variance in bids across time is attributed bidders learning about opponent signals.[27] The calculation of the theoretically optimal bids is possible because valuations are explained by the bidder per click valuation and every other variable present in the

[27] This is similar with EOS (2007), where they avoid dynamics by replacing the GSPA with the Generalized English auction. The latter, unlike GSPA, allows for bidders to learn within one auction.

equilibrium strategy is observed in the data.

By common intuition, the marketing effect does not only refer to no-paid-traffic hotels, but it explains marketing spending in all hotels in the same area; a basic assumption of the model is that marketing campaigns are organized at a geographical area level and not at the hotel level. All hotels in the dataset are located in the same geographical area thus, identifying bidder marketing spending in no-paid-traffic hotels implies identification of bidder marketing spending for all hotels in the geographical area.

Having identified bidder marketing spending for all hotels using only the observed bids for auctions in no-paid-traffic hotels, the bias of each bidder-hotel pair can be identified by the observed bids in hotels with paid traffic.[28] After accounting for the marketing effect, any distributional discrepancies between the observed bids and the theoretically optimal bids can be attributed to bidder bias, conditional on the signal distribution. By definition, marketing spending and bias are sufficient to identify the prior belief distribution of each bidder for each hotel (equation (1.3)). Finally, signal values explain the individual differences between observed bids and the theoretically optimal bids, conditional on bidder per click valuations, marketing effects, and biases.

1.7 Results

Table 1.1 displays the estimation of the prior parameters, α for each bidder-location combination across auctions, along with their standard errors. The large size of the sample and the estimator's efficiency leads to small standard errors. The latter provides confidence that bidders have statistically different priors among themselves for the same location. The numbers presented in the table

[28] As mentioned, bidders cannot be biased in auction in no-paid-traffic hotels.

represent the mean of the prior distribution of each bidder for each hotel. For example, if the estimated value of α in a hotel-bidder pair is 10, then it means that this bidder acts as her prior belief suggests that there are 10 times more clicks in slot 1 than in slot 2 in this given hotel on average, conditional on their organic click value. Note that the average ratio of organic clicks between slots 1 and 2 across the platform is 12:1. Results much higher than the average, e.g. 100, reflect either intensive marketing spending or excessive optimism. Respectively, a low number, e.g. 3, reflects either excessive pessimism or that the bidder bids purposely low in order to use spending power in other markets. Hotel 0 represents the location with no paid traffic. This location is used for the calculation of the marketing effect. For example, the fact that bidder 1 has a high number for hotel 0 means that she is marketing heavily in this geographical area.

The scale of the estimates seems to vary across bidders and hotels. However, the estimates are constructed to contain both the optimism and marketing effects in a multiplicative way, by definition (see equation (1.3)). This implies that if both the marketing and optimism effects for two bidder-hotel pairs differ tenfold, then the estimates would differ by two orders of magnitude. This explains, for example, the difference in the scale of estimates for bidder 1 and all other bidders.

Table 1.2 displays the real ratios of clicks between slot 1 and slot 2 in April 2021 as observed in the data, for reference and comparison with table 1.1. For example, the ratio of clicks between slots 1 and 2 for hotel 8 is 5.06. Table 1.1 suggests that bidder 1's prior for the ratio of hotel 8 is 11.04, while for bidder 2 it is 2.64. This suggests that bidder 1 is either optimistic about hotel 8 or spends heavily on marketing in the area, while the opposite holds for bidder 2.

Table 1.3 displays the estimated marketing effect for each bidder. Estimates greater than 1

	Bidder 1	**Bidder 2**	**Bidder 3**	**Bidder 4**	**Bidder 5**	**Bidder 6**	**Bidder 7**
hotel 0	42.75 (0.14)	1.89 (0.18)	2.3 (0.06)	1.11 (0.04)	1.64 (0.03)	18.1 (0.02)	10.9 (0.3)
hotel 1	3.62 (0.28)	4.61 (0.04)	0.45 (0.09)	2.12 (0.13)	0.35 (0.02)	0.11 (0.01)	1.36 (0.03)
hotel 2	105.25 (1.63)	5.52 (0.1)	1.21 (0.1)	4.02 (0.14)	6.95 (0.25)	3.32 (0.02)	13.5 (0.29)
hotel 3	17.06 (0.88)	1.68 (0.03)	0.58 (0.03)	0.82 (0.13)	0.74 (0.04)	4.89 (0.03)	4.45 (0.03)
hotel 4	9.19 (0.43)	0.14 (0.02)	9.92 (0.03)	0.85 (0.04)	0.83 (0.03)	7.56 (0.02)	3.28 (0.03)
hotel 5	16.77 (0.19)	6.06 (0.02)	2.37 (0.03)	1.82 (0.1)	1.24 (0.04)	4.52 (0.01)	1.28 (0.02)
hotel 6	4.84 (0.13)	4.43 (0.01)	1.38 (0.01)	1.29 (0.03)	1.19 (0.01)	1.96 (0.01)	1.21 (0.02)
hotel 7	10.6 (0.27)	2.19 (0.02)	0.42 (0.01)	0.91 (0.02)	0.34 (0.01)	3.07 (0.01)	3.36 (0.03)
hotel 8	11.04 (0.02)	2.64 (0.01)	0.43 (0.01)	0.13 (0.01)	0.66 (0.01)	1.08 (0.01)	1.95 (0.01)
hotel 9	17.8 (0.16)	1.11 (0.01)	1.42 (0.01)	2.25 (0.03)	0.14 (0.01)	2.19 (0.01)	1.57 (0.02)
hotel 10	25.27 (0.03)	5.35 (0.01)	0.76 (0.02)	2.15 (0.01)	0.28 (0.01)	2.88 (0.01)	1.29 (0.01)

Table 1.1: Bidder prior parameters (standard errors in parentheses)

hotel 0	**hotel 1**	**hotel 2**	**hotel 3**	**hotel 4**	**hotel 5**
11.38	13.41	12.36	16.12	12.66	11.73
hotel 6	**hotel 7**	**hotel 8**	**hotel 9**	**hotel 10**	
10.16	9.95	5.06	8.40	9.37	

Table 1.2: Click ratios for slots 1 and 2

suggest marketing spending. The larger the estimate is, the stronger is the marketing effect, i.e. bidders 1 and 6 are investing more heavily in marketing than the other bidders in this geographical area.

	Bidder 1	**Bidder 2**	**Bidder 3**	**Bidder 4**	**Bidder 5**	**Bidder 6**	**Bidder 7**
market	3.75 (0.114)	0.17 (0.083)	0.2 (0.074)	0.1 (0.031)	0.14 (0.012)	1.59 (0.109)	0.96 (0.115)

Table 1.3: Marketing effect (standard errors in parentheses)

Table 1.4 displays the estimated optimism index of each hotel-bidder pair.[29] These estimates describe the disturbance in prior belief parameter originating from optimism / pessimism bias. Given its multiplicative nature (see (1.3)) the closer to 1 the estimate is, the less biased a bidder is. Values higher than 1 suggest optimism, while values less than 1 suggest pessimism.

Bidders tend to be consistent in their biases in the sense that bidders are often optimistic about every hotel or pessimistic about every hotel. However, there are occasions of outlier distortions, i.e. a bidder can appear to be more optimistic about a certain hotel, which indicates constantly bigger than what the CTR and per click valuations would suggest. The estimation results for the prior parameter in table 1.1 can sometimes differ by an order of magnitude. This should not come as a surprise, since results for hotel 0 refer to auction with only organic traffic, i.e. marketing effect. In the model, the marketing effect interacts with the optimism/pessimism parameter in a multiplicative way. This is the reason that bidders with aggressive marketing campaigns (e.g. bidder 1, bidder 6) can have larger prior parameters for some hotels (e.g. hotel 2). The estimators are structured such that when the priors are adjusted using the marketing effect, they are compared with the real ratios observed for each hotel in table 1.2.[30] Therefore,

[29] I include the bidders that appear sufficiently often in each location, i.e. more than 1,000 instances.

[30] An example that displays the effect of the estimates of each table follows. The real ratio of clicks in hotel 2, as found in the data, is 12.36 (table 1.2). Bidder 7 displays a marketing effect equal to 0.96 (table 1.3) and a bias for

	Bidder 1	**Bidder 2**	**Bidder 3**	**Bidder 4**	**Bidder 5**	**Bidder 6**	**Bidder 7**
hotel 1	0.11 (0.002)	1.31 (0.141)	0.31 (0.097)	1.4 (0.006)	0.17 (0.028)	0.02 (0.015)	0.06 (0.044)
hotel 2	2.38 (0.263)	3.45 (0.17)	0.54 (0.084)	1.33 (0.254)	2.84 (0.123)	0.16 (0.041)	1.04 (0.088)
hotel 3	0.2 (0.096)	0.56 (0.085)	0.39 (0.121)	1.03 (0.106)	0.27 (0.038)	0.18 (0.007)	0.24 (0.011)
hotel 4	0.32 (0.006)	0.85 (0.04)	3.04 (0.182)	0.52 (0.039)	0.21 (0.008)	0.4 (0.054)	0.47 (0.062)
hotel 5	0.28 (0.119)	2.1 (0.446)	0.53 (0.127)	1.28 (0.111)	0.92 (0.085)	0.21 (0.01)	0.11 (0.003)
hotel 6	0.17 (0.005)	1.56 (0.117)	0.57 (0.12)	2.13 (0.149)	0.85 (0.074)	0.12 (0.045)	0.24 (0.064)
hotel 7	0.35 (0.007)	1.68 (0.124)	0.23 (0.107)	1.61 (0.115)	0.32 (0.08)	0.16 (0.005)	0.22 (0.008)
hotel 8	0.57 (0.171)	4.66 (0.852)	0.34 (0.098)	0.9 (0.011)	0.54 (0.22)	0.3 (0.073)	0.61 (0.104)
hotel 9	0.72 (0.086)	0.72 (0.148)	0.43 (0.009)	2.1 (0.091)	0.48 (0.13)	0.15 (0.076)	0.34 (0.052)
hotel 10	0.82 (0.031)	1.69 (0.22)	0.38 (0.017)	1.07 (0.3)	0.24 (0.005)	0.14 (0.014)	0.19 (0.002)

Table 1.4: Optimism effect (standard errors in parentheses)

the magnitude of the results in table 1.4 appears to be considerably smaller, with bidders being usually consistent in their biases. Any discrepancies within bidders can be attributed to imperfections of their prediction algorithms which create their biases.[31]

1.8 Counterfactual

In the counterfactual, I evaluate the decision of the platform to publish the ratios of organic to paid traffic for every hotel. In this case there is no uncertainty about the click through rate, and therefore the environment returns to EOS (2007) setup, marketing spending persists, and information asymmetry disappears. In this context, the clicks in slot 1 are known, hence the game is identical with the one studied by EOS. Bidders bid according to the formula derived by EOS:

$$b_i = v_i - \frac{c_k}{c_{k+1}}(v_i - b_{k+1}) \tag{1.11}$$

where i represents bidder id, and k represents slot id.

Results indicate that platform revenue increases by 7% when the ratios are revealed. As expected, bidders increase their bids the most for the auction where they demonstrated the highest level of pessimism, while the opposite holds for cases of optimism. Average bidder return decreases by 11% since competition for slot 1 increases and less biased bidders are not able

hotel 2 equal to 1.04 (table 1.4). Therefore, bidder 7 believes that slot 1 exhibits $12.36 \cdot 0.96 \cdot 1.04 = 12.34$ times more clicks than slot 2 for hotel 2, on average. This number is in accordance with the estimate of the prior parameter for this bidder-hotel pair found in table 1.1.

[31] There are instances in the data where the top 2 bidders bid up to 50 times higher than their counterparts in specific auctions. Commissions do not vary this much across bidders; hence it is unlikely that those bids come from an accurate information set.

to extract rents.

Figure 1.6 presents the relationship between optimism in the bidders' prior and the change in seller revenue after the introduction of the policy. I calculate an optimism index for each bidder-hotel combination using results from tables 1.2 and 1.4. If the number in table 1.4 is bigger than the corresponding number in table 1.2 then the respective bidder is optimistic for the CTR in the respective hotel (ratio ¿ 1). Similarly, a ratio smaller than 1 implies pessimism while if the ratio equals 1, then the bidder's prior is correct on average. In general, the closer is the ratio to 1, the less is the bidder's bias about the CTR.

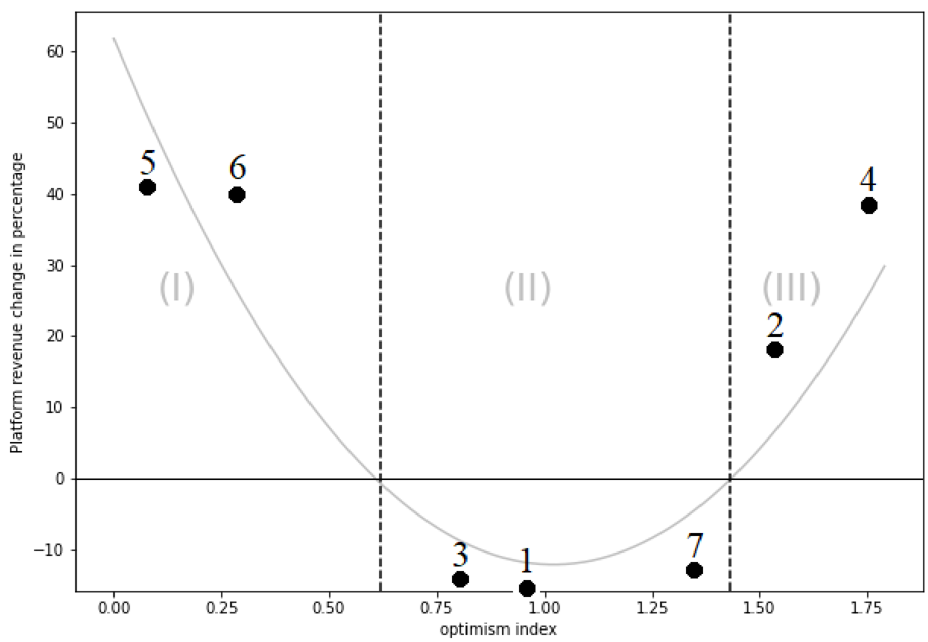

Figure 1.6: Platform revenue percentage change for different levels of optimism/pessimism

The figure includes the best fitting line for all observations of the ratios. The curve is U-shaped and can be interpreted by splitting it in 3 areas. Area (I) contains bidders who were significantly pessimistic in the pre-policy environment. The platform profits from these bidders when CTR becomes known. The reason for the latter is intuitive since these bidders start the auction with low priors on average. As the auction progresses, bidders in area (I) update their beliefs and understand that clicks in slot 1 are more than their initial belief suggests. However, the low priors lead bidders in area (I) to drop out early more often than they should. For this reason, these bidders increase their bids when CTR becomes known and subsequently platform revenue increases too.

Area (II) contains less biased bidders, with ratios close to 1. These bidders can extract rents in the pre-policy environment and lose this benefit when CTR becomes known. When bidders can be biased, more biased bidders can benefit from opponents that drop out early and can avoid overpaying for clicks in slot 1. These benefits stop when CTR becomes known. The latter makes the marginal return of increasing the bid to decrease, and therefore leads to smaller bids. The latter is translated to the difference for platform revenue being in the negative area. The result for area (III) is counterintuitive. Even though bidders in this area are optimistic, seller revenue increases when CTR becomes known. The optimism effect leads those bidders to avoid dropping out of the auction early.

However, as the auction progresses, bidders in area (III) observe other bidders dropping out and infer their signals. A basic assumption of the model is that bidders believe that their opponents have the same prior as they do. If a pessimistic bidder from area (I) drops out early in the auction, a bidder from area (III) interprets the latter as a bidder with the same prior with them dropping out. Given their belief that every bidder has the same prior with her, an early drop out can be interpreted only as a very low signal. These signals accrue as the auction progresses, leading bidders in area (III) to be more pessimistic than bidders in area (II), contrary to initial conditions. The equation that describes signal inference in the proof of Proposition 4 displays this dynamic, i.e. the inferred signal is decreasing in the bidder's prior parameter, α.

It is important to note that the previous analysis is platform-specific solely and applies solely to Tripadvisor. For example, there could be another platform where there are no pessimistic bidders. Hence, the absence of early dropouts would not allow optimistic bidders to adjust their beliefs downwards. This would translate to a revenue reduction from optimistic bidders after information about CTR is disclosed. Therefore, revenue analysis and bidder behavior vary with the data. This paper does not attempt to be provide generic results about CTR information disclosure. However, the paper attempts to be universal in the game theoretical aspect of the problem and equilibrium conditions. The parts of counterfactual analysis and policy recommendation are tailored to Tripadvisor and they might differ for other online platforms.

Figure 1.7 shows the change of platform revenue in percentage terms for each bidder separately. It should be noted that this graph does not perfectly correlate with how optimistic or pessimistic the bidder is estimated to be. Each bidder's payments depend on the next highest bid. Given that bidder rankings for each auction in the studied geographical area is not uniform, the revenue that the platform generates from each bidder is affected by the prior of the bidder that is usually one slot below.[32]

Figure 1.8 displays the estimated marketing effect and optimism index for every bidder-hotel observation. In addition, it includes the quadratic curve that fits this data best. The fact that pessimistic bidders appear to have higher marketing spending on average explains the positive outcome of the policy. In an environment of similar per click valuations, high marketing spending increases the likelihood of winning a high slot. As explained before, pessimistic bidders are

[32]The set of the top 3 bidders is the same in 46% of the auctions in the studied geographical area, e.g. bidder 3, bidder 1, bidder 6, ...

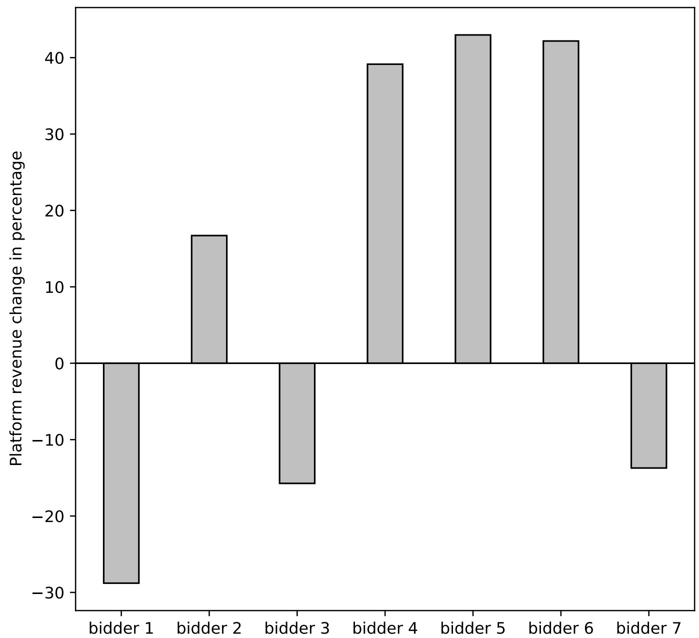

Figure 1.7: Platform revenue percentage change per bidder

expected to increase their bids after the introduction of the policy. More specifically, the policy boosts the spending of the bidders who spend the most on the auction on average, due to more intensive marketing. Therefore, it is straightforward that the policy's net effect should be positive for this platform, which is consistent with the counterfactual results.

The previous paragraphs describe how more biased bidders tend to underbid when there is asymmetric information and bidder bias in a generalized English auction environment. However, internet platforms perform the generalized second price auction, where bids are submitted simultaneously. Whereas the generalized English auction is a useful tool to examine the learning process of bidders, it is natural to ask how underbidding dynamics translate in a realistic setting. Consider a heavily biased bidder that enters an auction. More specifically, assume that the bidder has a prior belief about the return of slot 1 which is significantly higher than its actual return. Since bids are

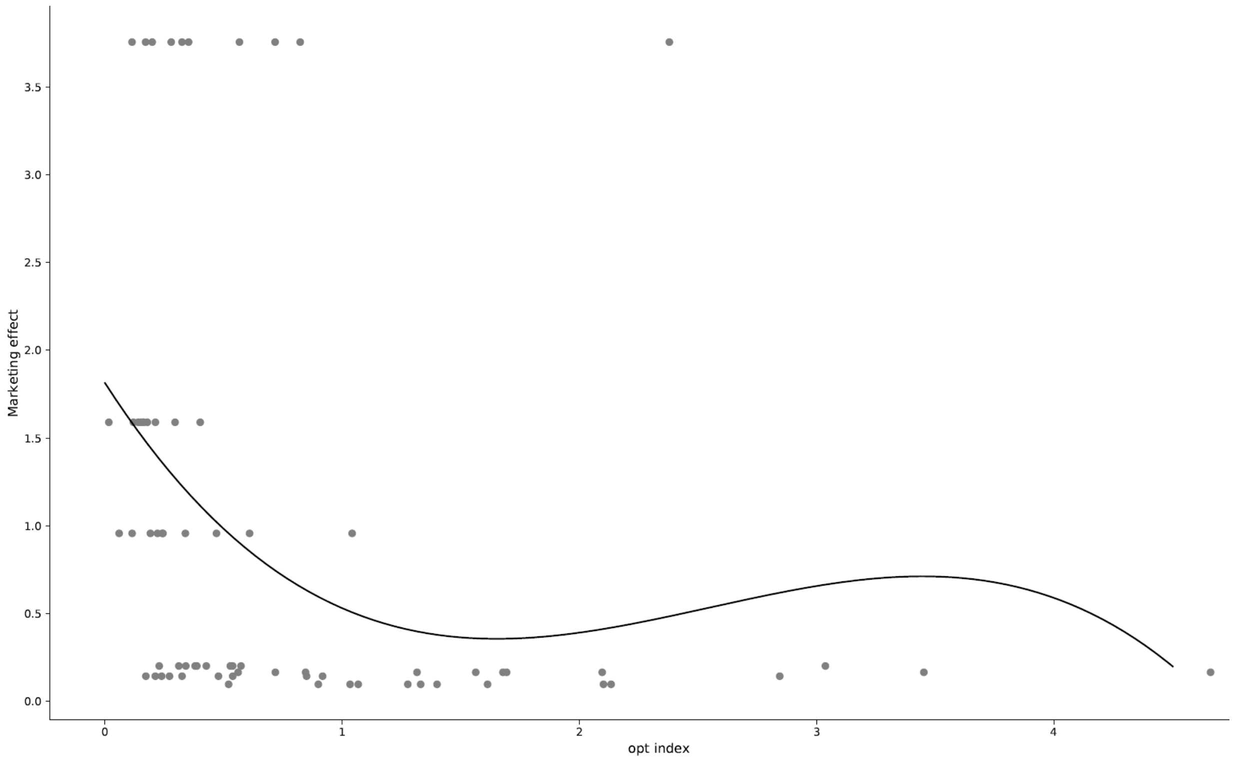

Figure 1.8: Marketing spending and Optimism

submitted simultaneously, it is likely that the bidder wins a high slot. The bidder has no access to her opponent's bids for confidentiality reasons.

In practice, bidders slightly change their bids when the auction is repeated in order receive information. When the bidder reduces her bid, she eventually receives a lower slot, which reveals the bid of the bidder who previously held this slot. Since this bidder is initially optimistic, she gradually understands that most opponents' bids are lower than her prior would suggest. Hence, the more optimistic a bidder initially is, the bigger this discrepancy becomes. The bidder quickly understands that slot 1 is not as desirable as believed, which results in a bid reduction. In addition, winning slot 1 and have perfect visibility on its return is not an option in a platform's scale. Bidders submit millions of bids daily, which makes the evaluation of each auction separately practically impossible.

In an environment that the click through rate remains the same across time, a bidder's knowledge would convert to the true value after a certain amount of iteration. However, in a platform's environment, user preferences and exogenous choices of the platform e.g. amount of paid traffic, change frequently. The latter implies that bidders are not able to realize their bias, making bidders susceptible to overbidding constantly.

In the appendix, figures A.2, A.3, A.4, A.5 show histograms that count the difference for platform revenue in percentage terms for each bidder separately and a histogram that counts the difference for platform revenue in percentage terms for every auction in the dataset in case the platform adopts the policy. Obviously, the amount paid to the platform from a certain bidder depends on actions of other bidders. The change in one bidder's payments is not equivalent with the effect of the implemented policy to this specific bidder. However, I abusively display the revenue per bidder changes in order to provide a broader picture where effects from opponent's actions are averaged. In accordance with figure 1.7, bidders 1 and 7 have the volume of their observations in the negative territory, while bidder 3 observations are more uniform but the clicks in the auctions with positive change are less valuable. Figures A.6, A.7, A.8, A.9, A.10 show histograms that count the difference for platform revenue in percentage terms for each hotel separately.

1.9 Conclusion

The Sponsored Search Auction is a significant monetization mechanism for several internet platforms. The literature contains important theoretical results on the possible equilibria of the SSA; however the underlying assumptions are questionable in practice. In a real setting, it is highly unlikely that all bidders are equally biased about the number of clicks that an ad in slot 1 will receive after the auction. A bidder needs to guess the number of users that will access a listing, the quality mix of clicks (organic/non engaged) and her opponent's beliefs. Since bidders are expected to have access to different amounts of resources, information asymmetry should arise in most internet platforms that hold SSAs.

Anecdotal evidence points to the direction of limited bidder information about the click through rates. The fact that user interaction and behavior fluctuate over time suggests that click through rates are also volatile. In addition, the variance in bidder ability to process the data from auction results amplifies the asymmetry. Therefore, the assumption that information about click through rates is asymmetric across bidders and that their prior beliefs can be biased becomes plausible. The asymmetric information environment and prior belief bias have important implications for bidder strategy and seller revenue. Intuitively, less biased bidders can take advantage of their more biased counterparts and extract rent using their information edge. I design a game of the generalized English auction with bidders receiving different signals that reflect their beliefs on the click through rate.

Furthermore, I model different bidder ability to interpret past auction data by allowing bidders to have different prior distributions about the click through rate. The latter allows bidders to be consistently optimistic or pessimistic about the value of slot 1 in a given geographical area on average. Then, I derive the Bayesian Nash equilibrium of this game in an identical private value environment. The equilibrium condition implies that bidders infer wrongly opponent signals when their priors are optimistic or pessimistic, making them susceptible underbidding. Then, I derive a consistent estimator for the prior distribution's parameters for each bidder-hotel combination based on past bids and bidder per click valuations. I use data on bids and valuations to estimate the prior parameters while accounting for over(under) - bidding due to marketing reasons. Estimation results suggest that bidders have indeed significantly different beliefs about the click through rates. Finally, I perform a counterfactual to assess whether disclosing CTR to bidders is beneficial for the seller in Tripadvisor's case. I find that the platform increases its revenue by 7% when the policy is implemented in the studied markets. The revenue increase is caused by bid adjustments from more biased bidders, which offset the bid reductions from less biased bidders that are not able to extract rents.

Limitations of the paper's Bayesian approach include the following. First, it cannot be applied successfully in platforms with many small sized bidders with limited processing capability. It is easier to justify that responding to opponents' actions and learning from their bids is possible with a few big players rather than many small players. Second, the equilibrium I derive applies solely in identical private value environments, which is the case for Tripadvisor. The equilibrium conditions differ in cases with significant differences in per click valuations, which I do not solve. Finally, my estimators need a large amount of data to converge, due to the large number of estimated parameters. My approach would not produce robust results in the case of a much smaller platform.

An interesting expansion would be to establish a general rule on the conditions that deem bidder bias favorable to the platform and develop a decision tool which allows the identification of the optimal information policy given data on past bids.

Chapter 2: Optimal Reserve Prices with Endogenous Demand: A Field Experiment

2.1 Introduction

Sponsored search auctions are used as a monetization mechanism by many online platforms. In general, online platforms blend their organic content with advertisements. Different parties are competing for the advertisement space on the platform, and therefore, user attention. The platform allocates the space by holding an auction which determines the price for each bidder. It is plausible to assume that the goal of the platform is to maximize its intertemporal revenue. Hence, the most interesting question to ask, is how to design the auction in order to generate the highest expected payoff to the seller across time.

The first theoretical results answer the question for the case of the sale of one item and independently distributed private values. Riley and Samuelson (1981) and Myerson (1981) show that if bidders are symmetric, the mechanism that maximizes the expected revenue of the seller is a second price auction with an appropriately chosen reserve price. Bulow and Roberts (1989) simplified the analysis by showing that the allocation problem is equivalent to the analysis of standard monopoly third-degree price discrimination. The latter introduced the marginal revenue of the bidder into discussion. Bulow and Klemperer (1996) provide a simple derivation of the result that the expected revenue from an auction equals the expected marginal revenue of the winning bidder(s).

The theoretical work extends to the allocation of multiple objects. Armstrong (2000) analyses optimal auctions of multiple objects when bidders have a binary distribution over their valuations for each object. Ben-Porath et al. (2014) discuss optimal allocation with costly verification from the principal's side. Maskin and Riley (1984) derive optimal mechanisms in settings with risk-averse bidders. Ausubel (2004) derives an efficient ascending-bid auction for homogeneous goods. With private values, this auction yields the same outcome as the sealed-bid Vickrey auction. Edelman and Schwarz (2010) derive the optimal auction design in Sponsored Search Auctions. They analyze the underlying dynamic game of incomplete information, and they establish an upper bound on the revenue of any equilibrium of any dynamic game in this environment. They show that a platform's optimal reserve price is independent of the number of bidders and independent of the rate at which click-through rate declines over positions.

Several papers discuss how the aforementioned theoretical results compare with empirical estimates of optimal reserve prices for different kinds of auctions, such as Bajari and Hortacsu (2003); Haile and Tamer (2003) and Tang (2011). McAfee and Vincent (1992) consider a common value auction model with bidder participation determined jointly by nature and by bidder optimization. Then, they derive a test statistic for establishing when it is optimal to raise the reserve price. Cramton et al. (2002) outline the differences between equation-based Transaction Evidence Pricing System and parity pricing for timber pricing in British Columbia. McAfee, Quan and Vincent (2002) compute the equilibrium bidding strategies and outcomes, and derive a lower bound on the optimal reserve price in a general auction model with affiliated signals, common components to valuations and endogenous entry.

A paper closely related to this research is Ostrovsky and Schwarz (2016). They present the results of a large field experiment on setting reserve prices in auctions for online advertisements, guided by the theoretical results from Edelman and Schwarz (2010). Their experiment was successful and consistent with the theory, revenues increased substantially after the new reserve prices were introduced. Additional papers have experimented with reserve prices in a first-price online auction (Reiley (2006)) or use and extend the framework proposed by Ostrovsky and Schwarz, confirming their findings (Sun et al. (2014) and Topinsky (2014)). Another related paper is Rafieian (2020). Rafieian examines the revenue gains from adopting a revenue-optimal dynamic auction to sequence ads. The paper uses theoretical framework to derive the revenue-optimal dynamic auction that captures both advertisers' strategic bidding and users' ad response and app usage. Rafieian documents significant revenue gains from using the revenue-optimal dynamic auction compared to the revenue-optimal static auction.

In the current paper, I allow the demand side to react to changes of the supply. In particular, I present empirical evidence that platform users engage less when the ad supply is limited. The latter happens because the ad space in the platform I study can be interpreted as organic content by the users. This creates a trade-off between increasing revenue by increasing the reserve price, versus higher demand by allowing more bidders to appear in the search results. Adding to Edelman and Schwarz (2010) environment, I allow click thru rate (CTR) to change when the number of winning bidders changes, following the intuition that when users see fewer offers they tend to engage less with the platform. I derive the optimal reserve prices computationally and evaluate their performance in a field experiment as in Ostrovsky and Schwarz (2016).

2.2 Auction Overview

The paper focuses on Tripadvisor's hotel meta-auction, which has the general characteristics of sponsored search auctions. Tripadvisor holds repeatedly simultaneous generalized second price auctions for different user searches. A user search is defined as combination of search variables such as hotel id (location id), length of stay, days to arrival and number of guests. First, the user sees a list of hotels that match her search, as displayed in figure 2.1. The order of appearance for hotels after step 1 is not related to the auction but is personalized for each different user. OTAs only compete for user attention within hotels. When a user clicks on a hotel, a limited number of travel agents (OTA) is displayed in a list form, as displayed in figure 2.2. Different positions have different desirability for OTAs. When a user clicks on an OTA's listing, she is redirected to the OTA's website, as displayed in 2.3. Then, the OTA pays the platform for sending the user to its website - "pay-per-click" rule. All positions are auctioned at once. However, bidders can submit only one bid. The highest bid wins the highest slot and so on, until the auction is repeated.

Winner determination:

The bidders who bid more than a preset reserve price, r , are ranked in terms of decreasing bids. Rarely, this is not the final order of appearance, since further adjustments take place to the rules.

When the order is determined, a bidder pays when it receives a click, regardless of whether the click leads to a booking. The pricing rule can vary, producing different variants of the auction

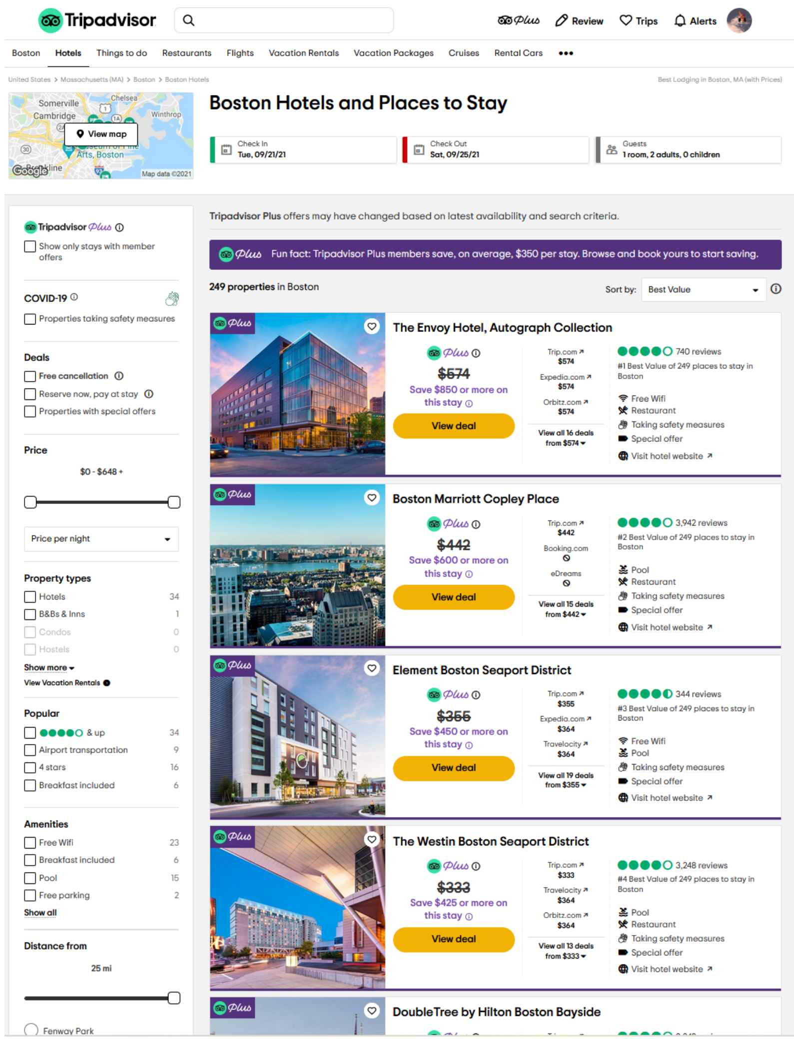

Figure 2.1: Hotel list

with price equal to:

1. The submitted bid $max(r, b)$; (GFPA)

2. The next highest bid $max(r, b_1)$; (GSPA)

3. A hybrid version: $max(b_1, \gamma \cdot b, r), \gamma \in (0, 1)$; (GHPA)

The pricing scheme used by Tripadvisor in 2021 is the second price payment, which is the revenue maximizing scheme as suggested in Dimitrellos (2020). In addition, using the result of Dimitrellos (2021), I assume that CTR is common knowledge, since it suggests higher platform

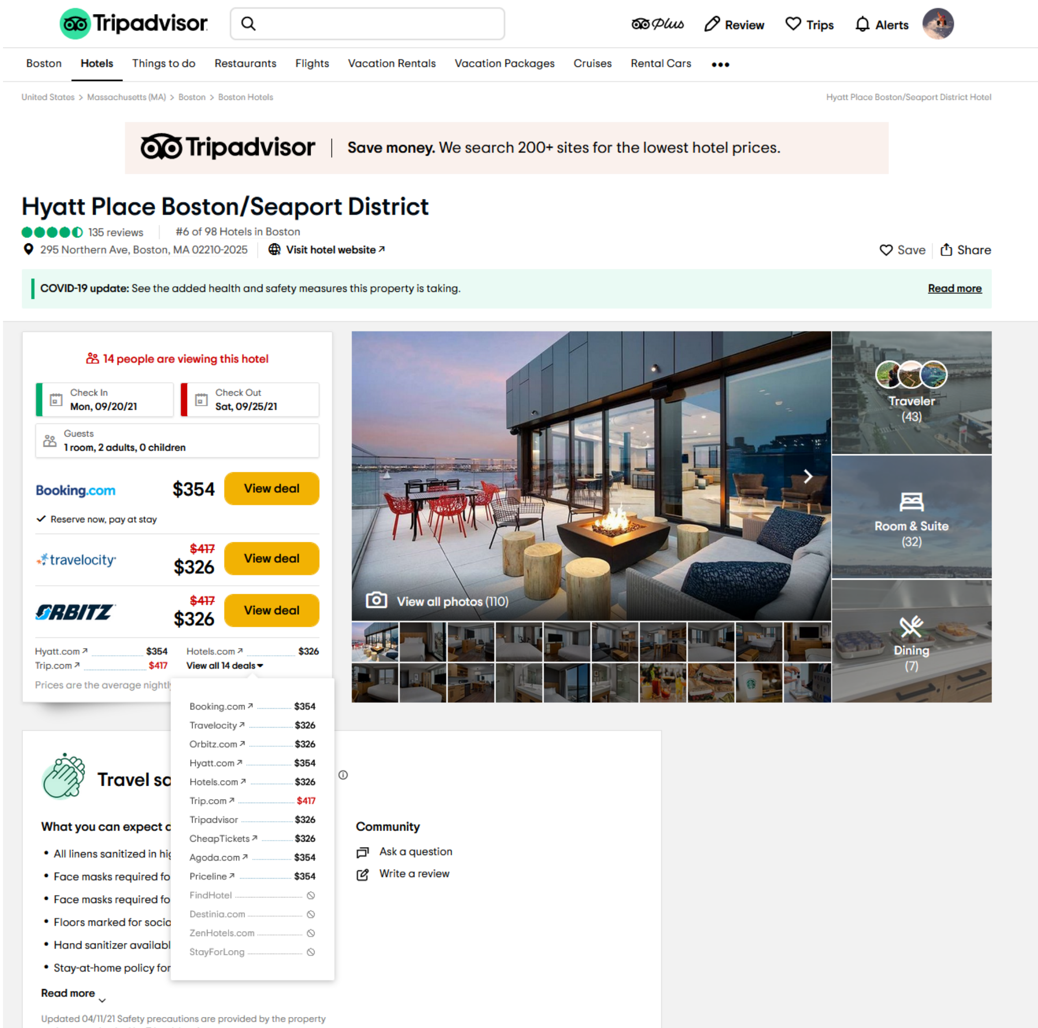

Figure 2.2: OTA display

revenue compared to the asymmetric information case. Bidders are allowed to update their bids twice a day. Each bidder is notified about the position she achieved in case of winning, their cost per click, and if she bid less than the reserve price. Theoretically, it is possible to infer the position of other winning bidders, since each auction's ranking is available to platform users. Given the platform's nature, it is plausible to assume that users prefer more OTA listings to fewer if they prefer more choices. This is not the case for platforms such as Google, where users do not enjoy seeing more ads and try to avoid them when possible.

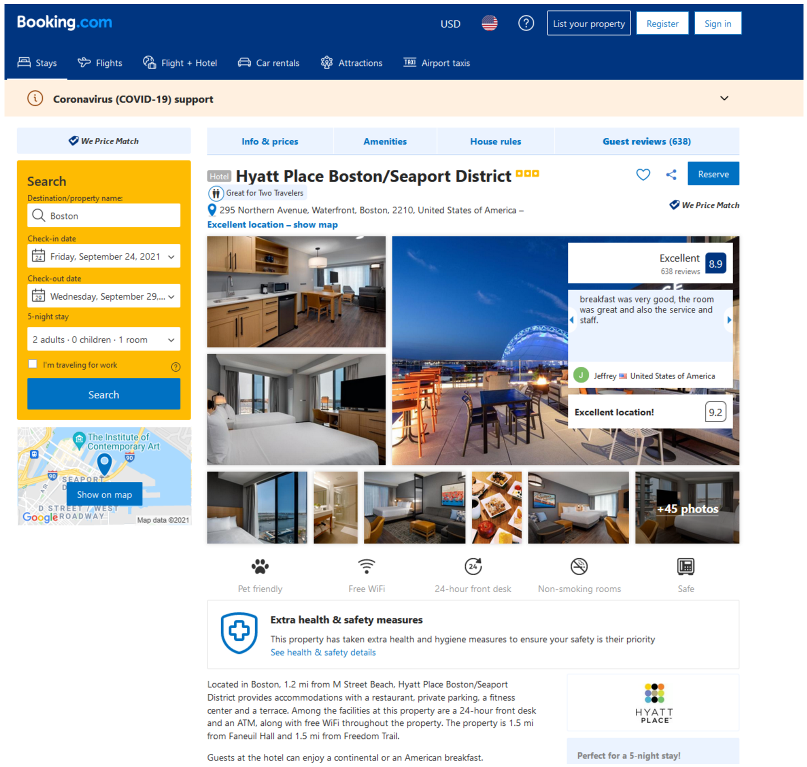

Figure 2.3: Redirection

2.3 Optimal Reserve Price

Given the previous results, the last aspect of the auction mechanism that the platform can modify are the reserve prices. Edelman and Schwarz (2010) consider a dynamic game of incomplete information used to sell sponsored search advertisements. They also consider a corresponding static game of complete information. They establish an upper bound on the revenue of any equilibrium of any dynamic game in this environment and they prove that the generalized English auction with a certain reserve price is an optimal mechanism. The current paper provides evidence that the generalized English auction environment performs better than alternative schemes and a natural question that occurs is how Edelman and Schwarz (2010) results are translated in the current auction's environment.

I expand the assumptions of Edelman and Schwarz by adding a responsive demand side. In the case of Edelman and Schwarz, the platform in discussion auctions ad slots, and display advertisements next to organic results. Users visit the platform in order to have access to its organic results and not advertisements. Therefore, an increase in reserve prices as suggested by Edelman and Schwarz, not only increases the platform revenue but, to the users' delight, decreases the amount of displayed advertisements and provides faster access to the organic results.

On the contrary, Tripadvisor does not auction advertisement slots, but organic content. A reduction in the displayed travel agents by an increase in the reserve price could lead to reduced user engagement due to less supply of organic content. I quantify the effect of reduced supply to the click through rate and I calculate the optimal reserve prices that balance the effect of increased per click revenue and decreased clicks caused by higher reserved prices. I provide evidence on the validity of the theoretical result by conducting a field experiment that compares current reserve prices with the theoretically proposed ones, in the fashion of Ostrovsky and Schwarz (2016).

2.4 Setup

First, I empirically estimate the demand's side reaction to the size of supply. Intuitively, a user that encounters a low amount of offers for a search has the option to look at other platforms in order to form a more educated opinion about the market price of her booking. This is equivalent with the assumption that users earn utility from variety. I also assume that user taste does not change when the auction mechanism changes, which allows its separate estimation.

The user is informed about the number of offers after the search is performed. After the search, the user sees the list of available offers (supply) and decides to click an offer in case the user is interested.[1] I estimate how much less the users click on an offer as supply is being reduced, i.e. the clicks to searches as the number of offers changes. It would be questionable to measure the change in click to search ratio when supply changes in different auctions, because auctions can display different characteristics. For example, auctions for bookings far in the future are less likely to be clicked compared to auctions for bookings in the near future.

Instead, I use data of supply changes within auctions, i.e. exogenous supply changes for an auction. I regress the logarithm of the click to search ratio on the logarithm of the ratio of the number of offers for this auction to the historically maximum number of offers for this auction. I find that when the number of offers is less than 9, the click to search ratio increases by 4.1% on average when the ratio of the number of offers for this auction to the historically maximum number of offers for this auction increases by 1%. I find no positive effect in the increase of supply, when the number of offers is 9 or more. Table 2.1 displays the regression results, where β denotes the coefficient of the ratio of the number of offers for this auction to the historically maximum number of offers.

The drawback of this estimation approach is that it potentially suffers from endogeneity issues. A common reason that an advertiser may not appear in the search results is that they have sold all their available rooms for this particular search. If popular advertisers run out of inventory first, then one part of the click reduction effect found above could be accounted to the fact that users do not see their favorite advertisers anymore and, therefore, click less. Indeed, specific bidders run out of inventory more often than others. However, this is not because

[1] Unlike Jeziorski and Segal, (2015) I assume that the number of clicks that each ad receives, conditional on its position and total number of ads, is not influenced by the identities of other advertisers that are showed on the screen.

Table 2.1: Supply Reduction Effect table

	$click/search\ ratio$
$const^{**}$	-0.29
	(0.011)
β^{**}	0.041
	(0.01)
Observations	2,802
$95\%\ C.I.\ for\ const$	$[-0.311, -0.269]$
$95\%\ C.I.\ for\ \beta$	$[0.022, 0.06]$

standard errors in parentheses
**: significant at 5% level

these advertisers' brand is stronger than their opponents'; it is discussed in the first chapter of this dissertation that no advertisers appear to have a better click to conversion ratio than their counterparts. The reason that specific bidders run out of inventory first is that these bidders win the highest slot more often, which leads to more clicks and conversions. When an advertiser is out of inventory, they are removed from the auction results and all bidders below them shift one slot higher. Hence, the first slot is always occupied and this is the reason that the method above estimates the supply reduction effect correctly.

Following Edelman and Schwarz, I calculate the reserve prices that maximize seller revenues in a direct revelation mechanism, using formula 5.12 from Krishna (2002). The optimal direct revelation mechanism can be characterized using the same technique as in a single object case, except that the probability of receiving an object is replaced with the expected number of clicks that a bidder receives. Denote $x_k(r)$ the expected number of clicks received by bidder k when the realized vector of bidder values is given by $r = (r_1, ..., r_N)$, and let $f(r)$ denote the pdf of vector r and $t_k(0)$ the expected payment of bidder k when her value is zero. Finally, I assume virtual valuation $\psi(s) = s - \frac{1-F(s)}{f(s)}$ is non-decreasing in s, in order to satisfy the regularity condition from Myerson (1981). Then, the seller revenue in a direct revelation mechanism is:

$$\sum_{k=1}^{N} t_k(0) + \int_r (\sum_{k=1}^{N} \psi(r_k) x_k(r)) f(r) dr \tag{2.1}$$

The main difference with Edelman and Schwarz analysis is that in this environment, the expected number of clicks depends on the total number of offers. Edelman and Schwarz provide an analytical solution for the maximization of seller revenue, where only bidders with positive virtual valuations are allocated a positive expected number of clicks, and bidders with higher virtual valuations are allocated higher positions. The aforementioned solution does not maximize seller revenue in this environment. The (negative) marginal effect of a lower reserve price in the average per click payment is outweighed by the (positive) marginal effect of increased clicks due to increased supply. The latter allows, in some cases, bidders with negative virtual valuations to be allocated a positive expected number of clicks, since their participation creates a positive externality for all bidders from increased user engagement.

2.5 Estimating the Distributions of Bidder Values

The seller revenue expression in 2.1 and bidder strategic behavior in Edelman, Schwarz and Ostrovsky (2007) (EOS) are a stylized representation of the platform's environment, however it serves as a useful approximation of reality. Iyengar and Kumar (2006) and Roughgarden and

Sundararajan (2007) follow a similar approach with EOS. In this section, I outline the estimation of the optimal reserve prices that maximize 2.1. The proposed algorithm consists of two parts. The first part estimates the distribution of bids and valuation and verifies that the regularity condition is satisfied. The second part numerically calculates the reserve prices that correspond to the target value of the virtual valuation for each auction. The company picked the set of auctions that serve as the input of the algorithm, with goal to use the proposed reserve prices in the subsequent experiment. The auctions belong to one of the top 10 user countries in terms of platform revenue. During the month before the experiment, the hotels chosen for the treatment group received around 32,000 clicks in this user-country, while the hotels in the control group receive around 225,000 clicks. These observations occur from user clicks during May 2021.

For each search in the sample, the two following moments were computed: The average bid, and the average standard deviation of the bids, where the average was taken over all searches. The bid of the highest bidder in every auction was excluded from the statistics, because the theory does not allow the researcher to pin it down. The number of bidders that participate in the auction does not need to be calculated, since the data provide full visibility on all bids, including those that win no slots because they are lower than the reserve price.

Next, as in Edelman and Schwarz, it was assumed that bidders' values were drawn from a lognormal distribution with a mean and a standard deviation to be estimated. Lognormal distribution ensures that the regularity condition is satisfied for the parameters that occur from Tripadvisor's auctions. Next, I simulate the two moments for possible true values of the mean and the standard deviation of the lognormal distribution of values. More specifically, for each combination of true values of the variables to be estimated, the algorithm creates five hundred draws of the vectors of bidder values. For each draw, equilibrium bids were computed, using the unique perfect Bayesian equilibrium of the Generalized English Auction in EOS (2007). The moments of interest were then computed, and averaged over all draws of bidder values.

For each search, the parameters of the distribution of bidder values were estimated by matching the observed moments to the simulated results. Finally, for each keyword, the theoretically optimal reserve price was numerically computed in two steps: First, I calculate the value of ψ that maximizes the seller's revenue. Note that the optimal values of ψ are negative because of the reduced supply effect, while in Edelman and Schwarz is always zero. Second, I calculate the reserve price that makes the virtual valuation equal to its optimal value, given the estimated distribution of values. The regularity condition ensures that the optimal value of the virtual valuation is unique and lies on the global maximum of the seller's revenue. The algorithm (in pseudocode) is presented below.

```
for all the auctions of interest do
    Calculate the ratio of clicks for consecutive bids;
    for all the parameter values that the valuation distribution can have do
        for i from 1 to 500 do
            Draw valuation sample, with dimension equal to the the number of bidders;
            Calculate bids(ratio,valuation) based on the theoretical equation;
        end
        Calculate mean, std of the calculated bids;
        Calculate moment distance between the calculated bids and the bids from data;
        Keep the parameter values that minimize this distance;
    end
    for i from 1 to 500 do
        Draw valuation sample, with dimension equal to the number of bidders;
        For each valuation calculate its ψ;
    end
    Choose the cutoff point for all ψ's [a], that maximizes equation 2.1, on average;
    The optimal reserve price, is the value that makes the psi function equal to the cutoff
      point;
    Save the optimal reserve price;
end
```

Algorithm 1: Optimal Reserve Price

[a] All ψ's smaller than this point are discarded

The existing reserve prices in Tripadvisor have been set years ago and are outdated since they do not reflect the current values of the auctions. More specifically, the correlation of current reserve prices to current bids is 0.18, while the correlation of the theoretically optimal reserve prices is 0.88, which reflects the fact that the optimal reserve price has a strong positive correlation with bidder valuations. Figure 2.4 displays the distribution of the current reserve prices, and Figure 2.5 displays the distributions of the theoretically optimal reserve prices.

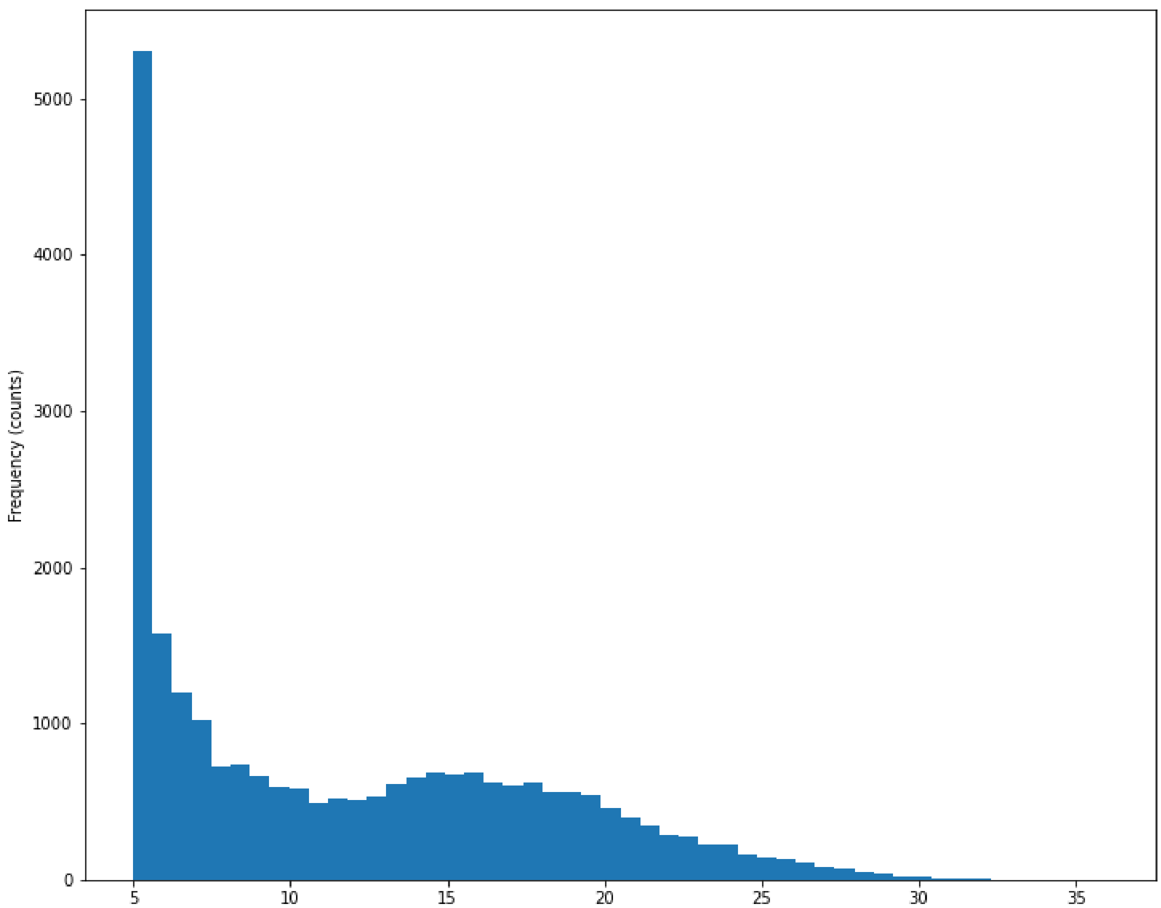

Figure 2.4: Current reserve prices

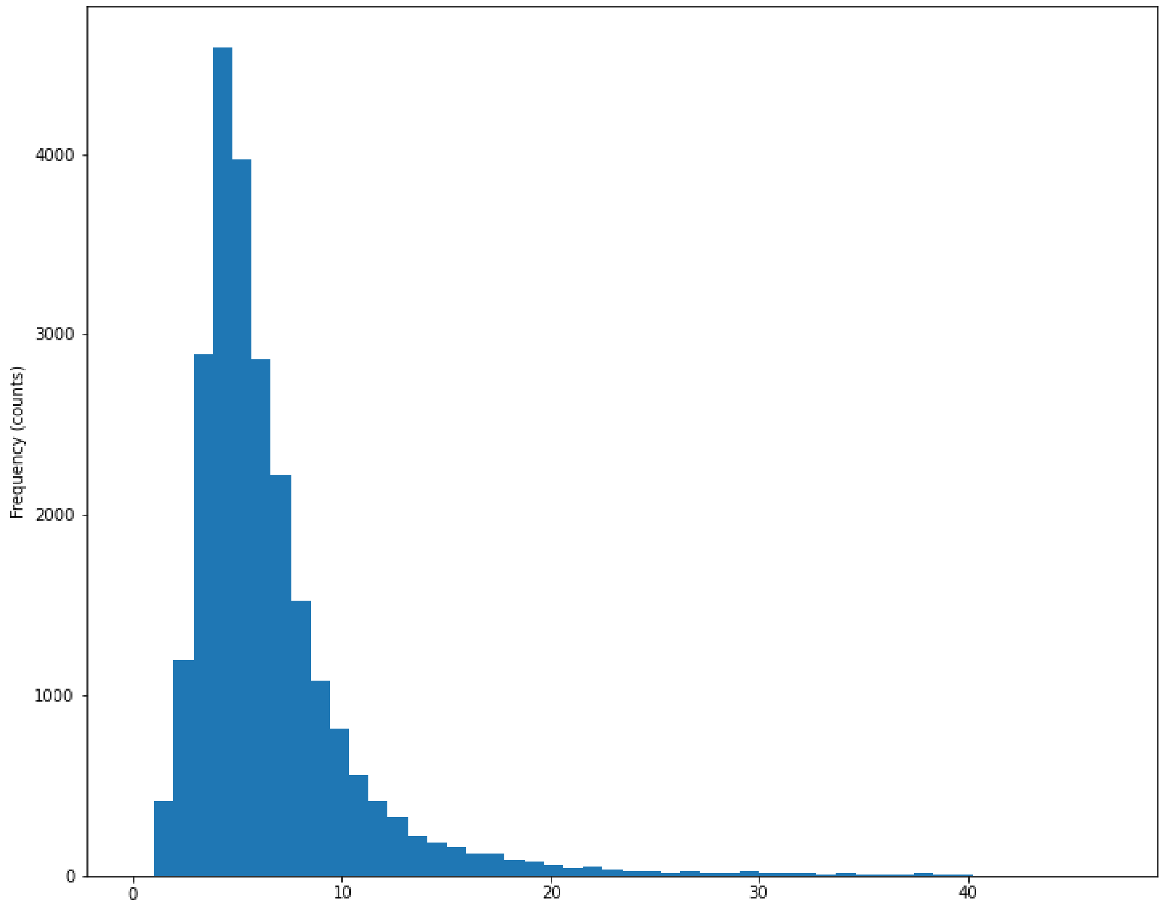

Figure 2.5: Theoretically optimal reserve prices

The details of the estimation process that affect the aforementioned results follow: First, I do not consider bidders that have a quality score. As in EOS (2007), I assume that a certain slot will receive a certain number of clicks regardless of the bidder in this slot. Second, I assume that click through rates are the same, given hotel-user country-device. In other words, I assume that click through rates do not change when variables like number of guests change. I use the average click though rate for every hotel-user country-device across other variables. Third, the per click value that bidders assign to each slot is assumed to be the same. Given a click, it is irrelevant for the bidder that it comes from a certain slot. If a user is interested in a listing, all past information on the auction (slot, opponent bidders) do not matter any more.

Furthermore, I used bidding data that are more than a month old, since bidders can be strategic if they know that their bid affect reserve prices in the short term. Finally, the platform has chosen not to announce the experiment to the bidders following the previous rational: If the bidders know that an auction belongs in the treatment group, then they can change their bids in the short term in order to favor the control group which contains lower reserve prices.

2.6 Experimental Results

In this section, I discuss the results of the experiment and analyze the characteristics of the control and treatment groups. The control group was chosen to contain more hotels than the treatment group in order to avoid significant revenue loss in case of a negative outcome. Following Ostrovky et al. (2016), I did not use different adjustment factors in the treatment group.[2] Furthermore, my dataset does not contain any hotels with searches of different order of magnitude compared to the mean. Hence, as expected, removing the top 0.1% of hotels does not create any statistically significant outcome.

2.6.1 Pre-Experiment Analysis

Before discussing experimental results, I present summary statistics on the differences between the control and treatment groups for May 2021; which is before the start of the experiment. The goal is to show that the two groups have no significantly structural characteristics which

[2]Ostrovsky et al. chose to divide the treatment group into further subgroups with various levels of reserve prices. Different subgroups had different adjustment factors and their reserve prices were set: (optimal reserve price)· (adjustment factor) + (10 cents)·(1 - adjustment factor). They did not find any systematic differences between different adjustment factors, hence I opted not using them.

would contaminate the analysis of the treatment effect. The variables I examine are the revenue per search, submitted bid, depth, and reserve price.[3] I normalize both the revenue per search and submitted bids to 1 over this time period.[4] Table 2.2 reports the summary statistics for both groups before the experiment. The average search happens in an auctions with 7.06 active slots with average pre-experiment reserve price around 8.9¢. The only statistically significant difference found between the two groups in the pre-experiment period is in the Depth variable. I find that an auction in the treatment group has 0.78 more winning bidders than an auction in the control group on average. This difference is small, therefore I consider the split to be plausible for post-experiment analysis. Finally, note that the optimal reserve prices are not necessarily higher than the existing ones on average. While there are several instances of significant increase, the reduced supply effect drives other optimal reserve prices down.

Table 2.2: Summary statistics and test of treatment–control balance

Variable	All	Treatment	Control	Difference	p-value
Revenue per search†	1	1.005	0.994	0.011	0.3368
	(1.0616)	(1.0820)	(1.0546)	(0.0064)	
Submitted bid	1	1.009	0.996	0.013	0.2864
	(1.0392)	(1.0551)	(1.0359)	(0.0062)	
Depth	7.061	7.277	6.497	0.78	0.0210**
	(4.1488)	(4.4122)	(3.5856)	(0.0187)	
Reserve price	8.898	8.785	8.957	-0.172	0.1397
	(6.6578)	(6.3353)	(7.0899)	(0.0266)	
Sample size	257,249	32,138	225,111		

† Revenue per search and Submitted bids are renormalized, to the average value of one across the overall sample. *, **, *** – significant at 10%, 5%, and 1% levels. Numbers in parentheses give the standard deviations for the statistics in Columns 1–3 and the standard errors for the differences in Column 4.

[3]Depth refers to the number of bidders that win a slot, and therefore are shown to users.

[4]The goal of the normalization is to keep revenue per search and bids confidential. I use the same formulas for normalization as in Ostrovsky and Schwarz (2016): Suppose that we observe n searches, each of them generating a revenue rps_i, $i = 1, 2..., n$. Then the normalized revenue per click is $\hat{rps}_i = cpc_i \frac{n}{\sum_{j=1}^{n} rps_j}$. Respectively, the normalized bid is $\hat{b_i} = b_i \frac{n}{\sum_{j=1}^{n} b_j}$.

2.6.2 Post-Experiment Analysis

I measure the treatment effect on the aforementioned variables (revenue per search, submitted bids, depth) using differences-in-differences estimates as in Ostrovsky and Schwarz (2016). I measure the post-treatment effect on the control group in order to account for external shocks, and I compare it with the post-treatment effect on the treatment group.[5]

The post-experiment data include the user effect: Users encounter a different amount of providers in the post-experiment period because of the new reserve prices. This leads to a change in their behavior, since increased supply may create additional clicks. The main variable to be measured, revenue per search, is affected by both the direct reserve price effect and the indirect demand response effect. I do not measure the magnitude of each effect in this analysis, but the net effect.[6] This is not problematic, since the variable of interest for a revenue maximizing platform is the total effect of a policy to its revenue.

The analysis' approach is to calculate the treatment effect for the full sample at first, and then to partition the full sample to subsamples based on variables of interest in order to identify the cases where the treatment effect is the strongest. The first partition is about the popularity of each hotel in terms of user searches. One would expect that bidders should react more promptly in more popular auctions, since they affect their revenue more. Then, I partition the subsample of popular hotels into two more parts based on the optimal reserve price. Intuitively, the depth of the auction should increase more in a case of a lower optimal reserve price since the user effect may

[5]The pre-experiment period is considered to be May 2021. The experiment was deployed in mid-June 2021. Data from June 2021 were discarded in order to allow some time for the bidders to internalize the changes and adapt their strategies. The post-experiment period was chosen to be July and August 2021.

[6]However, the estimated supply reduction effect (see Table 2.1) allows for the estimation of the users' response magnitude.

dominate the increased reserve price effect. Finally, I partition the subsample of popular hotels and high optimal reserve prices into 2 more subsamples based on auction depth. Theory suggests that an optimal reserve price generates additional seller revenue from every bidder, even those in the higher slots (see Edelman and Schwarz (2010)). Therefore, one would expect the auctions with higher depth to show a higher increase in revenue per search compared to the auctions of lower depth.

2.6.3 Results

The results for the full sample are reported in Table 2.3. The only statistically significant effect (at the 5% level) of the new reserve prices is an increase in the depth of the auction by 0.64 advertisers on average. This is expected, since the existence of additional advertisers tend to increase user interaction with all advertisers and therefore platform revenue. The algorithm that calculates the new reserve prices can internalize this externality, hence the increase in average depth. Results suggest that revenue per search increases by 10.54%, however this result is significant only at the 10% level.

The partition of the full sample in terms of hotel popularity offers a clearer picture. In hotels with more than 200 searches per month, the increase in revenue percent becomes 19.58% and it is significant at the 5% level. Bidders tend to be more attentive to hotels that generate the most bookings for them. Hence, it is more likely that a bidder adjusted her bid for an auction in this subsample in the optimal way as the theory suggests. The latter is also supported by a significant (at the 10% level) increase in submitted bids. Finally, the increase in depth is bigger than in the overall sample, adding 1.39 listings on average (significant at the 5% level). The significant increase in depth suggests that a part of the revenue per search increase comes from increased user interaction due to additional supply, highlighting the importance of this externality. Note that this subsample contains 36% of the searches, but generates 41.4% of revenues for the platform.

The subsample that contains the less popular hotels shows no statistically significant differences between the control and treatment groups. This can be attributed to the fact that these hotels individually receive very few clicks hence they create very few booking for the advertisers. It is plausible to conclude that an advertiser did not notice any changes in hotels that bring less than one click per day.

Table 2.3: Results (full sample, split by search volume)

	Full sample	$\geq$ 200 searches per month	< 200 searches per month
Δ-in-Δ Revenue per search	10.54%*	19.58%**	5.82%
t-statistic	[1.610]	[1.673]	[0.737]
p-value	(0.0537)	(0.0472)	(0.2305)
Δ-in-Δ Submitted bid	4.64%	13.73%*	2.57%
t-statistic	[1.223]	[1.541]	[0.201]
p-value	(0.1106)	(0.0616)	(0.4203)
Δ-in-Δ Depth	0.64**	1.39**	0.02
t-statistic	[1.855]	[1.930]	[-0.768]
p-value	(0.0318)	(0.0268)	(0.7789)
N. obs. in treatment group	77,859	66,633	11,226
N. obs. in control group	127,163	7,596	119,567
Fraction of total revenue	100%	41.4%	58.6%

*, **, *** – significant at 10%, 5%, and 1% levels. Changes in Revenue per search and Submitted bid are reported relative to the average revenue per search and average Submitted bid in the corresponding subsample before the experiment.

As Ostrovsky and Schwarz suggest, each percentage point of positive impact translates into potential improvements to search engine profits and revenues on the order of hundreds of millions of dollars per year. Furthermore, it is possible for the platform to notify the bidders about the changes in reserve prices once the policy becomes permanent. This will lead bidders to possibly update their behavior in less popular auction and further improve the results.

The fact that the treatment and control groups are imbalanced in terms of searches per month does not affect the validity of the aforementioned results. Table 2.4 shows a comparison of the hotels with more than 200 searches per month for both groups. It becomes clear that the hotels belonging in the two groups are not fundamentally different, hence the difference-in-difference approach is valid.

Table 2.4: Summary statistics of treatment and control groups for more that 200 searches per month

Variable	Treatment	Control
Number of searches†	1.004	0.984
	(1.0251)	(1.0441)
Submitted bid	0.990	1.031
	(1.0193)	(1.0685)
Per night price	1.001	0.999
	(1.0508)	(1.0847)
Hotel stars (5)	3.592	3.592
	(2.0731)	(2.0732)
Customer reviews (10)	7.416	7.429
	(6.1950)	(6.4075)
Sample size	66,633	7,596

† Number of searches, Submitted bids, and Per night prices are renormalized, to the average value of one across the overall sample for privacy reasons. Numbers in parentheses give the standard deviations for the means in Columns 1 and 2.

In the next two subsections, I further partition the subsample of popular hotels into more subsamples: auctions with high and low theoretically optimal reserve prices, and then auctions with high and low depth.

2.6.4 Results by Reserve Price Level

Another partition that is potentially interesting is to separate the sample in terms of high or low optimal reserve prices. As the optimal reserve price increases, the more possible becomes to differ more to the old reserve price (note that the old reserve prices had almost no correlation with submitted bids, while the new optimal reserve prices show a correlation of 0.88). Thus, one would expected a higher increase in revenue per search when the optimal reserve price is high. In contrast, a larger increase in depth is expected when the new reserve price is low, since the entry cost becomes lower for bidders. Table 2.5 reports the results, and they verify both claims: The increase in revenue per search surges to 30.06% in the subsample with higher optimal reserve prices. However, results are not as statistically significant as in Ostrovsky and Schwarz. The main reason for this is that in the current environment it is not always optimal to increase the reserve price. Some of the old reserve price were in higher level than optimal, while at the same time, increasing the reserve price decreases supply and therefore user engagement.

Consistent with theory, bids increase more in the auctions where reserve prices increase more (25.17% increase compared to 13.76%). The main reason for this is that when reserve prices increase some marginal bidders increase their bid to stay in the auction, applying pressure to bidders in higher slots. However, these findings are statistically significant only at the 10% level. Finally, there is a significant increase in depth for auctions with low optimal reserve prices; 1.61 more advertisers on average, significant at the 5% level. The corresponding increase in the auctions with high optimal reserve prices is just 0.67 additional advertisers per auction on average.

Table 2.5: Results (keywords with at least 200 searches per month, split by the level of estimated optimal reserve price

	Full subsample	$r^* \geq 10¢$	$r^* < 10¢$
Δ-in-Δ Revenue per search	19.58%**	30.06%*	20.06%*
t-statistic	[1.673]	[1.541]	[1.633]
p-value	(0.0472)	(0.0617)	(0.0512)
Δ-in-Δ Submitted bid	13.73%*	25.17%*	13.76%*
t-statistic	[1.541]	[1.499]	[1.488]
p-value	(0.0616)	(0.0669)	(0.0684)
Δ-in-Δ Depth	1.39**	0.67*	1.61**
t-statistic	[1.930]	[1.308]	[1.902]
p-value	(0.0268)	(0.0954)	(0.0286)
N. obs. in treatment group	66,633	13,537	53,506
N. obs. in control group	7,596	2,268	5,328
Fraction of total revenue	41.4%	11.72%	29.68%

*, **, *** – significant at 10%, 5%, and 1% levels. Changes in Revenue per search and Submitted bid are reported relative to the average revenue per search and average Submitted bid in the corresponding subsample before the experiment.

2.6.5 Results by the Number of Advertisers

I keep the subsample of popular hotels and auction with high optimal reserve prices and I partition it further in terms of auction depth. Theory suggests that the effect of imposing the optimal reserve price on revenue should be higher in percentage terms for the auctions with less participants. Ostrovsky and Schwarz partially verify this claim with their findings. My results are shown in Table 2.6. Auctions with more participants increase their revenue by 25.78%, while auctions with less participants increase their revenue by just 23.96%. Note that auctions that reduce their depth are ”punished” by the users, which explains the difference of outcome when I partition based on auction depth. This analysis does not reach the same conclusion with Ostrovsky and Schwarz for two reasons. First, in the current case, a low depth implies lower user interaction. The platform’s revenue is heavily affected by users who click less after their searches.

Second, the dataset becomes significantly smaller in this partition. Treatment and control groups represent only 8.62% and 3.1% of total revenue respectively. The small size of the dataset leads to higher variance and larger confidence intervals for the estimates. Consequently, there is no estimate which is statistically significant at the 5% level.

Table 2.6: Results (keywords with at least 200 searches per month and the estimated optimal reserve price of at least 10 cents, split by the average number of advertisers)

	Full subsample	depth > 7.5	depth < 7.5
Δ-in-Δ Revenue per search	30.06%*	25.78%*	23.96%
t-statistic	[1.541]	[1.301]	[1.207]
p-value	(0.0617)	(0.0966)	(0.1137)
Δ-in-Δ Submitted bid	25.17%*	19.70%	21.71%
t-statistic	[1.499]	[1.214]	[1.193]
p-value	(0.0669)	(0.1123)	(0.1164)
Δ-in-Δ Depth	0.67*	0.38	-0.67*
t-statistic	[1.308]	[1.205]	[1.341]
p-value	(0.0954)	(0.1142)	(0.0900)
N. obs. in treatment group	13,537	9,146	4,391
N. obs. in control group	2,268	1,161	1,107
Fraction of total revenue	11.72%	8.62%	3.1%

*, **, *** – significant at 10%, 5%, and 1% levels. Changes in Revenue per search and Submitted bid are reported relative to the average revenue per search and average Submitted bid in the corresponding subsample before the experiment.

2.7 Conclusion

The results of the experiment suggest that incorporating the effect of reduced supply when applicable can lead to substantial increases in the platform's revenue. The realized increase (e.g. almost 20% for popular searches) is substantially larger than the improvements found in existing literature, highlighting the importance of incorporating demand's response in reserve price changes.[7] Increased user engagement as a result of increased supply does not only benefit the marginal bidders entering in the lower slots, but it creates a positive externality for the bidders in higher slots, that receive the majority of clicks.

This paper describes a way to measure how user engagement changes as supply decreases. In addition, it develops an algorithm to numerically calculate the optimal reserve prices with endogenous user behavior. The findings of the paper can incentivize other platforms who use sponsored search auctions to allocate organic content, to study the effect of reduced choice to user engagement. This way, a platform can be able to balance the positive effect of increasing a low existing reserve price (as in Ostrovsky and Schwarz) with the negative effect of providing less options to their customers. The results of the experiment confirm the theoretical finding that the introduction of optimal reserve prices while considering the reduced supply externality has led to a significant increase in seller revenues.

[7] The comparison is merely suggestive, since different platforms have different levels of optimality in their existing reserve prices.

Chapter 3: The Tragedy of Commons in the Taxi Industry: A Case for a Central Dispatcher

3.1 Introduction

How would a Social Planner affect a spatial market with search frictions? The peculiar structure of spatial markets creates a number of externalities that have an important impact on the market's efficiency. The long literature on matching and network search gives insights on how agent actions affect market supply and demand. In addition, existing results establish the concept of dynamic spatial equilibrium in spatial markets in the presence of frictions. This paper addresses the equilibrium effects of externalities, in order to give economic insights on market policies that could achieve more efficient allocations. I do not aim to prove the superiority of central planning against the competitive equilibrium, as in certain circumstances the Central Dispatcher could produce strictly suboptimal outcomes compared to the market. However, being able to internalize already known information gives a measurable advantage to the policymaker in mitigating external effects. I estimate the magnitude of the aforementioned frictions by adding additional detail to existing models, such as congestion and endogenous matching efficiency. The main forces that reduce the social surplus are business stealing, a congestion externality and an endogenous demand externality. First, a central planner can internalize the effect of supplier's decisions on other suppliers. The literature suggests significant business stealing when suppliers are identical, which leads to excess market entry. Second, the central planner can internalize the effect of a supplier choices on market congestion. When congestion is high, matching efficiency can be affected. Market participants are unable to internalize their individual impact on congestion, while a central planner can compensate for this impact. Finally, when market participants are a small fraction of the total supply, they take demand as given, ignoring the effect of plenteous supply on consumer beliefs. I make these additions to the existing benchmark model and I estimate their effects in order to calculate the optimal market allocation in terms of social surplus.

The industry used as a spatial market paradigm in this paper is the taxi industry. The taxicab industry remains a critical component of the transportation infrastructure despite the fierce competition that it faces from the ridesharing industry. Existing regulations in multiple cities maintain the status of the taxicab industry as a significant factor of the market. The lack of centralized control in urban taxi markets supports the claims of market inefficiency. Taxi drivers make their search decisions based on the maximization of individual profit, ignoring the aggregation of their choices, which determines the supply in any location. Buchholz (2020) develops an empirical model of the search frictions resulting from price regulation. There is additional literature that studies the benefits of dynamic pricing for ride-hail services: Hall et al. (2015), Castillo et al. (2017). Price regulation fails to allow different locations to become equally attractive to drivers and leads to lack of coordination of demand and supply. Buchholz discusses the effect of price regulation on demand and supply surplus. Buchholz shows that when prices are allowed to vary by time, location and distance, daily net surplus raises by $194,000 and there are 31,000 additional daily taxi-passenger matches. Finally, Buchholz solves for the market equilibrium in the case of a search-frictionless market.

This paper expands the model proposed by Buchholz by, first, deriving the optimal centralized decisions in the case where search frictions are still present and second, by adding endogenous congestion and endogenous demand to the existing model.[1] Congestion is modeled as a sigmoid function of the number of taxis present in a location. The functional form is selected in order to represent the empirical observation of the existence of a critical point in the effect of the number of vehicles on traffic speed. I allow congestion to affect the travel cost of vehicles, matching efficiency and consumer demand. Travel cost has been observed to be higher in high traffic conditions, and demand for taxis is expected to decrease when congestion is higher. I allow demand to change with respect to the number of taxis present in a location, to capture the change in consumer beliefs about match probabilities when plenty of taxis are searching in the area. These extensions allow the central dispatcher to internalize a more complete set of market frictions. I keep price regulation in both the baseline model and the counterfactual, since price deregulation has been studied extensively by Buchholz. I perform a reduced form test of the significance of the Tragedy of the Commons, i.e. business stealing. Results support the existence of business stealing, and hence they indicate that there is room for policy on driver strategy.

The counterfactual studies the effect of a Central Dispatcher. The Central Dispatcher makes decisions instead of drivers to maximize social surplus. Specifically, in the benchmark model, when a driver is searching in a location and fails to match with a passenger, they must choose the next search location. In this counterfactual, this decision lies with the Central Dispatcher. The drivers have to comply with the Central Dispatcher's instructions regardless of their idiosyncratic

[1]In contrast with Buchholz's frictionless counterfactual, I derive the optimal allocation while taxis can still fail spotting nearby passengers.

shocks. Two different algorithms are used in order to approximate the optimal allocation due to the high dimensionality of the state space. The first algorithm uses a greedy criterion by optimizing only with respect to the next time period. The second algorithm uses an intertemporal approach based on an approximation of the objective function. This approximation is obtained by fitting the original objective function onto its estimation via a one layer neural network. The quality of the fit obtained by the latter indicates that this algorithm's result is not far from the optimum. Theoretically, I prove that the estimation process is unbiased. Results suggest that the Central Dispatcher increases social surplus by 3.13% while benefiting both drivers and passengers by 2.06% and 3.21% respectively. In monetary terms, this translates to an increase of social surplus of $798 thousand per shift. In addition, traffic speed increases by approximately 10%, with the strongest effect for Mid and Upper Manhattan, and daily matches increase by 4.8% on average. These results indicate the importance of coordination between taxis and the need for increased planning in their actions.

The rest of the paper is organized as follows. Section 2 presents existing literature in detail. Section 3 introduces the data, and section 4 discusses reduced form evidence that supports the business stealing hypothesis. Section 5 presents the structural model and section 6 discusses the estimation results. Finally, section 7 introduces the Central Dispatcher counterfactual, section 8 reports the results and section 9 concludes.

3.2 Literature

This paper is built upon the literature on network effects and search and matching. This literature has focused on the New York City taxi industry because of the availability of rich data. The Taxi and Limousine Commission of New York City releases detailed taxicab records. Researchers have modeled driver behavior at different levels of detail over the years. Over time, models have become more detailed by endogenizing features of the industry such as its dynamic nature, driver beliefs, the matching mechanism, and consumer demand. Brennan (2014) provides a rich description and valuable insights on the Taxi market of New York City.

First, Cairns and Liston-Hayes (1996) present an aggregate model of the taxi industry in order to evaluate several ways of regulating the market. Their model is not agent-based and demand is treated as a function in equilibrium, at the aggregate level. First, the authors solve the monopoly case and derive the monopoly number of taxis, number of hours per taxi and fare. Next, they solve for the same variables in the social optimum. The social surplus maximizing solution implies negative profit for taxis and therefore it cannot be achieved. The authors also solve for the second best solution, where they add a zero profit condition for the taxis. In the second part of the paper, Cairns and Liston-Heyes argue that an equilibrium cannot be achieved in an unregulated market. This is because of the positive search cost of passengers, so a taxi driver can always bargain for a slightly higher price than the supposed equilibrium price. Then, the authors use their model to show that regulating just the fare leads to a worse equilibrium than the second best because of excess entry. Finally, the authors suggest that the most efficient type of regulation is to fit both the price and the number of taxis, given that regulating working hours is not realistic.

Lagos (2003) is the benchmark model for dynamic spatial analysis of the taxicab market. While the level of detail in this model is small, it has the basic mechanisms of demand, supply and a spatial dimension that more complicated future papers present, e.g. Buchholz (2017). The main assumptions of the model are:

1. Drivers know other drivers' chosen strategies.

2. The model achieves a steady state.

3. Demand is exogenous and does not respond to endogenous shocks, such as changes in waiting time.

The model setup is as follows: There are distinct locations which are characterized by their demand, which can vary among locations. The matching mechanism is frictionless, and the number of meetings is $m_i = min[taxis_i, passengers_i]$. The author assigns a value to the driver for each location that depends on the matching probability, the expected revenue and the value of waiting. In equilibrium, there is no arbitrage, and passenger movement is characterized by a steady state. These conditions allow for a simple closed form solution for the value function in equilibrium, from which the author derives the number of taxis in each location. The final allocation depends on the aggregate market tightness (cabs/passengers) and has two main potential outcomes, one with excess supply in most locations and one with excess demand. The main policy result is that changes in fares and the number of taxis do not improve the total number of matches in a significant way.

In an attempt to study further the effects of regulation in the industry, Frechette, Lizzeri and Salz (2018) develop a dynamic general equilibrium model of the cab industry. The basic question they study is the effects of matching frictions and regulatory limitations on efficiency. The market characteristics that create the inefficiencies are first, taxi drivers' lack of knowledge about passenger location and vice versa, and second, the fixed prices per mile, imposed by regulation.

The latter makes the market clear not through price adjustment but through waiting time for passengers and search time for taxis. The model assumes that drivers have complete information on their current period search time, costs and revenues per hour. Drivers choose when to start and end their shift but not their locations. This is justified by the fact that in equilibrium the search time in each area adjusts to make drivers indifferent. Demand is estimated in a reduced form way. The market forces that lead to equilibrium are the following:

1. Waiting time determines demand through the demand function.

2. Supply and demand determine the waiting and search times through the matching function.

3. Drivers' start and stop decisions determine supply, using revenues and costs as inputs.

4. Search time determines revenues to make drivers indifferent in terms of location.

As mentioned before, demand is estimated in a reduced form fashion. The first step is to estimate the matching function, which takes demand, supply and exogenous variables as inputs and gives the waiting and searching times as outputs. The authors estimate the matching function with simulation, with the objective of matching the simulation result to the average observed taxi search time. Next, they invert the estimated matching function, to get the value of the demand. The demand function is determined by a regression of the estimated demand on time fixed effects and waiting times. The authors use shift change shocks as an instrument in order to deal with the endogeneity, of waiting time with respect to demand. Then, given the demand function and using the matching functional form, one can calculate waiting and search times. Obviously, these calculations have to take place in a simultaneous equation environment as one cannot estimate demand without knowing the waiting time and vice versa.

The supply side is estimated using a more structural approach. Drivers choose starting and end times to maximize their utility. The cost is assumed to be a function of the driver's hours worked and exogenous characteristics. There is a random utility shock and an expected value of future utility from continuing to drive. If the random shocks are assumed to be Type I extreme value, then one can derive a closed form expression of the probability of a driver stopping his/her shift at a given time. The starting decision is modeled in a similar way. Given how many drivers start and stop their shift at a given time, the authors derive the supply of taxis. Finally, revenues are determined by search time in order to make all locations equally attractive. This system of four main equations comprises the dynamic equilibrium of the market. The main findings after estimating the model are the following. First, taxis that operate within a fleet follow the 5am to 5pm shift pattern more stringently than owner-operated taxis. Second, the owner-operated taxis have higher costs but less convex cost functions, which leads to smoother stopping behavior. Third, the effect on daily entry decisions is highest for earnings increases in the first hours of the shift. This is because the probability of actually receiving these earnings decreases as the shift becomes longer due to increasing costs and repeated exposure to the outside option. Finally, a more effective matching mechanism increases the welfare of both drivers and passengers, although it reduces market thickness.

I study a search and matching framework with dynamic oligopoly in the tradition of Buchholz (2020), who analyzes a dynamic spatial equilibrium and discusses market inefficiencies in terms of waiting times and spatial mismatch. The author focuses on driver decisions rather than demand formulation. Buchholz develops a model which accounts for cab location in great detail. The flow of demand, matches and vacancies in every location are described by a state transition matrix. The basic characteristics of the model follow:

- An urn-ball based matching function is used to determine the matches in a location at a certain moment, given the number of passengers and taxis in this location.

- Each driver decides his/her own location based on his/her beliefs on the distribution of other drivers, rather than each other driver separately, to minimize the computational burden. By assuming T1EV random shocks, the probability of each driver to go to a certain location can be expressed in closed form. This decision takes place only if the taxi is vacant. Otherwise, the driver has to go to the passenger's destination.

- The driver is ”inactive” when not vacant, in that there are no decisions or shocks until the passenger drop off.

- Demand follows a Poisson distribution that depends on price and time fixed effects, but not on changes in waiting time.

- The number of taxis in each location is determined by the multiplication of the previous state with the transition matrix.

The solution for the equilibrium follows Buchholz (2020), building on the notions of Markov perfect equilibrium by Ericson and Pakes (1995), and of non-stationary oblivious equilibrium as suggested by Weintraub et al. (2008). First, each state is derived from the previous state and the transition matrix. Second, each driver maximizes his expected payoff by choosing a location when vacant. Finally, agents' expectations are rational, so there is no arbitrage. Then, the number of vacant taxis, the number of searching passengers, and the matching function are estimated. The estimates show that the New York market achieves about $4.2 million in daily surplus, about a third of which is realized as consumer surplus. Counterfactual analysis reveals that allowing prices to vary by time, location or distance can enhance allocative efficiency given the presence of search frictions, offering daily net surplus gains at least $194 thousand and 31,000 additional matches daily.

A challenging part of the modeling was the congestion specification. Mangrum and Molnar (2017), exploit the introduction of a new class of restricted taxi licenses in New York City to provide a causal estimate of the impact on congestion from the addition of taxis to the city. They document a large spike in taxi cab activity north of the restriction boundary, driven entirely by entry from the new restricted license taxis (green) and partially offset by retrenchment from traditional yellow taxis, which face additional localized competition.

Finally, Wong (2018) presents a dynamic spatial matching game model to study the effects of improved matching. He estimates the price and waiting time elasticities on the demand side to predict the response of net demand for taxicabs under different regimes. The key points of the model are:

- There are both street-hail and e-hail taxicabs
- E-hail cabs can pick up passengers found on the street
- Drivers act based on their beliefs about search time at different locations
- Time is discrete (20-minute intervals)
- In equilibrium, drivers maximize their expected utility and their beliefs are rational
- Demand is exogenous

Each street-hail driver has beliefs over the probability of getting matched with a passenger in each location. A vacant driver chooses which location to drive towards, taking into account the chance to be matched with a passenger on the way, the expected revenue of this match and the expected value of the trip's destination. In addition, the driver takes into account the probability of not matching on the way and the value of reaching his destination vacant. Based on the latter, the driver chooses the driving destination.

Each e-hail passenger takes into account the probability of matching electronically with a passenger, the expected revenue of this matching, the possibility of rejecting it, and the probability of returning late to his depot after accepting, which will result in a late fee. In case the driver is not matched electronically, then his problem reduces to the street-hail driver's problem.

The matching function's form is assumed to be Cobb-Douglas. The estimation method consists of two loops. The outer loop solves for the matching function's parameter and the inner loop solves for driver beliefs. The algorithm stops when both loops converge. Summarizing, this paper discusses search frictions in the taxicab market by using an agent-based model with respect to the drivers' decisions, as in my project. However, Wong attributes the frictions to a failure of drivers to find passengers and he claims that a centralized matching platform will solve the issue. My project discusses the possibility that drivers will choose the same actions even with a better matching mechanism because they end up stealing business from their competitors. In addition, demand is exogenous in Wong's paper while demand reacts to supply in my project. Finally, Wong discusses an environment that includes e-hail drivers, while my project does not allow different ways that a driver matches with a passenger.

3.3 Data

The main data source is the TLC's Taxicab Passenger Enhancements Project (TPEP), which creates an electronic record on every yellow cab trip. For each trip, it records a unique identifier for the driver and the cab license (medallion). It also regards the mile distance and time duration of the trip, the fare, the tip and any surcharges, and the geo-spatial start and endpoint of the trip. I use data from 2013, when the Uber presence was insignificant. Although the industry is different now, the analysis maintains external validity because business stealing, traffic and demand externalities still exist in the ridesharing market. The analysis focuses on Monday through Thursday, since activity on these days is almost identical whereas activity on weekends is characterized by peculiar features.

The calculation of each cab's location at any time from the data is straightforward. The model in subsequent sections describes the market in discrete time, so I organize the data in 20 minute intervals, and state variables will be assumed to change only when the time interval changes and not continuously. In terms of location, I partition New York City into the same 8 areas as in Frechette, Lizzeri and Salz (2018). The latter provide evidence that 8 areas describe the heterogeneity across areas almost as well as a finer partition with 16 areas. The number of matches between cabs and passengers can be straightforwardly computed from the data for each time interval and area.

I examine data on yellow taxi activity in December 2013 and provide a quantitative description of pickup and dropoff locations. More particularly, I examine matching data of the following dates: 2-5 December 2013 and 9-12 December 2013. Next, I calculate the mean values for each variable within these dates to smooth day-specific anomalies and create a dataset of an average weekday. For the purposes of this description I split New York City into 4 locations, as shown in Figure 3.1.

- Location 1: Lower Manhattan, south of W 39th street.
- Location 2: Midtown Manhattan, between W 59th street and W 40th street.
- Location 3: Upper Manhattan, between W 124th street and W 60th street.
- Location 4: The rest of the metro area.

Figure 3.1: Locations map for NYC Metro Area

Table 3.1 summarizes pickup and dropoff intensity over the designated locations. The difference in the percentages between pickups and dropoffs in location IV can be explained by the fact that in the boroughs outside Manhattan green cabs compete with yellow cabs in pickups. A person living in location IV has two ways to transport to Area I (green and yellow taxis) but only one option to return (yellow taxis).

Table 3.1: Spatial frequencies

	Location I	Location II	Location III	Location IV
Yellow pickups	42.4%	22.8%	26%	8.8%
Yellow dropoffs	38.5%	22.3%	25.6%	13.6%

December, 2013

Total number of pickups: 13,741,515

Total number of dropoffs: 13,733,541

3.4 Reduced Form Test

In this section, I develop a reduced-form test of the hypothesis that the taxicab market displays a Tragedy of the Commons. The main externality studied in this paper is the business stealing that occurs between taxi drivers. More specifically, a driver chooses actions based on the maximization of personal utility. When a driver acts to acquire an additional passenger, she fails to internalize that this passenger is taken from another driver. If business stealing exists in this industry, the Central Dispatcher could increase social surplus by internalizing the crossover effect between different drivers' actions. I present a statistic that takes different values depending on whether business stealing exists. I prove that under certain assumptions the market is characterized by business stealing. The result of the reduced-form test justifies measuring the extent of business stealing in a structural model.

In the literature, the Tragedy of the Commons is associated with excessive entry and business stealing, see Dixit and Stiglitz (1977), Mankiw and Winston (1986) as in Berry and Waldfogel (1999). The latter mention that the logic of free entry dictates that firms enter as long as the private benefit accruing to an entrant exceeds fixed costs. A Tragedy of the Commons would significantly reduce the social surplus in the taxi industry, given that entry costs are high while prices are fixed, so consumers do not benefit from increased entry. When the products are substitutes, the business stolen from incumbents creates a wedge between private and social benefits of entry. Berry and Waldfogel argue that the main determinant of whether there is business stealing is whether demand grows at sufficient rate as new firms enter the market. Do firms just split a fixed pie (business stealing) or do they add new customers (market expansion)? In the current context, taxis are perfect substitutes, but on the other hand there is a market expansion effect. When new cabs are added in a location, passenger waiting time is shorter and this attracts new customers. A concave matching function ensures that the latter effect is weaker than the direct new entry effect, which is linear. A simple example will demonstrate this point.

Consider 2 locations. There is a number T of taxis to be allocated. In the first location, the matching function is $m(N) = \sqrt{N}$, meaning that if there are N taxis in this area, then $\sqrt{N}$ taxi-passenger matches occur. It is obvious that the total matches increase as new taxis enter, but in a concave way. In the second location, the matching function is $m(N) = log(N)$, which is also concave. In order to keep the example simple, one can assume that all passengers are identical, so a driver cares only about his/her own number of matches. The results do not change in absence of this assumption. The competitive equilibrium ensures an equal number of matches per taxi in both areas so that drivers are indifferent between locations, i.e.:

$\frac{\sqrt{N_1}}{N_1} = \frac{log(T-N_1)}{T-N_1} \Rightarrow$

$$T - N_1^{comp} = \sqrt{N_1^{comp}} \cdot log(T - N_1^{comp})$$

The socially optimal solution, meanwhile, maximizes the total number of matches:

$$max_{N_1}(\sqrt{N_1} + log(T - N_1)) \Rightarrow$$

$$T - N_1^{opt} = \sqrt{N_1^{opt}} \cdot 2$$

By comparing the solutions, one can see that $N_1^{comp} < N_1^{opt}$ (and $N_2^{comp} > N_2^{opt}$). The intuition is the following: Free markets with price controls have too much crowding of sellers where demand is highest, so the outcomes depends heavily on the curvature of demand. In the competitive equilibrium each driver maximizes his own benefit, which increases in a concave way with respect to N in each location, but declines with respect to $\frac{1}{N}$. Given that the matching functions are concave, the latter effect is stronger, and the effects cancel out each other only if the matching functions are linear. This fact pushes for a more uniform distribution of taxis among locations than is optimal. The social benefit declines with respect to $\frac{1}{2\sqrt{N_1}}$ in location 1 and $\frac{1}{N_2}$ in location 2, which dictates the placement of more taxis in location 1 than location 2, in contradiction with the competitive equilibrium.

For the purposes of my test, I must establish assumptions under which a concave matching function over the total area implies concave matching functions over all locations. Consider I different locations and an initial total number of taxis N_1, an intermediate number of taxis $N_2 > N_1$, and a final total number of taxis $N_3 > N_2$. Define as m_1 the total number of matches when there are N_1 taxis, and m_2, m_3 respectively.

Assumption 4.1: A matching function's domain is $\mathbb{R}$, even though the number of cabs belongs to $\mathbb{N}$

Assumption 4.2: A matching function can be either globally concave, linear, or globally convex.

Define α_1, α_2 and γ_1, γ_2 as: $N_2 = \alpha_1 \cdot N_1$, $N_3 = \alpha_2 \cdot N_2$ and $m_2 = \gamma_1 \cdot m_1$ and $m_3 = \gamma_2 \cdot m_2$. Note that by definition, $\alpha_1, \alpha_2 > 1$. Then, for each location i denote the number of taxis in this location as N_{1i}, N_{2i}, N_{3i}. Then similarly $N_{2i} = \alpha_{1i} \cdot N_{1i}$, $N_{3i} = \alpha_{2i} \cdot N_{2i}$ and $m_{2i} = \gamma_{1i} \cdot m_{1i}$ and $m_{3i} = \gamma_{2i} \cdot m_{2i}$.

The total matching function is concave if and only if $\frac{\gamma_1}{\alpha_1} > \frac{\gamma_2}{\alpha_2}$. First, I prove that in this case, there must be at least one location with a concave matching function. Then, I show that if one location has a concave matching function, then all locations must have concave matching functions. So:

$$\alpha_1 \cdot N_1 = N_2 = \sum_i N_{2i} = \sum_i \alpha_{1i} \cdot N_{1i} \Rightarrow$$

$$\alpha_1 = \frac{\sum_i \alpha_{1i} \cdot N_{1i}}{\sum_i N_{1i}}$$

Thus, α_1 is a weighted average of the α_{1i}'s. The same holds for $\alpha_2, \gamma_1, \gamma_2$. Suppose that no location has a concave matching function, i.e. $\frac{\gamma_{1i}}{\alpha_{1i}} \leq \frac{\gamma_{2i}}{\alpha_{2i}} \quad \forall i$. By the weighted average property, it follows that $\frac{\gamma_1}{\alpha_1} \leq \frac{\gamma_2}{\alpha_2}$, which contradicts the concavity property of the total matching function. Therefore, there must be at least one location with a concave matching function.

Given that there is at least one location with a concave matching function, I next prove that under some additional assumptions, all other locations must have concave matching functions as well.

Assumption 4.3: The expected revenue from a passenger, given that the passenger is picked up by a driver and given the location, does not depend on the number of taxis or passengers in the location.

Assumption 4.4: Drivers choose locations maximizing their expected benefit.

Assumption 4.5: The matching functions are continuous and differentiable on their domains.

In order to simplify notation, I work with two locations. The proof is similar for any number of locations. If N_{11}, N_{21}, N_{31} are the number of taxis in the first location, then $N_1 - N_{11}, N_2 - N_{21}, N_3 - N_{31}$ are the taxis in the second location. Suppose that the first location has a concave matching function, f. Assume that the matching function in the second area is linear, $\gamma \cdot N_2$. Then, by assumptions 4.3 and 4.4 we have:

$$\frac{f(N_{11})}{N_{11}} \cdot c_1 = \frac{\gamma \cdot (N_1 - N_{11})}{(N_1 - N_{11})} \cdot c_2$$

$$\frac{f(N_{21})}{N_{21}} \cdot c_1 = \frac{\gamma \cdot (N_2 - N_{21})}{(N_2 - N_{21})} \cdot c_2$$

$$\frac{f(N_{31})}{N_{31}} \cdot c_1 = \frac{\gamma \cdot (N_3 - N_{31})}{(N_3 - N_{31})} \cdot c_2$$

where c_1, c_2 are the expected revenues from picking up a passenger from locations 1 and 2 respectively. The latter equations are equivalent to:

$$\frac{f(N_{11})}{N_{11}} = \frac{f(N_{21})}{N_{21}} = \frac{f(N_{31})}{N_{31}} = \gamma \cdot \frac{c_2}{c_1} \coloneqq k$$

Then, define β_1, β_2 as: $N_{31} = N_{21} + \beta_2$, $N_{21} = N_{11} + \beta_1$. Then,

$$\frac{f(N_{11}+\beta_1+\beta_2)}{N_{11}+\beta_1+\beta_2} = \frac{f(N_{11}+\beta_1)}{N_{11}+\beta_1} = \frac{f(N_{11})}{N_{11}} = k \Rightarrow$$

- $f(N_{11}) = k \cdot N_{11}$
- $f(N_{11} + \beta_1) = k \cdot (N_{11} + \beta_1) = k \cdot N_{11} + k \cdot \beta_1 = f(N_{11}) + k \cdot \beta_1$
- $f(N_{11} + \beta_1 + \beta_2) = k \cdot (N_{11} + \beta_1 + \beta_2) = f(N_{21}) + k \cdot \beta_2$

Thus,

$$\frac{f(N_{11}+\beta_1)-f(N_{11})}{\beta_1} = k$$

$$\frac{f(N_{11}+\beta_1+\beta_2)-f(N_{11}+\beta_1)}{\beta_2} = k$$

By the Mean Value Theorem, $\exists x_1 \in (N_{11}, N_{11} + \beta_1)$ and $\exists x_2 \in (N_{11} + \beta_1, N_{11} + \beta_1 + \beta_2)$ such that $f'(x_1) = f'(x_2) = k$, which is a contradiction, as from the concavity of f it must be that f' is strictly decreasing. This way, linearity of the matching function in location 2 is rejected.

Convexity of the matching function in location 2 leads also to contradiction, as taxis are not concentrated in one location. Therefore, given one location with a concave matching function, all other locations must have concave matching functions under assumptions 4.3-4.5. Under assumptions 4.1-4.5, therefore, there must be a concave matching function in all locations if the aggregate matching function is concave. Summarizing, if a test applied to aggregate data implies a concave total matching function, then under the aforementioned assumptions this implies that the matching functions are concave in all areas, which implies a business stealing effect.

The variation needed for the test comes from the addition of new vehicles over time. I include Ridehailing app vehicles in my test. The test covers the time interval from Jan. 2015 to Jun. 2017, when Ridehailing app data is available. I choose this time period instead of the period used for the model's estimation because the shock of ridesharing becoming available provides the required variance to test the hypothesis. A graphical representation of the number of matches and vehicles can be found in Figures 3.2 and 3.3 respectively. I show data from 2010-2019, instead of 2015-2017 to present the long trend of ridesharing entering.

At a glance, one can see that the number of vehicles increased threefold while trips increased less than twofold. In order to construct an intuitive test for concavity, I regress the number of vehicles and its square on the number of trips per vehicle. The choice of the independent variable is valid, as the number of vehicles increases monotonically over time. If trips per driver decline as the total number of vehicles rises, then this indicates a business stealing effect. Figure 3.4 reveals an obvious downtrend in trips per vehicle with respect to the number of vehicles.

The regression is performed as follows:

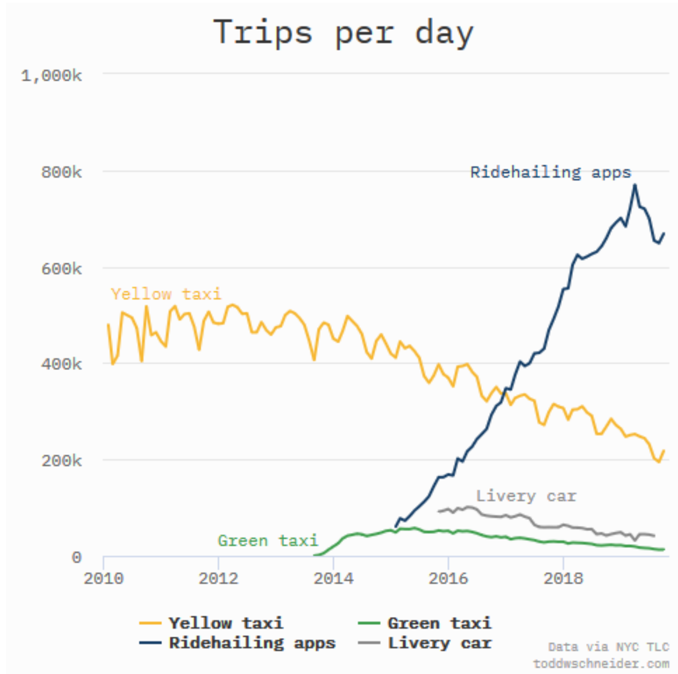

Figure 3.2: Average trips per day

$$TripsPerVehicle_t = \beta_0 + \beta_1 \cdot Vehicles_t + \epsilon_t$$

After deriving an estimator for β_0, β_1, I perform the following test:

$H_0 : \quad \beta_1 = 0$, (linearity)

$H_a : \quad \beta_1 < 0$, (concavity)

The regression results are presented in Table 3.2.

The null hypothesis is rejected at a 99% confidence level, providing strong evidence in favor of concavity. The fact that the R^2 is more than 0.76 shows that only a small part of the variation is not explained by the reduced form model, adding credibility to the test.

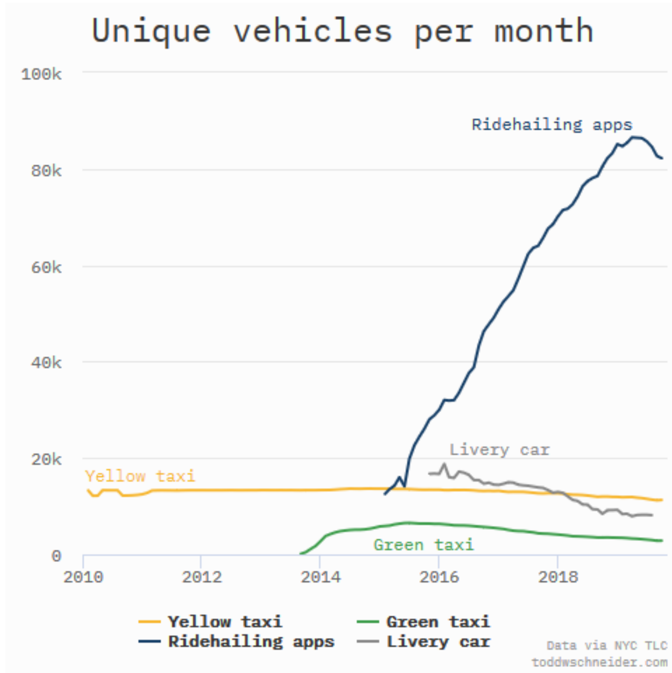

Figure 3.3: Average vehicles per month

It is useful to rule out that the reduction of trips per yellow taxi is not a consequence of their facing more efficient competitors (i.e. lower price). The data suggests that New York is the only city where taxis are slightly cheaper, than Uber cars, even with a 20% tip for the taxi driver included. The only case where Uber is cheaper is in zero traffic conditions where the speed of the car is more than 30 miles/hour, an unrealistically large number for New York City. Figure 3.5

Table 3.2: Concavity Test table

	TripsPerVehicle
β_1(**)	-13e-05
	(11.3e-6)
β_0(**)	34.86
	(0.75)
Observations	42
R^2	0.7635

standard errors in parentheses
(**): significant at 99% confidence level

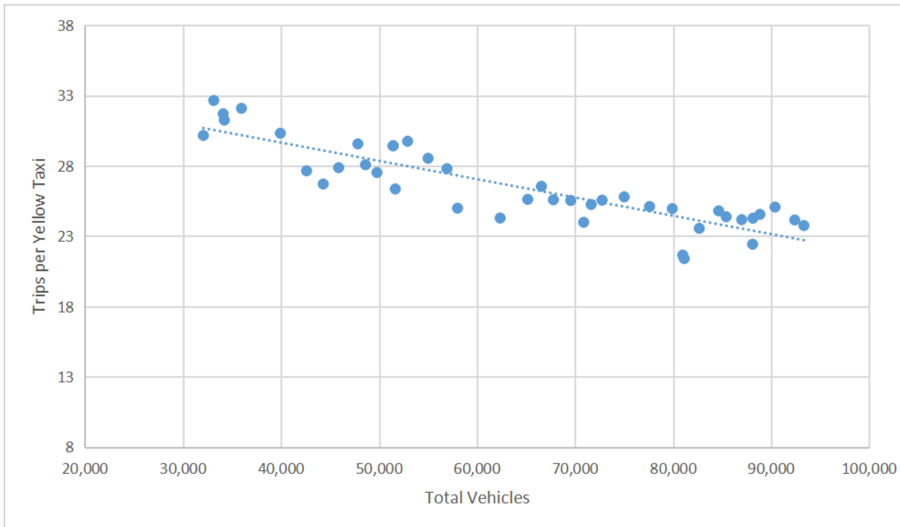

Figure 3.4: Average trips per yellow Taxi

summarizes these findings.

Taking all the aforementioned evidence into account, one can conclude that the matching function is concave with high confidence. Under the assumptions mentioned before, we can conclude that the matching function is concave in all locations,suggesting business stealing effect and a Tragedy of the Commons.

3.5 Model

The structural model extends the search and matching model developed by Buchholz (2020). Time and space are discretized. Drivers make spatial decisions among L locations over T periods in a day shift. The city is designed as a graph, where the vertices are the locations in the city and the edges are the routes from one location to the other. Consider L distinct locations and TC utility maximizing agents (taxis). Time is discrete and finite, with T periods. Given the distribution of taxis over locations, passengers make decisions on how to commute. Specifically, given the distribution of vacant taxis in every location at time t, N_{it} $\forall i \in 1, ...L$, and the traffic

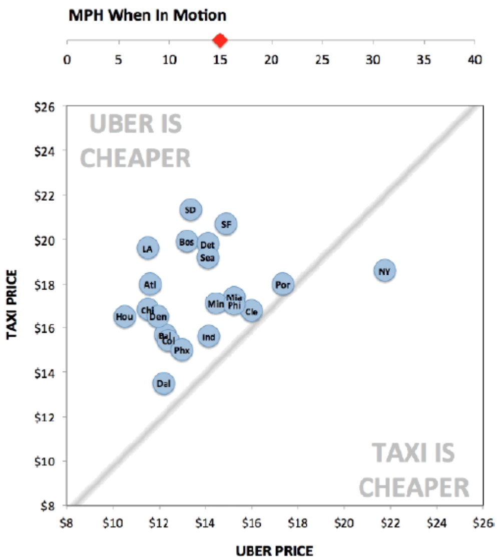

(a) 15 MPH

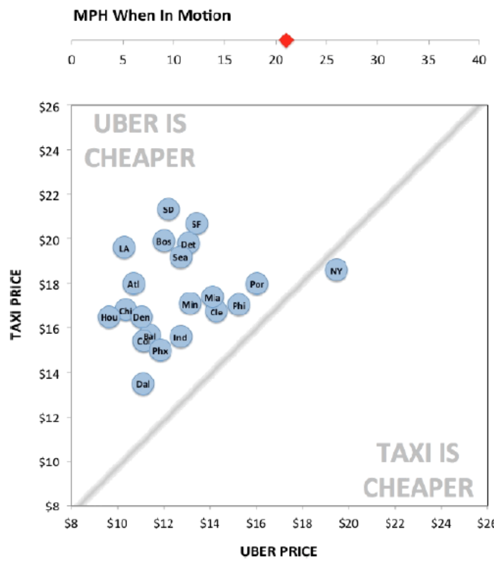

(b) 21 MPH

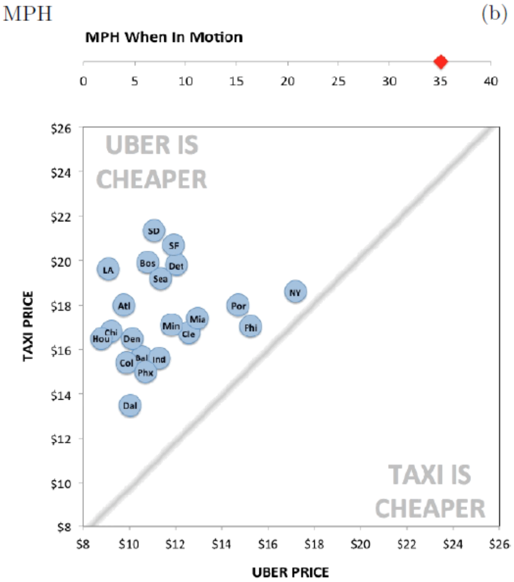

(c) 35 MPH

Figure 3.5: Taxi-Uber price comparison, 20% tip for taxi included

speed in every edge, tr_{ijt} $\forall i, j \in 1, ...L$, passengers enter location i waiting to commute to location j, $cons_{ijt}$. Taxis and passengers are matched in each location according to a matching function $f(N_{it}, cons_{ijt})$ and m_{ijt} rides occur from location i to location j. A taxi matched with a passenger that wants to commute to location j is full for the duration of their trip τ_{ij}, and receives a payment p_{ij}. After reaching their destination, the taxi drops the passenger off and becomes vacant again. Vacant taxis that failed to match with passengers make decisions on their location to look for passengers in future periods. Commuting from location i to location j, incurs a cost, $cost_{ijt}$, for either a vacant or full taxi. After matches are determined, the traffic speed on each edge is updated.

3.5.1 Congestion

Traffic speed is modeled to depend on the number of taxis traveling along a particular route, as well as, time and location characteristics. The effect of non-taxi vehicles is expressed by route and time of day fixed effects, which capture rush hour traffic and the routes that non-taxi vehicles take to commute. The latter implies the assumption that non-commercial vehicle drivers' decisions are independent of taxi drivers' decisions. More specifically, I assume that the congestion created by non-commercial vehicles for a given location and time of weekday is exogenous. Up to a certain number of vehicles, traffic speed is determined by exogenous factors, e.g. traffic lights, weather, etc. After a critical point of vehicles, there is not enough time for all of them to cross traffic lights, congestion occurs, and traffic speed reduces to the speed of the queues at traffic lights. Wang et al. (2016), suggest that there is a peak density that a roadway can sustain at an uncongested state, and if this density is surpassed then the roadway will fall into a congested traffic state. This density is known as the critical density, or KC. I use a sigmoid function to express the traffic speed as a function of the number of vehicles in a location to capture the critical density effect. The convexity of a sigmoid function at its lower levels provides the required slow increase at a low number of vehicles, until a critical point where the function becomes concave. The critical point captures the aforementioned critical density effect. I choose the hyperbolic tangent function to express these dynamics. Traffic speed is calculated as follows:

$$tr_{ijt} = tanh(d_t + c_{ij} \cdot Ns_{ijt}) \quad (I)$$

where d, c are time and location fixed effects respectively, and Ns_{ijt} is the number of taxis traveling from location i to location j at time t, either full or vacant. I assume that congestion on the route from location i to location j is affected only by taxis commuting from i to j. This assumption ignores the fact that different routes may overlap. However, this assumption gives a reasonably simple function, while the congestion model's predictive accuracy is high (see Estimation section).

3.5.2 Cost and Revenue

Price per mile is regulated in the New York City taxi market. Therefore, the price paid by the consumer for commuting through an edge,P_{ij}, depends only on the length of the edge. In the price calculation I have included the $1 rush hour surcharge when applicable, and the New York State Congestion Surcharge of $2.50 for all trips that begin, end or pass through Manhattan south of 96th Street. The only part of the Metered Fare that I omit is that when the vehicle travels below 12 mph, per mile pricing stops and the charge is 50 cents/minute, as if the vehicle travels at 12 mph. Drivers pay fuel costs for each trip, regardless of being full or vacant. The cost of traveling through edge ij at time t is given by:

$$cost_{ijt} = \frac{cost\_coef \cdot g \cdot c(tr_{ijt}) \cdot distance(i,j)}{tr_{ijt}} \quad (II)$$

where:

- g: The price of gas/gallon measured in USD.
- $c(tr_{ijt})$: The mean gas consumption of the taxi, measured in gallons/hour, as a function of speed
- $distance_{ij}$: The distance between locations i, j measured in miles
- tr_{ijt}: The traffic speed on edge ij at time t, measured in miles/hour
- $cost\_coef$: a numerical coefficient used for adjustment

I calculate the function c using data from *fueleconomy.gov*. Data indicates that up to 30 mph, miles per gallon are sufficiently described by the equation $mpg = speed(mph) + 5$, therefore $c(tr_{ijt}) = \frac{tr_{ijt}}{tr_{ijt}+5}$ gallons/hour. Hence, the driver's profit for a ride from location i to location j at time t is:

$$\Pi_{ijt} = P_{ij} - cost_{ijt} \quad (III)$$

3.5.3 Matching

At the start of each period t, N_{it} vacant taxis in location i search for passengers. Given the number of vacant taxis in a location, traffic conditions, prices and time, a number of passengers appear in location i wanting to commute to location j, $cons_{ijt}$, $\forall j \in 1, ..., L$. The number of passengers in location i who want to commute to location j that manage to match with a taxi at time t is $m_{ijt} = f(N_{it}, cons_{ijt})$. Therefore, the total number of matches that occur in location i at time t is $m_{it} = \sum_j m_{ijt}$. Note that when a vacant taxi meets a passenger, it is obliged by law to serve the passenger and has no right to refuse service. I use the same matching functional form as Buchholz, with the difference that matching efficiency depends on traffic conditions. The reason is that when the traffic is heavy, taxis are able to search in a smaller fraction of a location's total area and therefore encounter fewer passengers.

Buchholz notes that simulating every intersection in New York City to obtain the number of matches would create a computational burden. Instead, Buchholz proposes a functional form that is derived from an urn-ball matching problem, where N balls are randomly placed in $cons$ urns, and a match occurs only for the first ball placed in any urn. The latter implies the following assumptions:

- An area is represented as a collection of points (urns)
- All taxis are identical and have identical searching strategies
- Passengers are uniformly distributed over points
- Passengers do not move while looking for a taxi

Given that the number of passengers is a Poisson random variable, the expected number of matches traveling from location i to location j, given the distribution parameters and matching efficiency, is:

$$m_{ijt} = E[N_{it} \cdot (1 - (1 - \frac{1}{loc_{ijt} \cdot N_{it}})^{cons_{ijt}}) | \lambda_{ijt}] = N_{it} \cdot (1 - \exp(-\frac{cons_{ijt}}{loc_{ijt} \cdot N_{it}})) \quad (IV)$$

where loc_{ijt} is the inverse matching efficiency in location i for a trip to location j at time t and it is given by:

$$loc_{ijt} = \sqrt{\frac{a_i}{tr_{ijt}}} \quad (V)$$

where a_i is a location parameter. Given the assumption of identical taxis, the probability that a taxi in location i matches with a passenger going to location j is:

$$p_{ijt} = \frac{m_{ijt}}{N_{it}} \quad (VI)$$

and thus the probability that a taxi is matched with a passenger in location i is $p_{it} = \sum_j p_{ijt}$.

3.5.4 Passengers

In this model, demand for trips from location i to location j at time t is a random variable that follows a Poisson distribution. The parameter of the distribution is determined endogenously as follows:

$$\lambda_{ijt} = \alpha_t^{time} \cdot (\alpha_i^{location} + \alpha^{price} \cdot log(P_{ij}) + \alpha_i^{supply} \cdot N_{it} + \alpha_i^{speed} \cdot tr_{ijt}) \quad (VII)$$

where P_{ij} is the price of a trip from location i to location j, N_{it} is the number of vacant taxis in location i at time t, and tr_{ijt} is the traffic speed. The parameters $\alpha_t^{time}, \alpha_i^{location}, \alpha^{price}, \alpha^{supply}, \alpha^{speed}$ capture the effect of time, location, price, number of taxis and traffic speed. Their dimensions are $1xT$, $1xL$, $1x1$, $1x1$ and $1x1$ respectively. Therefore the number of passengers that appear in location i at time t and want to commute to location j is:

$$cons_{ijt} \sim Poisson(\lambda_{ijt}) \quad (VIII)$$

Equation (VII) implies that prices, taxi supply, and traffic speed are known to potential passengers.

3.5.5 Drivers' Decisions

A driver observes the state of the market, and if vacant makes a decision on where to commute in order to search for passengers. A driver at time t can be in one of the following situations:

1. Having a passenger and transiting to a location
2. Being vacant and transiting to a location
3. Being vacant and looking for passengers in the area where they are located

Drivers in situations 2 and 3 can match with passengers in the area where they are located. Drivers in situation 2 that fail to match with a passenger continue their trip to their destination. Drivers in situation 3 that fail to match with a passenger decide where to search for passengers in the future. Given the state of the market, i.e. the distribution of taxis over locations at time t and their expectations about the number of passengers λ_{ijt}, drivers in situation 3 choose to commute to the location with the maximum expected value. As in Buchholz, the expected value of location i at time t is:

$$V_{it} = E[\textstyle\sum_j (p_{ijt} \cdot (\Pi_{ijt} + \delta^{\tau_{ij}} \cdot V_{j(t+\tau_{ij})})) + (1 - \sum_j (p_{ijt})) \cdot E_{\epsilon_{j,a}}[max_j \{V_{t+\tau_{ij}} - cost_{ijt} + \epsilon_{j,a}\}]] \quad (IX)$$

where:

- p_{ijt}: The probability that a vacant taxi at location i matches with a passenger going to location j at time t, described in equation (VI).
- Π_{ijt}: The profit of a driver serving a passenger commuting from i to j at time t, described in equation (III).

- δ: Driver's discount factor $\in (0, 1)$.

- τ_{ij}: The number of time periods needed to commute from location i to location j.

- $cost_{ijt}$: The gas cost of commuting from location i to location j at time t, described in equation (II).

- $\epsilon_{j,a}$: Driver a-specific i.i.d. shock, drawn from a Type-I extreme value distribution.

Hence, when unmatched and called to make a decision, an unmatched driver chooses to commute to:

$$j^{\star} = argmax_j \{V_{t+\tau_{ij}} - cost_{ijt} + \epsilon_{j,a}\} \quad (X)$$

Therefore, given the distribution of the shocks, the probability of a driver choosing to commute from location i to location j at time t is:

$$P_i[j|V_t] = \frac{exp(E_{V_{j,t+\tau_{ij}}}[V_{j,t+\tau_{ij}} - cost_{ijt}]/\sigma_\epsilon)}{\sum_k exp(E_{V_{k,t+\tau_{ik}}}[V_{k,t+\tau_{ik}} - cost_{ikt}]/\sigma_\epsilon)} \quad (XI)$$

The processes of matching and driver decisions determine the transition process, which functions as follows: At time t, there are vacant taxis that reside in location i and search for passengers, EL_{it}; there are vacant taxis that are commuting to another location and transit through location i, ET_{it}; and there are full taxis that travel through or to location i, FT_{it}. The taxis that can match with a passenger are those belonging in the sets EL_{it} and ET_{it}. Those who match with a passenger at time t will belong to the set $FT_{i,t+1}$ in the next period, as they will have a passenger and start their trip to their destination. Those taxis belonging to EL_{it} that fail to match decide where to look for a passenger next period. Those who choose to stay in i will belong to $EL_{i,t+1}$. Those who choose to commute to another location j will belong to $ET_{k,t+1}$, where k is the next station of their trip. The taxis that belong to ET_{it} will belong to $ET_{k,t+1}$, where k is the next station of their trip, or to $EL_{i,t+1}$, if i was the destination of their trip. The full taxis, FT_{it}, will belong to $FT_{k,t+1}$, if k is the next station of the passenger's travel, or to $EL_{i,t+1}$ if i was the destination of the passenger's trip.

The timing of events within a period, as in Bian (2019), is presented in Figure 3.6: t_0 : Drivers

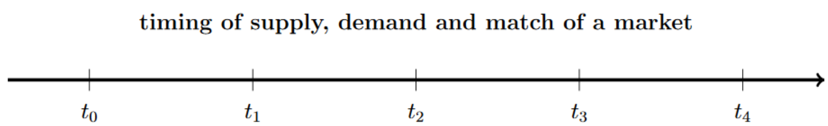

Figure 3.6: Timing

arrive and become supply. t_1 : Passengers arrive given drivers' decisions. t_2 : Drivers and passengers are matched. t_3 : Unmatched passengers exit the market and take the subway (or walk). t_4 : Employed drivers deliver passengers and unemployed drivers make search decisions for the next period.

I assume that each driver ignores the impact of his decisions on congestion, matching efficiency and demand, as they think themselves too small compared to the whole market.

3.6 Estimation Results

This section discusses the estimation results of the spatial model. Buchholz explains in detail the identification of the parameters used in the original model. The data suffice for the identification of the features I add in order to create the enriched model. In particular, I account for endogeneity issues for the effect of congestion on demand and matching efficiency, and the effect of taxi supply on demand. First, meteorological conditions create the required variance in congestion while they are assumed to have no direct effect on demand and the matching process. For instance, a snowy day slows the traffic even if the number of cars on the road is the same, as drivers reduce their speed in bad weather, while demand and matching efficiency are not directly affected. Passengers are expected to commute to their workplaces early in the morning despite the weather conditions and drivers do not face additional difficulties in seeing hailing passengers due to the weather conditions. Another instrument I use to identify taxi supply is the variation over time in gas prices. Increased gas prices discourage drivers from driving while they do not directly affect demand or the matching process. The effect of taxi supply on demand can be identified by the variance in the number of matches created by the different numbers of active drivers across days and time of day. I calculate the number of active drivers by counting the different Driver IDs in the matching data, given the time interval. I assume that if a driver has not matched for two consecutive hours, then the driver is not active at this time interval. The estimated parameters of the model are the following:

- Demand coefficients
- Matching coefficients
- Cost coefficients
- Congestion coefficients

The aforementioned parameters are sufficient to calculate all relevant market outcomes, i.e. the number of matches, the number of taxis, demand, driving costs and congestion in each location, according to the equations provided in Section 5. Note that the prices for each trip are determined

Table 3.3: Estimation results: Demand parameters

	Lower Manhattan	Mid Manhattan	Upper Manhattan	Boroughs
$\alpha^{location}$	17.17 (0.95)	10.75 (0.05)	10.94 (0.22)	7.71 (0.32)
α^{supply}	4.95 (0.005)	5.04 (0.06)	4.99 (0.02)	5.05 (0.1)
α^{speed}	5.49 (0.45)	5.73 (0.04)	4.54 (0.12)	3.63 (0.03)
α^{price}	0.91 (0.02)			
	Morning	Early Afternoon	Late Afternoon	
α^{time}	0.42 (0.0001)	1.30 (0.1)	1.63 (0.01)	

by TLC and are not subject to change, and I treat them as constants in the estimation. Demand coefficients are comprised of time effects, location fixed effects, a price parameter[2], a supply parameter and a parameter related to traffic speed. The parameters for supply and traffic speed are allowed to have different values for each location while the price coefficient is assumed not to vary among different locations. Data suggests that time effects can be grouped into three categories: morning ($1 \leq t \leq 11$), early afternoon ($12 \leq t \leq 22$) and late afternoon ($23 \leq t \leq 33$). There appears to be little variance in time effects within these groups, so for parsimony I model time effects to be $1x3$ dimensional. The demand parameters are plugged into equation (VII) to calculate the expected number of passenger arrivals in each time interval at each location. Table 3.3 reports the estimated values of the demand parameters, with standard errors in parentheses.

The matching parameters reflect the effect of congestion on the matching efficiency of each location. I define matching efficiency in the same way as Buchholz. The estimation results on matching efficiency are reported in Table 3.4, and suggest that matching efficiency is affected most by congestion in Lower Manhattan, i.e. matching efficiency decreases faster in Lower Manhattan than in other locations as congestion increases. The cost coefficient is one-dimensional

[2] I identify the demand elasticity with respect to price as in Buchholz (2020) by re-estimating the model for September, 2012, following a change in the regulated tariff by TLC.

Table 3.4: Estimation results: Matching/Congestion efficiency parameters

	Lower Manhattan	Mid Manhattan	Upper Manhattan	Boroughs
$\alpha^{matching}$	0.52 (0.003)	0.19 (0.002)	0.35 (0.001)	0.31 (0.01)

Table 3.5: Estimation results: Traffic speed parameters

	Lower Manhattan	Mid Manhattan	Upper Manhattan	Boroughs
c ($x10^{-4}$)	-8.64 (0.65)	-8.1 (0.57)	-11 (0.77)	-0.05 (0.098)
	$t = 8$	$t = 16$	$t = 24$	$t = 32$
d ($x10^{-3}$)	8.4 (2.16)	7.5 (2.29)	6.8 (2.29)	6.2 (2.11)

and captures the bias of the cost formula presented in equation (II) based on the fuel cost of driving. The aforementioned bias comes from the fact that the exact way that fuel consumption depends on traffic speed is not known. The cost coefficient is found to be:

$$cost\_coef = 1.3286 \quad (0.02)$$

Finally, the congestion coefficients describe the way that traffic speed depends on time effects, location effects, and the number of taxis traveling between any pair of locations, as presented in equation (I). I estimate these coefficients offline after the estimation of the rest of the parameters. This is possible because traffic speed can be inferred directly from the data using the location coordinates and the timestamps, and then fitted to the number of taxis estimated to travel between any pair of locations. The association of each taxi trip with the corresponding spatial areas I define is achieved through the point-in-polygon matching procedure outlined in Brophy (2013). Table 3.5 reports the estimation results, presenting a subset of the time parameters because of their large number.

Given the estimated demand parameters, I calculate the consumer and producer surplus as in Buchholz, by integrating under the demand function. I keep Buchholz's assumption that while there is no perfect substitute for taxis, even the highest-value consumers have limits to their willingness to pay. In order to avoid extraordinarily large consumer valuations, I use a maximum willingness to pay of $100 for a trip. I present the realized demand, supply, matches and traffic speed in Section 8 in order to be easily comparable with the counterfactual results.

3.7 Central Dispatcher

I now address is about the benefit of a Central Dispatcher on the social surplus of the market. In this paper's context, a Central Dispatcher is defined as an agent who at time t has knowledge of:

- All the model's parameters
- The location and direction of all taxis up to time t
- The condition of each taxi (vacant or full) up to time t

The Central Dispatcher has rational expectations on:

- The number of passenger that will arrive at time t
- The number of matches that will take place at time t for each allocation of taxis

An analytical solution to the Central Dispatcher's problem is unfeasible due to the exponential increase of market states over time. I use two separate algorithms to obtain approximate solutions to the first best. An approximate solution can serve as a lower bound of the Central Dispatcher's ability to increase the social surplus, as it considers a subset of the available policies. The first approach is a greedy algorithm, where the Central Dispatcher optimizes only with respect to the next period. This is clearly suboptimal, as the optimal path may require some unattractive short term decisions, which are disregarded in this paradigm. Figure 3.7 provides an illustrative example. While the shortest path from node A to node F is ABDF, with a total length of $4+1+11 = 16$, a greedy algorithm chooses to transit to C when on node A as 2 is smaller than 4, while ignoring the rest of the graph. The path selected by the greedy algorithm is ACEDF, with a total length of $2+3+4+11 = 20$, which is longer than the optimal path. The second algorithm performs an intertemporal optimization, but it uses an approximate version of the objective function in order to reduce the dimensionality of the state space.

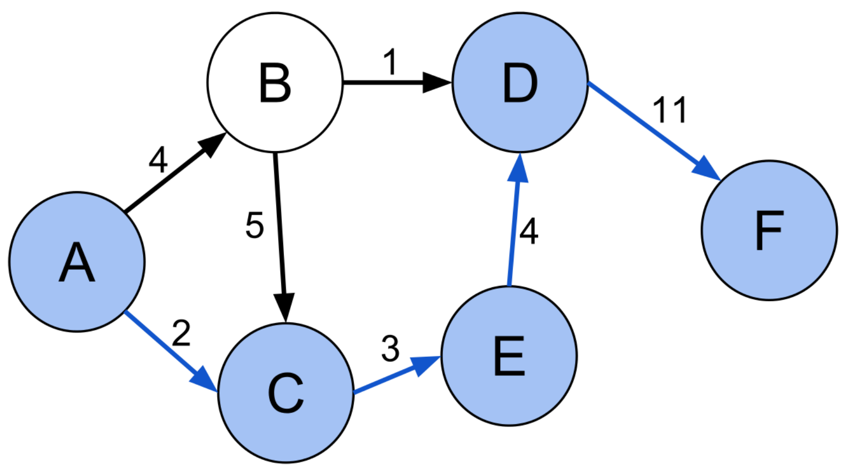

Figure 3.7: Optimal path

3.7.1 Algorithm 1: Greedy optimization

The Central Dispatcher makes decisions on behalf of the drivers to maximize social surplus. Specifically, in the benchmark decentralized model, a searching driver who fails to match in a particular location in this period then chooses the next search location. In the counterfactual, this decision is made by the Central Dispatcher. The drivers have to comply with the Central Dispatcher's instructions regardless of their shocks and personal interest. The Central Dispatcher modifies the transition dynamics as follows (XII):

1. $EL_{it} \Rightarrow$

    - Matched, added to FT_{t+1}
    - Not matched, added either to $EL_{i(t+1)}$ or $ET_{k(t+1)}$ for $k \neq i$, *decided by the Central Dispatcher*

2. $ET_{it} \Rightarrow$

    - Matched, added to FT_{t+1}
    - Not matched, continues to next location k automatically, added to $ET_{k(t+1)}$ or to $EL_{k(t+1)}$, if k is the last station of their trip

3. $FT_{it} \Rightarrow$

    - Continues to the next station k, added to $FT_{k(t+1)}$
    - Reaches destination k, drops off the passenger, added to $EL_{k(t+1)}$

The Central Dispatcher chooses drivers' search locations to maximize social surplus defined as the sum of producer and consumer surpluses, $TW = PS + CS$. The social surplus from time t onward given a taxi allocation S_t and a decision by the Central Dispatcher x_t is:[3]

$$TW_{t...,S_t,x_t} = PS_{t...,S_t,x_t} + CS_{t...,S_t,x_t} =$$

$$PS_{t,S_t,x_t} + PS_{t+1...,S_{t+1},x_{t+1}|S_t,x_t} + CS_{t,S_t,x_t} + CS_{t+1...,S_{t+1},x_{t+1}|S_t,x_t} \quad (XIII)$$

Equation $(XIII)$ is expressed as a recursion in order to facilitate its solution through backwards induction. In other words, given all allocations after period t: $S_t, S_{t+1}, ..., S_T$, I solve for S_t and

[3]The difference between CS_{t,S_t,x_t} and $CS_{t...,S_t,x_t}$ is that the former is the producer surplus time t only, while the latter is the consumer surplus from time t onward until T.

x_t to maximize welfare at time t, as the allocations after t are already determined, and S_t is calculated by the transition rules.

The producer surplus added at time t is:

$$PS_t = \sum_{i,j}(m_{ij(t+1)} \cdot (P_{ij} - cost_{ij(t+1)}) - x_{ijt} \cdot cost_{ijt}) \quad (XIV)$$

where the producer surplus is the sum of the total profit that comes from matches minus the cost of the Central Dispatcher's decisions, x_{ijt}, which determine the number of taxis that relocate from location i to location j at time t. The consumer surplus added at time t is:

$$CS_t = \sum_{i,j} \int_{P_{ij}}^{P_{max}} D_{ij(t+1)}(p)dp \quad (XV)$$

where D_{ijt} is the expected consumer demand specified by equation (VII), and P_{max} is the choke price as defined by Buchholz. The Central Dispatcher decides on x_{ijt} at time t with the constraint that the number of decisions in a location equals the number of unmatched vacant taxis that are not transiting through the location, i.e.:

$$\sum_j x_{ijt} = \sum_j (1 - p_{ijt}) \cdot EL_{it} \quad (XVI)$$

In order to achieve an analytical solution, I approximate S_t using a continuous extrapolation. Without this, the differentiation of the objective function would be impossible, as the number of taxis in each state can only be an integer number. In other words, the objective function in the greedy case is the sum of the current flow producer and consumer surpluses ($PS_t + CS_t$) instead of the intertemporal flow ($E[\sum_{t=1}^{T} \delta^t \cdot (PS_t + CS_t)]$), which would give the optimal solution. Hence, the Central Dispatcher solves for every time t, given the state at $t+1$:

$$\frac{d(PS_t(S_t|S_{t+1})+CS_t(S_t|S_{t+1}))}{dx_{ijt}} = 0 \quad \forall i \neq j \quad (XVII)$$

$$x_{iit} = \sum_j (1-p_{ijt}) \cdot EL_{it} - \sum_{j\neq i} x_{ijt} \forall i \quad (XVIII)$$

$$S_{t+1} = transit(S_t, x_t) \quad (XII)$$

Equations $(XVII), (XVIII), (XII)$ solve for S_t and $x_{ijt} \quad \forall i, j$. More specifically, the greedy algorithm used by the Central Dispatcher to obtain the optimal policy is:

This scheme provides a suboptimal allocation over time because of the planner's short sightedness.

Guess an allocation for the final time period, S_T, as in Buchholz;
for $t = T - 1$ *to* 1 **do**
  Observe demand for time t;
  Solve for S_t, x_t given S_{t+1} using $XVII, XVIII, XII$ and $I - VII, XIV, XV$;
  Update t, S_t;
end

Algorithm 2: Greedy Algorithm

It can be used as a lower bound of the improvement caused by the addition of the Central Dispatcher, in case it increases the social surplus compared with the competitive allocation. Note that the Central Dispatcher is able to internalize the externalities of business stealing, as the Central Dispatcher maximizes total surplus instead of each driver's individual profit. Moreover, the algorithm also accounts for the externality from congestion as in equations $(II), (IV), (V)$, and the externality from endogenous demand as in equation (VII).

3.7.2 Algorithm 2: Value Function Approximation

The optimal intertemporal spatial allocation of taxis in terms of Social Surplus can be obtained by solving the following problem[4]:

$$max_{S_t} \sum_{t=1}^{T} E[TW(S_t, x_t)]$$

$$s.t. \quad S_{t+1} = transit(S_t, x_t) \quad \forall t$$

where:

- S_t : The spatial allocation of taxis at time t
- x_t : The planner's decision vector at time t, i.e. instructions to vacant taxis
- TW : The Social Surplus function for period t given the taxi allocation at time t, S_t
- $transit$: The state transition function that gives S_{t+1}, given S_t and x_t

Given the large number of possible allocations at a given time and the number of possible solution paths through time, which grows exponentially with respect to time, the solution of the latter problem is computationally impossible.

Instead, I propose an approximate solution based on Value Function Approximation. The output of this algorithm is a function that takes as input an allocation S and a time period t, and gives as output the estimated Social Surplus that occurs given that the optimal path will be taken after time t, given the initial allocation S_t. The estimation uses a neural network in order

[4]The Central Dispatcher is assumed not to discount time during the day.

to achieve a more effective fit and is generalized through time using backward induction. The algorithm starts at the last period, T. No decisions are made at T, and the Social Surplus added at T is given by $TW(S_T)$ $\forall S_T$. The function TW is known but is highly nonlinear because of the need to calculate the number of passengers, traffic and matches for a given S_t. The first step of the algorithm is to produce an estimator $\hat{V}_T$ for the social surplus at time T as a simple expression of S_T. Then, for $t = T - 1$ the function to be estimated will depend on $TW(S_{T-1})$ and $\hat{V}_T$. The result of the latter estimator is a function that describes the sum of social surplus to be added for every $t \geq T - 1$. The Value Function Approximation Algorithm is shown below.

Randomly sample K allocations $S_1, S_2, ..., S_K$ from the set of all possible allocations;
Calculate $TW_T(S_k)$ $\forall k = 1, 2, ..., K$;
Using the dataset $(S_k, TW(S_k))$ estimate TW_T: $\hat{V}_T$;
for $t = T - 1$ *to* 1 **do**
    Randomly sample K allocations $S_1, S_2, ..., S_K$;
    Randomly sample K decisions $x_1, x_2, ..., x_K$;
    Calculate $S'_k = transit(S_k, x_k)$, $\forall k = 1, 2, ..., K$;
    Calculate $V_t(S_k, x_k) = TW_t(S_k, x_k) + \hat{V}_{t+1}(S'_k)$, $\forall k = 1, 2, ..., K$;
    Using the dataset $(S_k, x_k, V_t(S_k, x_k))$ estimate V_t: $\hat{V}_t$;
    for $k = 1$ *to* K **do**
        Randomly sample L decisions $x_1, x_2, ..., x_L$ for S_k;
        Calculate $\hat{V}_t(S_k) = max_{x_l} \hat{V}_t(S_k, x_l)$;
        Using the dataset $(S_k, \hat{V}_t(S_k))$ estimate $\hat{V}_t(S_k, argmax(x_l))$: $\hat{V}_t$;
    end
end

Algorithm 3: Value Function Approximation Algorithm

I perform the estimations in steps 3, 9, 13 using a neural network with one hidden layer. The layers are made of nodes. A node is a place where computation happens, loosely patterned on a neuron in the human brain, which fires when it encounters sufficient stimuli. A node combines input from the data with a set of coefficients, or weights, that either amplify or dampen that input, thereby assigning significance to inputs with regard to the task the algorithm is trying to learn; e.g.

which input is most helpful is classifying data without error? These input-weight products are summed and then the sum is passed through a node's activation function, to determine whether and to what extent that signal should progress further through the network to affect the ultimate outcome. If the signal passes through, the neuron has been activated. More specifically, first each neuron adds up the value of every other neuron from the input column (i.e. spatial allocation) it is connected to. Before being added, each value is multiplied by another variable called "weight", w, which determines the connection between the two neurons. Each connection of neurons has its own weight, and those are the only values that will be modified during the learning process. Moreover, a bias value may be added to the total value calculated. The bias is a value not coming from a specific neuron and is chosen before the learning phase, but can be useful for the network. After the summations, the neuron finally applies an "activation function", f, to generate the obtained value. The activation functions that are used in literature include the identity function, which makes the output linear with respect to inputs, and sigmoid functions, e.g. $tanh(x)$, $\frac{1}{1+exp(-x)}$, etc. The activation function serves to turn the total value calculated previously to a numerical output. Figure 3.8 presents a graphical representation of the neural network.

Lemma 1: *The estimator $\hat{V}_t$ equals V_t on average $\forall t$.*

Proof: See Appendix.

After deriving all the estimated functions $\hat{V}_t(S,x)$ $\forall t < T$ and $\hat{V}_t(S)$ $\forall t \leq T$, I obtain the optimal path using the following algorithm.

The set $\{x^1, ..., x^{T-1}\}$ is the set of optimal decisions and the set $\{S_1, ..., S_T\}$ is the set of the

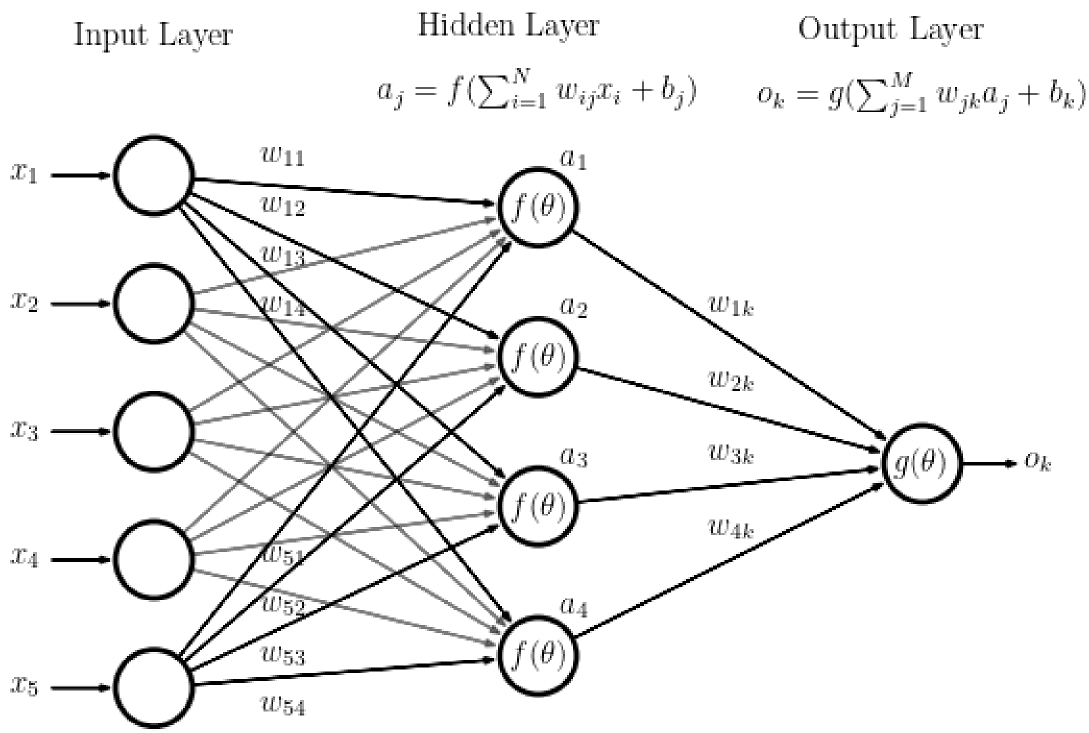

Figure 3.8: One Hidden Layer Neural Network

Read S_1 as the exogenous initial allocation ;
for $t = 1$ *to* $(T-1)$ **do**
  Randomly draw M samples of feasible decisions given $S_t : x_1, x_2, ..., x_M$;
  Calculate $x^t = argmax_{x_m} \hat{V}_t(S, x_m)$;
  Calculate $S_{t+1} = transit(S_t, x^t)$;
end

Algorithm 4: Optimal Path Algorithm

corresponding allocations. In conclusion, the Value Function Approximation algorithm is feasible and gives an approximation of the optimal solution. The latter outperforms the greedy algorithm's result. However, the greedy algorithm is relevant, since it provides a lower bound on the performance of the optimal solution and it is a tool for understanding the mechanisms behind the welfare gain (see Results section).

3.8 Results

This section discusses the Central Dispatcher counterfactual results for both algorithms used and compares them with the competitive allocation. I present the effects on social surplus, profit, matches, taxis per location, traffic speed and passengers over time. Briefly, the Central Dispatcher acts optimally in terms of social surplus by using the neural network based algorithm. The dynamics that lead to different results among the algorithms are:

1. The longest trip in these settings is the trip from Lower Manhattan to the Boroughs. This trip takes on average 2 time periods to be completed. The greedy algorithm, by construction, takes into account only the current and next period returns, therefore it does not allocate taxis from downtown to the Boroughs as it cannot internalize the benefit of this choice. Hence, there are fewer matches and vacant taxis over time in the Boroughs in the greedy algorithm's results.

2. In the competitive equilibrium a relatively low amount of taxis tend to choose to relocate to the boroughs, because of the low density of passengers in this location. The latter leads to decreased demand due to the supply shortage. As demand decreases, supply decreases even further. The Central Dispatcher is able to internalize this dynamic, and more taxis are allocated to the Boroughs. Demand increases and therefore matching becomes easier than in the competitive case. Note that matches in the Boroughs are the most profitable, as most of the trips end in downtown, hence are longer. In the competitive equilibrium supply in the Boroughs is low, as one driver by himself is unable to affect demand.

3. Passenger demand is structured such that relatively many passengers ask to commute between

Upper and Mid Manhattan. In addition, demand is dense in both locations. This leads to excess supply in this route in the competitive equilibrium. The Central Dispatcher internalizes the business stealing effect in this route due to excess supply, and redirects drivers to Lower Manhattan and the Boroughs. This is evident in the results, where the Central Dispatcher achieves a lower amount of matches and dispatches fewer vacant taxis in Upper and Mid Manhattan, while the opposite holds for Lower Manhattan and the Boroughs.

4. The results produced by the greedy algorithm display more variance. This happens because the greedy algorithm treats every period as the second last period. The greedy algorithm makes extreme decisions in order to maximize the next period return, which leads to non-efficient allocations in the future periods. The greedy algorithm alternates between more and less profitable periods.

5. Traffic speed increases in Mid and Upper Manhattan under the Central Dispatcher. The reason for this is that business stealing decreases in Mid and Upper Manhattan and therefore fewer drivers remain unmatched in these locations. The latter leads to fewer vacant drivers searching in Mid and Upper Manhattan in order to pick up passengers, leading the total number of vehicles under the critical number discussed in the congestion section.

The metric that is the basis of comparison between the different algorithms is the Social Surplus over time. Social Surplus is comprised of the producer surplus and the consumer surplus. Figures B.1, B.2 and B.3 provide the social surplus, producer surplus and consumer surplus over time respectively. The neural network algorithm achieves a 14.93% increase of social surplus compared to the competitive allocation. The percentage of welfare improvement from internalizing matching frictions is in accordance with the results of Kalouptsidi et al. (2019) and Brancaccio et al.(2020). The neural network algorithm achieve a higher social surplus outcome consistently over time. The greedy algorithm achieves a 8.14% increase of social surplus compared to the competitive allocation but the increase is not consistent over time.

Next, the results on matches for each location over time follow. Figures B.4, B.5, B.6 and B.7 present the matching outcomes for each allocation in absolute numbers. Note that the matching numbers presented by the competitive algorithm are the estimation results of the competitive equilibrium and not the matching data. However, the matching results and the data look similar. Any discrepancies are attributed to real dynamics not captured by the spatial model. As explained before, the Central Dispatcher achieves a small number of matches in Mid and Upper Mahnattan while the opposite holds in Lower Manhattan and the Boroughs.

Next, the results on vacant taxis for each location over time follow. Figures B.8, B.9, B.10 and B.11 present the vacant taxis distribution for each allocation in absolute numbers. Following the aforementioned intuition, the neural network algorithm achieves a lower average number of vacant taxis over time compared to the competitive allocation, as its matching process is more efficient and business stealing decreases.

Next, Figures B.12, B.13, B.14, and B.15 present the number of passengers looking for a taxi in each location over time. The latter is not related to matched passengers in each area. It describes only the number of consumers that attempt to get a taxi regardless of matching or not. The differences between the outcomes among the algorithms are that increased supply in Lower Manhattan and the Boroughs leads to increased demand in the Central Dispatcher case.

Finally, Figures B.16, B.17, B.18 and B.19 present the traffic speed for each allocation in absolute numbers. As expected, the Central Dispatcher is able to internalize the effects of traffic in driving cost and demand leading to less congestion under the neural network and greedy algorithms compared to the competitive equilibrium in the case of Upper and Mid Manhattan. This is achieved by avoiding having large numbers of vacant taxis looking for passengers in these locations. More specifically, the Central Dispatcher minimizes the number of taxis that fail to match in Upper and Mid Manhattan, internalizing that congestion effects are more intense in this area. This can be controlled by dispatching as many taxis to high congestion areas as needed in terms of predicted demand. The latter translates to a 7.5% increase in traffic speed in Mid Manhattan and a 10.6% increase in traffic speed in Upper Manhattan when the Central Dispatcher is added. The traffic speed achieved by the Central Dispatcher is not substantially different compared to the competitive equilibrium in the Boroughs because of the large surface area. The large surface area leads to a high critical density for congestion. The latter cannot be surpassed even when the Central Dispatches increases the number of drivers in the Boroughs.

3.9 Conclusion

Business stealing and congestion externalities decrease the efficiency of the taxi market. Drivers and customers search decisions are found to be non optimal and find each other less often. A Central Dispatcher can internalize the externalities and design an allocation that increases both driver and passenger utility on average. This paper adds congestion and endogenous demand to the standard framework created by Buchholz. I estimate the structural model using daily spatial and time data on taxi trips provided by the TLC. I identify congestion by the time needed to complete a trip shown by the data as a function of the the number of taxis traveling between locations. Matching efficiency as a function of congestion is identified by the variation created during rush hours. I identify the demand curve by the price variation among different trips and the welfare metrics by the demand elasticity, same way as in Buchholz.

I show that the Central Dispatcher can calculate an allocation that increases the social surplus in a timely manner using value function approximation. I use an one layer neural network in order to achieve a far more successful fit on the value function compared to analytical methods ,i.e. the R^2 increases from 75% to more than 95% by the use of the neural network. I show that the allocation suggested by the Central Dispatcher achieves a welfare gain of approximately $800 thousand per shift. A more sophisticated model of congestion that takes into account the number of private vehicles and a thorough modeling of alternative ways of transportation would possibly provide an even better allocation in terms of social surplus. A future expansion of the model could include decisions taken by drivers working for ridesharing firms like Uber and Lyft. This comes with a challenging part of modeling the ridesharing firms' pricing decision and matching mechanism. However, work done by Shapiro (2018), Bian (2019) and others provide useful insights on these puzzles.

Appendix A: Information, Bias, and Revenues in Sponsored Search Auctions

A.1 Paid Traffic Unpredictability

Dep. Variable:	paid_traffic_perc	**R-squared:**	0.030
Model:	OLS	**Adj. R-squared:**	0.003
Method:	Least Squares	**F-statistic:**	1.103
No. Observations:	5700	**AIC:**	4.880e+04
Df Residuals:	5541	**BIC:**	4.986e+04
Df Model:	158	**Log-Likelihood:**	-24243.

	coef	**std err**	**t**	**P>—*t*—**	**[0.025**	**0.975]**
y_lag1	0.0075	0.013	0.558	0.577	-0.019	0.034
y_lag2	-0.0269	0.013	-2.004	0.045	-0.053	-0.001
y_lag3	-0.0207	0.013	-1.545	0.122	-0.047	0.006
slot1	0.0524	0.088	0.595	0.552	-0.120	0.225
ratio	-0.0031	0.052	-0.060	0.952	-0.106	0.099
mean_bid	0.0104	0.011	0.961	0.337	-0.011	0.031
b1	-0.0214	0.036	-0.591	0.554	-0.092	0.050
b2	0.0052	0.028	0.184	0.854	-0.050	0.061
b3	0.1338	0.233	0.573	0.567	-0.324	0.591
b4	0.1202	0.116	1.034	0.301	-0.108	0.348
b5	0.0546	0.081	0.673	0.501	-0.104	0.214
b6	-0.1315	0.177	-0.742	0.458	-0.479	0.216
b7	6.08e-13	1.24e-12	0.491	0.624	-1.82e-12	3.04e-12
b8	0.0246	0.042	0.579	0.563	-0.059	0.108
b9	-0.0531	0.097	-0.546	0.585	-0.244	0.138
b10	0.0065	0.010	0.661	0.509	-0.013	0.026
b11	-0.0065	0.249	-0.026	0.979	-0.496	0.483
b13	0.0111	0.191	0.058	0.953	-0.362	0.385
b14	0.0007	0.001	0.661	0.509	-0.001	0.003
b15	0.0007	0.001	0.661	0.509	-0.001	0.003
b16	0.0007	0.001	0.661	0.509	-0.001	0.003
b17	-0.0065	0.049	-0.132	0.895	-0.103	0.090
b18	0.0066	0.009	0.725	0.468	-0.011	0.024
b19	0.0096	0.023	0.409	0.683	-0.036	0.056
b20	0.0305	0.046	0.667	0.505	-0.059	0.120
hotel_fe_1	-3.1795	3.233	-0.983	0.326	-9.518	3.159
hotel_fe_2	-3.8674	3.234	-1.196	0.232	-10.208	2.473
hotel_fe_3	-2.7876	3.233	-0.862	0.389	-9.126	3.551
hotel_fe_4	-4.6878	3.234	-1.449	0.147	-11.028	1.653
hotel_fe_5	-3.6254	3.234	-1.121	0.262	-9.965	2.714
hotel_fe_6	-1.7097	3.233	-0.529	0.597	-8.048	4.629
hotel_fe_7	-5.0246	3.235	-1.553	0.120	-11.366	1.316

hotel_fe_8	-4.5615	3.234	-1.410	0.158	-10.902	1.779
hotel_fe_9	-5.9552	3.235	-1.841	0.066	-12.297	0.387
hotel_fe_10	-5.2754	3.235	-1.631	0.103	-11.618	1.067
hotel_fe_11	-0.0674	3.233	-0.021	0.983	-6.405	6.271
hotel_fe_12	-0.2301	3.233	-0.071	0.943	-6.568	6.108
hotel_fe_13	-4.4964	3.235	-1.390	0.165	-10.838	1.845
hotel_fe_14	-5.0427	3.235	-1.559	0.119	-11.385	1.299
hotel_fe_15	-2.4505	3.233	-0.758	0.449	-8.789	3.888
hotel_fe_16	-6.3034	3.236	-1.948	0.052	-12.648	0.041
hotel_fe_17	-1.2441	3.233	-0.385	0.700	-7.583	5.094
hotel_fe_18	0.3088	3.233	0.096	0.924	-6.029	6.647
hotel_fe_19	-3.4924	3.234	-1.080	0.280	-9.833	2.848
hotel_fe_20	-4.7021	3.235	-1.454	0.146	-11.044	1.639
hotel_fe_21	-2.7897	3.234	-0.863	0.388	-9.129	3.550
hotel_fe_22	-5.6374	3.236	-1.742	0.082	-11.981	0.707
hotel_fe_23	-2.8778	3.234	-0.890	0.374	-9.218	3.462
hotel_fe_24	-5.1515	3.235	-1.593	0.111	-11.493	1.189
hotel_fe_25	-8.6118	3.239	-2.658	0.008	-14.962	-2.261
hotel_fe_26	-2.3567	3.234	-0.729	0.466	-8.696	3.982
hotel_fe_27	-3.4834	3.234	-1.077	0.281	-9.823	2.856
hotel_fe_28	-3.8055	3.234	-1.177	0.239	-10.146	2.535
hotel_fe_29	-0.3244	3.233	-0.100	0.920	-6.663	6.014
hotel_fe_30	-4.3203	3.235	-1.336	0.182	-10.662	2.021
hotel_fe_31	-1.5820	3.233	-0.489	0.625	-7.921	4.757
hotel_fe_32	0.2584	3.233	0.080	0.936	-6.080	6.597
hotel_fe_33	-3.1973	3.234	-0.989	0.323	-9.537	3.142
hotel_fe_34	-5.5163	3.236	-1.705	0.088	-11.859	0.827
hotel_fe_35	-5.0415	3.235	-1.559	0.119	-11.383	1.300
hotel_fe_36	-4.3601	3.234	-1.348	0.178	-10.701	1.980
hotel_fe_37	-0.7942	3.233	-0.246	0.806	-7.132	5.544
hotel_fe_38	-0.2722	3.233	-0.084	0.933	-6.610	6.066
hotel_fe_39	-4.8910	3.234	-1.512	0.131	-11.232	1.450
hotel_fe_40	-1.6306	3.233	-0.504	0.614	-7.969	4.708
hotel_fe_41	-3.1043	3.234	-0.960	0.337	-9.444	3.235
hotel_fe_42	-1.7315	3.233	-0.536	0.592	-8.070	4.607
hotel_fe_43	-1.5018	3.234	-0.464	0.642	-7.841	4.837
hotel_fe_44	-6.8106	3.236	-2.105	0.035	-13.154	-0.467
hotel_fe_45	-2.2089	3.233	-0.683	0.495	-8.547	4.129
hotel_fe_46	-5.7887	3.235	-1.789	0.074	-12.131	0.554
hotel_fe_47	-3.1168	3.234	-0.964	0.335	-9.457	3.223
hotel_fe_48	-1.0098	3.233	-0.312	0.755	-7.348	5.328
hotel_fe_49	-8.7593	3.238	-2.705	0.007	-15.108	-2.411
hotel_fe_50	-4.2028	3.234	-1.300	0.194	-10.543	2.137
hotel_fe_51	-2.7407	3.234	-0.848	0.397	-9.080	3.598
hotel_fe_52	-4.3130	3.235	-1.333	0.182	-10.654	2.028

hotel_fe_53	-2.5496	3.233	-0.789	0.430	-8.888	3.789
hotel_fe_54	-2.5821	3.234	-0.799	0.425	-8.921	3.757
hotel_fe_55	-2.6718	3.234	-0.826	0.409	-9.011	3.667
hotel_fe_56	-2.8456	3.233	-0.880	0.379	-9.184	3.493
hotel_fe_57	-0.2950	3.233	-0.091	0.927	-6.633	6.043
hotel_fe_58	-3.1938	3.234	-0.988	0.323	-9.533	3.145
hotel_fe_59	-5.2785	3.235	-1.632	0.103	-11.621	1.064
hotel_fe_60	-0.5457	3.233	-0.169	0.866	-6.884	5.792
hotel_fe_61	-6.9551	3.237	-2.148	0.032	-13.302	-0.608
hotel_fe_62	-6.6129	3.236	-2.044	0.041	-12.956	-0.270
hotel_fe_63	-2.3345	3.233	-0.722	0.470	-8.673	4.004
hotel_fe_64	-1.6884	3.233	-0.522	0.602	-8.027	4.650
hotel_fe_65	-4.4919	3.234	-1.389	0.165	-10.833	1.849
hotel_fe_66	-7.2143	3.237	-2.229	0.026	-13.560	-0.868
hotel_fe_67	-5.6724	3.235	-1.753	0.080	-12.014	0.670
hotel_fe_68	-3.8609	3.234	-1.194	0.233	-10.202	2.480
hotel_fe_69	-4.5035	3.234	-1.392	0.164	-10.844	1.837
hotel_fe_70	-2.7309	3.234	-0.845	0.398	-9.070	3.608
hotel_fe_71	-3.3697	3.235	-1.042	0.298	-9.711	2.971
hotel_fe_72	-7.9938	3.238	-2.469	0.014	-14.342	-1.646
hotel_fe_73	-1.6235	3.233	-0.502	0.616	-7.962	4.715
hotel_fe_74	-3.2562	3.234	-1.007	0.314	-9.596	3.083
hotel_fe_75	0.5139	3.233	0.159	0.874	-5.824	6.852
hotel_fe_76	-7.7207	3.237	-2.385	0.017	-14.066	-1.376
hotel_fe_77	-1.0610	3.233	-0.328	0.743	-7.399	5.277
hotel_fe_78	-1.6944	3.233	-0.524	0.600	-8.033	4.644
hotel_fe_79	-3.5138	3.234	-1.086	0.277	-9.854	2.826
hotel_fe_80	-6.8369	3.236	-2.112	0.035	-13.182	-0.492
hotel_fe_81	-1.5173	3.233	-0.469	0.639	-7.856	4.821
hotel_fe_82	-2.1774	3.233	-0.673	0.501	-8.516	4.161
hotel_fe_83	-3.8533	3.234	-1.191	0.234	-10.193	2.487
hotel_fe_84	2.6396	3.234	0.816	0.414	-3.700	8.979
hotel_fe_85	-4.5866	3.235	-1.418	0.156	-10.928	1.755
hotel_fe_86	-4.3359	3.234	-1.341	0.180	-10.676	2.004
hotel_fe_87	2.9544	3.234	0.914	0.361	-3.385	9.294
hotel_fe_88	-1.1035	3.233	-0.341	0.733	-7.442	5.235
hotel_fe_89	-1.7961	3.233	-0.556	0.579	-8.135	4.542
hotel_fe_90	-3.2275	3.234	-0.998	0.318	-9.567	3.112
hotel_fe_91	-1.2604	3.233	-0.390	0.697	-7.599	5.078
hotel_fe_92	-2.4078	3.234	-0.745	0.457	-8.747	3.931
hotel_fe_93	-2.7547	3.234	-0.852	0.394	-9.094	3.585
hotel_fe_94	-1.8398	3.233	-0.569	0.569	-8.178	4.498
hotel_fe_95	6.7491	3.238	2.085	0.037	0.402	13.096
hotel_fe_96	-4.8068	3.235	-1.486	0.137	-11.148	1.534

hotel_fe_97	-2.5166	3.233	-0.778	0.436	-8.855	3.822
hotel_fe_98	-4.1952	3.235	-1.297	0.195	-10.537	2.146
hotel_fe_99	2.8566	3.233	0.883	0.377	-3.482	9.195
time_fe_4	1.2798	1.565	0.818	0.414	-1.789	4.348
time_fe_5	-1.7951	2.007	-0.894	0.371	-5.730	2.140
time_fe_6	0.2074	1.661	0.125	0.901	-3.048	3.463
time_fe_7	-0.3190	1.682	-0.190	0.850	-3.616	2.978
time_fe_8	0.8469	2.418	0.350	0.726	-3.894	5.588
time_fe_9	0.9282	1.657	0.560	0.575	-2.321	4.177
time_fe_10	-1.7797	1.659	-1.073	0.283	-5.031	1.472
time_fe_11	0.3955	1.669	0.237	0.813	-2.876	3.667
time_fe_12	-3.8478	1.698	-2.266	0.023	-7.177	-0.519
time_fe_13	-1.4952	1.938	-0.772	0.440	-5.294	2.304
time_fe_14	-0.9184	1.587	-0.579	0.563	-4.029	2.193
time_fe_15	-0.8625	2.277	-0.379	0.705	-5.326	3.601
time_fe_16	1.1471	2.036	0.563	0.573	-2.845	5.139
time_fe_17	-0.5482	2.203	-0.249	0.804	-4.867	3.771
time_fe_18	-0.5981	1.648	-0.363	0.717	-3.829	2.632
time_fe_19	-0.7438	1.675	-0.444	0.657	-4.027	2.540
time_fe_20	-2.7364	2.149	-1.273	0.203	-6.950	1.477
time_fe_21	2.6116	2.203	1.185	0.236	-1.708	6.931
time_fe_22	0.1819	2.375	0.077	0.939	-4.475	4.838
time_fe_23	0.0982	1.682	0.058	0.953	-3.200	3.396
time_fe_24	-0.9251	1.785	-0.518	0.604	-4.425	2.575
time_fe_25	3.2172	1.686	1.908	0.056	-0.089	6.523
time_fe_26	1.0815	1.801	0.601	0.548	-2.449	4.612
time_fe_27	-0.8333	1.493	-0.558	0.577	-3.761	2.095
time_fe_28	-1.2051	1.535	-0.785	0.432	-4.214	1.804
time_fe_29	1.8154	1.659	1.094	0.274	-1.437	5.067
time_fe_30	-0.2202	1.615	-0.136	0.892	-3.386	2.946
time_fe_31	3.0358	1.643	1.848	0.065	-0.184	6.256
time_fe_32	1.2599	2.329	0.541	0.589	-3.305	5.825
time_fe_33	2.1782	1.699	1.282	0.200	-1.153	5.510
time_fe_34	-1.4047	2.293	-0.612	0.540	-5.901	3.091
time_fe_35	-0.2190	1.552	-0.141	0.888	-3.262	2.824
time_fe_36	-1.1705	1.525	-0.768	0.443	-4.160	1.819
time_fe_37	-2.0910	1.702	-1.228	0.219	-5.428	1.246
time_fe_38	-0.5965	1.499	-0.398	0.691	-3.536	2.343
time_fe_39	1.3372	1.681	0.795	0.426	-1.959	4.633
time_fe_40	-2.3244	1.993	-1.166	0.244	-6.232	1.584
time_fe_41	-0.3130	1.543	-0.203	0.839	-3.337	2.711
time_fe_42	0.1409	2.102	0.067	0.947	-3.979	4.261
time_fe_43	-2.5176	1.698	-1.483	0.138	-5.847	0.811
time_fe_44	-1.6177	1.541	-1.050	0.294	-4.638	1.403
time_fe_45	1.9838	1.552	1.278	0.201	-1.060	5.027

time_fe_46	0.7624	1.652	0.462	0.644	-2.476	4.001
time_fe_47	1.5294	1.646	0.929	0.353	-1.698	4.756
time_fe_48	-3.6461	1.669	-2.184	0.029	-6.919	-0.374
time_fe_49	-0.1445	1.719	-0.084	0.933	-3.514	3.225
time_fe_50	1.9528	1.674	1.167	0.243	-1.328	5.234
time_fe_51	-0.6879	1.560	-0.441	0.659	-3.745	2.369
time_fe_52	3.3500	1.547	2.166	0.030	0.318	6.382
time_fe_53	-2.4731	1.639	-1.508	0.131	-5.687	0.741
time_fe_54	1.5950	2.008	0.794	0.427	-2.341	5.531
time_fe_55	0.2765	1.576	0.175	0.861	-2.813	3.366
time_fe_56	-1.9586	1.572	-1.246	0.213	-5.040	1.123
time_fe_57	0.7511	1.984	0.379	0.705	-3.138	4.641
time_fe_58	1.7930	1.644	1.091	0.275	-1.429	5.015
time_fe_59	2.3668	1.515	1.562	0.118	-0.603	5.337

Omnibus:	2414.941	**Durbin-Watson:**	1.999
Prob(Omnibus):	0.000	**Jarque-Bera (JB):**	291.083
Skew:	0.011	**Prob(JB):**	6.20e-64
Kurtosis:	1.893	**Cond. No.**	1.73e+19

Table A.1: Panel Regression Results

A.2 Bids and Learning

Figure A.1 shows how bids evolve over a month (60 auctions) for a hotel. The bid axis is unnumbered for confidentiality reasons. The observed behavior matches the model design: Bidders change their bids over time to discover their opponents' bids and therefore infer their valuations. In the model, this is depicted as signal inference. However, it is straightforward to see that the moving average of each bidder does not change significantly. This suggests that bidders' ability of predicting CTR does not improve over time. In the model, this is depicted by keeping the prior distribution for each bidder the same across time.

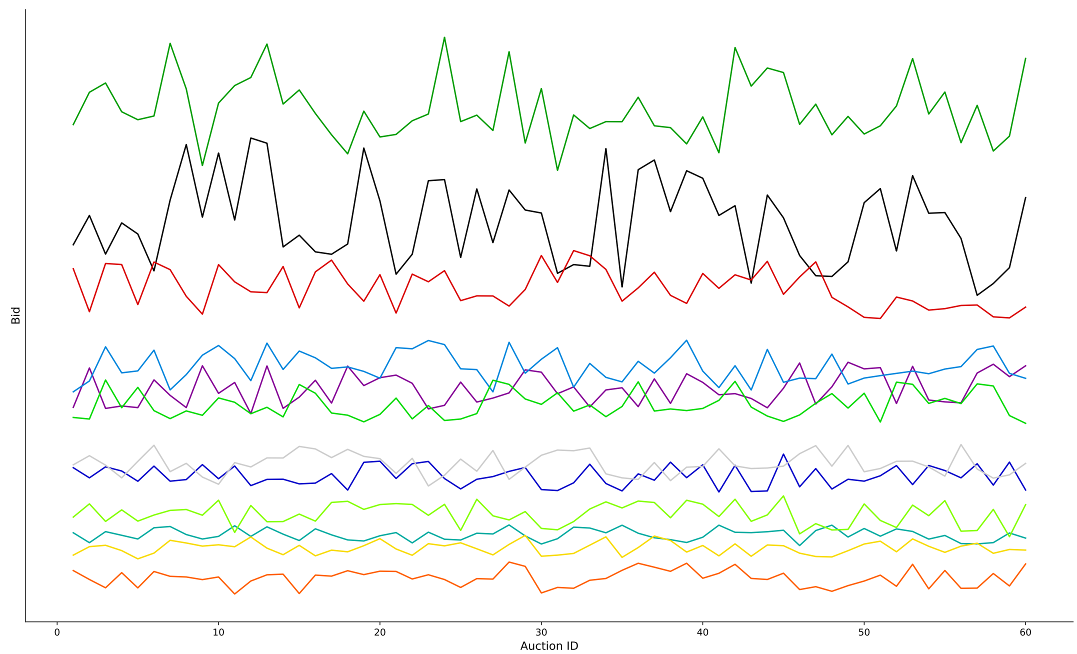

Figure A.1: Bids across time within hotel

A.3 Model Solution

Proposition 1 *The dropout price in the two bidder Generalized English Auction with asymmetric information is:*

$$p_i^* = \frac{(\alpha - 4)c_2 + 2s_i}{\alpha \cdot c_2 + 2s_i} \cdot v + \frac{4c_2}{\alpha \cdot c_2 + 2s_i} \cdot r$$

Proof: Denote the opponent bidder's quitting price on this equilibrium as p_{-i}^*. The latter is a function of the bidder valuations, the signal of bidder $-i$ and the known parameters. Then for bidder i, the utility of quitting at price p is:

$$\underline{u_i} = c_2 \cdot (v - r)$$

The utility of staying while bidder $-i$ quits at p, is:

$$\bar{u_i}(p) = (v - p) \cdot \mathbb{E}[c_1 | s_i, p_{-i}^*(s_{-i}) = p]$$

It is not hard to see that dropping out before the price level that makes the bidder indifferent between slot 2 and the expected value of slot 1 is a weekly dominated strategy, when the opponent bidder drops out when they are indifferent. In addition, it is not profitable to stay for higher prices than the indifference price since if the bidder wins it means that she has lower expected valuation than her opponent. This is because the opponent drops out at the indifference point, and the fact their signal is higher means that the winner bidder is worse off getting slot 1 than slot 2 at this price, in expectation.

In equilibrium, that is p_i^*, the following holds:

$$(v - p_i^*) \cdot \mathbb{E}[c_1 | s_i, s_{-i} = p_{-i}^{*-1}(p_i^*) = s_i] = c_2 \cdot (v - r)$$

The calculation of $\mathbb{E}[c_1|s_i, s_i]$ occurs from the calculation of the posterior distribution. In this case, when the signal comes from a Weibull(λ, β) and the prior is an Inverse Gamma (a, b), then the posterior distribution after receiving n independent signals is an Inverse Gamma $(a + n, b + \sum_{i=1}^{n} x_i^{\beta})$.

To ensure the existence of a Bayesian Equilibrium, bidders treat all signals equally. The latter ensures that if one signal increases, the utility of the bidder who owns the signal is affected the most. Given that there are 2 bidders in this case:

$$\mathbb{E}[c_1|s_i, s_i] = \frac{\alpha \cdot c_2 + 2 \cdot s_i}{4}$$

Therefore:

$$(v - p_i^*) \cdot \frac{\alpha \cdot c_2 + 2 \cdot s_i}{4} = c_2 \cdot (v - r) \Longrightarrow$$

$$p_i^* = \frac{(\alpha - 4) c_2 + 2 s_i}{\alpha \cdot c_2 + 2 s_i} \cdot v + \frac{4 c_2}{\alpha \cdot c_2 + 2 s_i} \cdot r$$

A.3.1 Example - Three bidders

I provide an additional example with 3 bidders to display the effect of signal inference and belief update that bidders perform when an opponent bidder drops out.

- 3 bidders
- Same valuation per click, $v_1 = v_2 = v_3 = v$
- Three slots, with click rates c_1, c_2, c_3.
- c_2, c_3 are known to all bidders, while c_1 is not known.
- Each bidder i gets a signal s_i from a Weibull distribution, $W[1, c_1]$.
- The signals that bidders get are conditionally independent, given c_1.
- Reserve price: r

Timing of the game:

1. Each bidder receives an independent private signal from $W[1, c_1]$.
2. All bidders update their beliefs on c_1 based on their private signal.
3. A clock showing the current price, which continuously increases over time, starts at the reserve price r.
4. As the current price continuously increases a bidder drops out - denote with p_3. This bidder gets the last slot (slot 3) and pays r for each click in this slot.

5. The remaining bidders infer the signal of the bidder who dropped out.

6. The remaining bidders update their beliefs on c_1 based on their inference of the signal of the bidder who just dropped out and their private signal, using the Bayesian rule.

7. The clock continues from p_3.

8. As the current price continuously increases a bidder drops out - denote with p_2. This bidder gets the last available slot (slot 2) and pays p_3 for each click in this slot.

9. The clicks in slot 1 are realized and the remaining bidder gets slot 1 and pays p_2 for each click in slot 1.

In a symmetric equilibrium in this case, a bidding decision is not pivotal when one opponent bidder leaves. In case the latter were true a contradiction occurs.[1]

Proposition 2 *The symmetric, increasing equilibrium strategy profile for the bidders in the three*

[1] In this environment, a bidder drops out when they are indifferent between getting slot 3 at the reserve price and slot 2 at the current price in case an opponent bidder quits. The increasing equilibrium function assumption dictates that the bidder assumes that she has the same signal with the quitting bidder at this price, as before $(s_{-i} = p^{*-1}_{-i}(p^*_i) = s_i)$. But the bidder cannot win slot 1 in case they stay at this price. This is because they were indifferent between staying and quitting, exactly like the bidder who dropped out. The fact that the third bidder is still in the auction means that he has the highest signal, because of the increasing equilibrium function assumption. Therefore, at the indifference point, the bidder is indifferent between slot 3 and slot 2, and not between slot 3 and entering a game for slot 1 or slot 2, because this game is lost already. The latter fact makes the equilibrium function to depend solely on slot 3 and slot 2 characteristics, which are not stochastic. In other words, the equilibrium function is the same as EOS (2007) and does not depend on the signal. This is a contradiction since the increasing equilibrium function assumption dictates a strictly increasing bidding function in signals.

In this edge case where all the valuations are equal, a bidder has no hope of winning slot 1 when they do not have the highest signal. If the pivotal bidding decision is about slot 2 and slot 3, the signal does not matter, hence all bidders drop out at the same price. This price is the price where slot 2 is equally profitable with slot 3 at the reserve price and thus, does not depend on the signal. This argument leads to the exploration of the following possibility: the pivotal bidding strategy to be when the two opponents drop out simultaneously and the remaining bidder gets slot 1.

bidder Generalized English Auction with asymmetric information is:

$$[(p^*(s_i), p^{\#}(s_i, s_0, p_0)), (p^*(s_j), p^{\#}(s_j, s_0, p_0)), (p^*(s_i), p^{\#}(s_k, s_0, p_0))]$$

where,

$$p_i^* = \frac{(\alpha - 4)c_2 + 2s_i}{\alpha \cdot c_2 + 2s_i} \cdot v + \frac{4c_2}{\alpha \cdot c_2 + 2s_i} \cdot r$$

$$p_i^{\#} = \frac{(\alpha - 5)c_2 + 2s_i + s_0}{\alpha c_2 + 2s_i + s_0} \cdot v + \frac{5c_2}{\alpha c_2 + 2s_i + s_0} \cdot p_0$$

given s_0 and p_0 at the start of the second stage, i.e. game history.

Proof: For bidder i, the utility of quitting at price p is:

$$\underline{u_i} = c_3 \cdot (v - r)$$

The expected utility for bidder i when staying in the auction while the two opponent bidders j, k quit and getting slot 1 is:

$$\bar{u_i}(p) = (v - p) \cdot \mathbb{E}[c_1 | s_i, p_j^*(s_j) = p_k^*(s_k) = p]$$

In the symmetric, strictly increasing equilibrium strategy profile: $(p_i^*, p_j^*, p_k^*) = (p^*, p^*, p^*)$. Then, in equilibrium, the expected utility of bidder i staying in the auction when bidders j, k drop out is equal with the utility of bidder i dropping out:

$$(v - p^*) \cdot \mathbb{E}[c_1 | s_i, s_j = p^{*-1}(p^*(s_i)) = s_i, s_k = p^{*-1}(p^*(s_i)) = s_i] = c_3 \cdot (v - r)$$

As previously,

$$\mathbb{E}[c_1 | s_i, s_i, s_i] = \frac{\alpha \cdot c_2 + 3 \cdot s_i}{5}$$

Therefore:

$$(v - p_i^*) \cdot \frac{\alpha \cdot c_2 + 3 \cdot s_i}{5} = c_3 \cdot (v - r) \implies$$

$$p_i^* = \frac{\alpha c_2 - 5c_3 + 3s_i}{\alpha \cdot c_2 + 3s_i} \cdot v + \frac{5c_3}{\alpha \cdot c_2 + 3s_i} \cdot r$$

This is not the full strategy profile since the game continues when the bidder with the lowest signal drops out. After one bidder drops out, the remaining two bidders enter to the subgame of section 1. The ”reserve price” p_0 of the subgame is the price that the lowest signal bidder dropped out. The bidder who dropped out was indifferent between slot 1 and slot 3 at p_0, therefore, the remaining bidders have higher indifference prices than r_2 between slot 1 and slot 2. The reasons for this is that the remaining bidders have higher signals than the bidder who dropped out. Some modifications needed before using the expression derived in section 1 are:

- Bidders evaluate the expected number of clicks in slot 1 not by only using their signal, but the signal of the bidder who dropped out, s_0.
- The reserve price is $r = p_0$

Entering the subgame, the utility of quitting at price p for bidder i and getting slot 2 is:

$\underline{u_i} = c_2 \cdot (v - p_0)$

The utility of staying while bidder $-i$ quits at p, is:

$\bar{u_i}(p) = (v - p) \cdot \mathbb{E}[c_1 | s_i, p^{\#}_{-i}(s_{-i}) = p, s_0]$

Because of symmetry, $p^{\#}_{-i} = p^{\#}_{i}$, hence $\mathbb{E}[c_1 | s_i, p^{\#}_{-i}(s_{-i}) = p, s_0] = \mathbb{E}[c_1 | s_i, s_i, s_0]$, in equilibrium.

The expected value of c_1 after calculating the posterior beliefs is:

$$\mathbb{E}[c_1 | s_i, s_i, s_0] = \frac{\alpha \cdot c_2 + 2 \cdot s_i + s_0}{5}$$

Therefore:

$$(v - p^{\#}_{i}) \cdot \frac{\alpha \cdot c_2 + 2 \cdot s_i + s_0}{5} = c_2 \cdot (v - p_0) \Longrightarrow$$

$$p^{\#}_{i} = \frac{(\alpha - 5)c_2 + 2s_i + s_0}{\alpha c_2 + 2s_i + s_0} \cdot v + \frac{5c_2}{\alpha c_2 + 2s_i + s_0} \cdot p_0$$

In conclusion, the symmetric, increasing equilibrium strategy profile for the bidders is:

$$[(p^*(s_i), p^{\#}(s_i, s_0, p_0)), (p^*(s_j), p^{\#}(s_j, s_0, p_0)), (p^*(s_i), p^{\#}(s_k, s_0, p_0))]$$

given s_0 and p_0 at the start of the second stage, i.e. game history.[2]

This case with equal valuations is a special case of the more general problem, where valuations can vary. In the general case, the bidder with the highest signal does not always win slot 1. This major difference changes the pivotal bidding decision in the case of three bidders. When an opponent bidder drops out when the bidder is indifferent, it does not mean that they have the same signal.

A.3.2 Comparison with the Game with Known CTR

In this subsection, I compare the Seller's revenue in the cases of asymmetric information about CTR and when CTR is common knowledge. In addition, I calculate for which signal is the Bidder's expected return maximized. I cover the two-bidder case since the algebra in the three-bidder is considerably tedious.

A.3.2.1 Expected Return, given the Signal

The expected return of a bidder with signal s in the aforementioned equilibrium, ex ante, is:

$$R^e(s) = \int_0^\infty \underbrace{Ret(s,x)}_{\text{Return, given opponent's signal}} \cdot \underbrace{f(x)}_{\text{Distribution of the signal}} \cdot dx =$$

$$= \int_0^s \underbrace{c_1 \cdot (v - (v - \tfrac{4c_2(v-r)}{\alpha c_2 + 2x}))}_{\text{Return of slot 1, given opponent's signal}} \cdot f(x) \cdot dx + \int_s^\infty \underbrace{c_2 \cdot (v-r)}_{\text{Return of slot 2, given opponent's signal}} \cdot f(x) \cdot dx =$$

[2]Stage is defined in a specific auction, between bidder dropouts. I do not assume that different auctions happen simultaneously. In any case, the clock setting of the auction is artificial, just to transit to the generalized English auction environment.

$$= c_2(v-r)(1-F(s)) + \int_0^s \frac{4c_1c_2(v-r)}{\alpha c_2+2x} f(x)dx =$$

$$= c_2(v-r)e^{-\frac{s}{\alpha c_2}} + \frac{4c_1(v-r)}{\alpha} \cdot \frac{\sqrt{e}}{2} \cdot [\underbrace{E_1(\tfrac{1}{2})}_{\text{Exponential integral}} - E_1(\tfrac{2s+\alpha c_2}{2\alpha c_2})]$$

Assuming that the mean of the bidders' prior is equal to the amount of clicks in slot 1, i.e. $\alpha c_2 = c_1$:

$$R^e(s) = c_2(v-r) \cdot (e^{-\frac{s}{c_1}} + 2\sqrt{e} \cdot [E_1(\tfrac{1}{2}) - E_1(\tfrac{2s+c_1}{2c_1})])$$

The signal s^*, that maximizes the expected return for the bidder is:

$$\frac{dR^e(s)}{ds}(s^*) = 0 \Longrightarrow s^* = \frac{3}{2}c_1$$

It occurs that the game benefits slightly optimistic bidders. This result might seem counterintuitive. The expected revenue of the bidder is an increasing function of the bidder's signal until $\frac{3}{2}c_1$ and decreasing afterwards. A reasonable expectation would be to for the expected revenue of the bidder to maximize at the true value of clicks, c_1, but this is not the case. When the opponent's signal is lower than the true value, c_1, it does matter whether the bidder's signal is equal to c_1 or more since the opponent drops out before this point. The same holds when the opponent's signal is more than $\frac{3}{2}c_1$. The difference becomes apparent when the opponent's signal lies between c_1 and $\frac{3}{2}c_1$. Consider the example when the opponent's signal equals $c_1 + \epsilon$, when ϵ is positive and very small and the bidder's signal equals the true value, c_1. Bidder rationality suggests that the bidder with the signal equal with c_1 drops out before her opponent because according to her information she is better off getting slot 2 at the reserve price than winning slot 1 against a bidder with signal $c_1 + \epsilon$, in expectation. In this case, although rational, the bidder is wrong. The reason is the quality of her information. In this Bayesian setting, bidders start with their prior and update it with theirs and their opponents' signals.

In a 2-bidders game, the initial prior has a significant weight in the posterior because the bidder can only process two signals. Because of the prior distribution used in this example, the bidder with signal c_1 believes she is better off dropping out before the bidder with signal $c_1 + \epsilon$. The reason is that this bidder does not know the number of clicks in slot 1, and she uses her prior and available signal to infer it. The prior used in this example leads the bidder to underestimate the number of clicks in slot 1. She would be better off not to drop out before the bidder with signal $c_1 + \epsilon$, given that the real number of clicks is c_1. But this is not a profitable deviation in the game's setting, since the bidder does not know that the number of clicks in slot 1 equals their signal.

A.3.2.2 Seller's revenue

The expected revenue of the seller is the sum of revenues from the two slots:

$$R^e_{seller} = c_2 r + \int_0^\infty c_1 \cdot \left(v - \frac{4c_2(v-r)}{\alpha c_2 + 2x}\right) \qquad \underbrace{g(x)}_{\text{Distribution of lowest signal}} \qquad dx$$

The distribution of the lowest signal comes from the first order statistic of two independent draws from a Weibull$(1, c_1)$ with pdf equal to f.

$$g(x) = [1 - F(x)] \cdot f(x) = \frac{1}{c_1} e^{-\frac{2x}{c_1}}$$

The integral has not a closed form expression and has to be calculated numerically. Assuming that the mean of the bidders' prior is equal to the amount of clicks in slot 1, i.e. $\alpha c_2 = c_1$:

$$R^{e}_{seller} = c_1 v(1 - \frac{2.38}{\alpha}) + c_2 r \cdot 3.38$$

Switching to the known CTR case, where c_1 is known to the bidders, a natural equilibrium is the one where both bidders drop out at $p = v - \frac{c_2}{c_1}(v - r)$. The slots are allocated randomly, and bidders receive equal returns. Therefore, the seller's revenue is:

$$R^{known}_{seller} = c_2 r + c_1(v - \frac{c_2}{c_1}(v - r)) = (c_1 - c_2) \cdot v + 2c_2 r$$

Then:

$$R^{known}_{seller} - R^{e}_{seller} = 1.38 \cdot c_2(v - r) > 0$$

This result is expected since bidders extract rents in the asymmetric information case.

A.4 Estimation Proofs

Proposition 4 *The signal that each bidder receives in each auction can be calculated given the history of the game within an auction and has the following form:*

$$s_i = c(c_2, n, t) \cdot \alpha_i + d(c_2, c_t, b_i, b_t, t, n, r)$$

where:

$c(c_2, n, t) = -\frac{c_2}{n-t+1} + \frac{c_2}{n^2} + \sum_{j=2}^{t-1} \frac{c_2}{(n-j)\cdot(n-j+1)}$

$d(c_2, c_t, b_i, b_t, t, n, r) = \frac{n+2}{n-t+1} \cdot \frac{v-b_{t-1}}{v-b_t} \cdot c_t - \frac{1}{n} \cdot \frac{n+2}{n} \cdot \frac{v-r}{v-b_1} \cdot c_1 - \sum_{j=2}^{t-1} \frac{1}{n-j} \cdot \frac{n+2}{n-j+1} \cdot \frac{v-b_{j-1}}{v-b_j} \cdot c_j$

Proof: Equation (1.4) states that

$$p_i^t = v - \frac{(n+2) \cdot c_t}{\alpha_i \cdot c_2 + (n-t+1)s_i + \sum_{j=1}^{t-1} s_j^i(b_j)} \cdot (v - b_{t-1})$$

The bidder drops out at p_i^t which signifies b_t. Solving with respect to s_i:

$$s_i = \frac{n+2}{n-t+1} \cdot \frac{v-b_{t-1}}{v-b_t} \cdot c_t - \alpha \cdot \frac{c_2}{n-t+1} - \frac{\sum_{j=1}^{t-1} s_j^i(b_j)}{n-t+1}$$

The term $s_j^i(b_j)$ denotes the signal that bidder i infers for any bidder j that dropped out before her. The inference of the opponents' signals is wrong since bidder i believes that her opponents have the same prior with her. Using the history of the game before bidder i's drop out $\{b_1, ..., b_{t-1}\}$, I solve for the inferred bids for bidder i.

$$\sum_{j=1}^{t-1} s_j^i(b_j) = s_{t-1}^i(b_{t-1}) + \sum_{j=1}^{t-2} s_j^i(b_j) \overset{(1.4)}{=}$$

$$\overset{(1.4)}{=} \frac{n+2}{n-t+2} \cdot \frac{v-b_{t-2}}{v-b_{t-1}} \cdot c_{t-1} - \alpha \cdot \frac{c_2}{n-t+2} - \frac{\sum_{j=1}^{t-2} s_j^i(b_j)}{n-t+2} + \sum_{j=1}^{t-2} s_j^i(b_j) =$$

$$= \frac{n+2}{n-t+2} \cdot \frac{v-b_{t-2}}{v-b_{t-1}} \cdot c_{t-1} - \alpha \cdot \frac{c_2}{n-t+2} + \frac{n-t+1}{n-t+2} \cdot \sum_{j=1}^{t-2} s_j^i(b_j) \overset{(1.4)}{=} \dots$$

$$\overset{(1.4)}{=} \frac{n+1-t}{n} \cdot s_1^i(b_1) + \frac{n+1-t}{n-1} con_1 + \frac{n+1-t}{n-2} con_2 + \dots + \frac{n+1-t}{n+2-t} con_{t-2} + con_{t-1}$$

where

$$con_j = \frac{n+2}{n-j+1} \cdot \frac{v-b_{j-1}}{v-b_j} \cdot c_j - \alpha \cdot \frac{c_2}{n-j+1}, \quad j > 1$$

The solution for s_1 is straightforward by using equation (1.5) and solving for the signal of the first drop out, given bidder i's beliefs:

$$s_1^i(b_1) = \frac{n+2}{n} \cdot \frac{v-r}{v-b_1} - \alpha \cdot \frac{c_2}{n}$$

Hence, the signal of bidder i that drops out at b_t, winning slot t, given history $\{b_1, ..., b_{t-1}\}$ is:

$$s_i = con_t - \sum_{j=1}^{t-1} \frac{1}{n-j} con_j$$

where

$$con_j = \frac{n+2}{n-j+1} \cdot \frac{v-b_{j-1}}{v-b_j} \cdot c_j - \alpha \cdot \frac{c_2}{n-j+1}, \quad j > 1$$

$$con_1 = \frac{n+2}{n} \cdot \frac{v-r}{v-b_1} - \alpha \cdot \frac{c_2}{n}$$

One can observe that the variable which describes the bidder's prior, α, enters the equation linearly. Hence, by separating the terms which contain α, it occurs:

$$s_i = \underbrace{(-\frac{c_2}{n-t+1} + \frac{c_2}{n^2} + \sum_{j=2}^{t-1} \frac{c_2}{(n-j)\cdot(n-j+1)})}_{c(c_2,n,t)} \cdot \alpha +$$

$$+ \underbrace{\frac{n+2}{n-t+1} \cdot \frac{v-b_{t-1}}{v-b_t} \cdot c_t - \frac{1}{n} \cdot \frac{n+2}{n} \cdot \frac{v-r}{v-b_1} \cdot c_1 - \sum_{j=2}^{t-1} \frac{1}{n-j} \cdot \frac{n+2}{n-j+1} \cdot \frac{v-b_{j-1}}{v-b_j} \cdot c_j}_{d(c_2,c_t,b_i,b_t,t,n,r)}$$

Proposition 5 *The estimator for* $\alpha_{i,h}$

$$\hat{\alpha_{i,h}} = -\frac{\sum_{m=1}^{n} \tilde{d_{i,h}^m} \cdot c_{i,h}^m}{\sum_{m=1}^{n} (c_{i,h}^m)^2}$$

is consistent, i.e.

$$\lim_{n\to\infty} \hat{\alpha_{i,h}} = \alpha_{i,h}$$

with residuals for bidder i *over auctions indexed by* m*:*

$$\hat{\epsilon}_i^m = \frac{1}{M} \sum_{m=1}^{M} \tilde{d_{i,h}^m} + \hat{\alpha_{i,h}} \cdot c_{i,h}^m$$

Proof: The estimator is derived by the maximum likelihood estimation method. Given a bidder i and a hotel h, there is a set of M auctions where bidder i participates in location l. Using the result of Proposition 4, I write:

$$\underbrace{\tilde{d_{i,h}^{m}}}_{d_{i,h}^{m}-c_1} = -c_{i,h}^{m} \cdot \alpha_{i,h} + \underbrace{\tilde{s_{i,h}^{m}}}_{s_{i,h}^{m}-c_1}$$

Given that $s_{i,h}^{m}$ follows distribution $W[1, c_1, 0]$, then $\tilde{d_{i,h}^{m}}$ follows distribution $W[1, c_1, -c_1 - c_{i,h}^{m} \cdot \alpha_{i,h}]$, where the third parameter is the location of the distribution. The goal is to find the parameter $\alpha_{i,h}$ that maximizes the likelihood that $\tilde{d_{i,h}^{m}}$ follows $W[1, c_1, -c_1 - c_{i,h}^{m} \cdot \alpha_{i,h}]$. Note that the data allow for the calculation of $c_{i,h}^{m}$ and $\tilde{d_{i,h}^{m}}$. The signals that a bidder receives across auctions are independent, hence the likelihood that $\tilde{d_{i,h}^{m}}$ follows $W[1, c_1, -c_1 - c_{i,h}^{m} \cdot \alpha_{i,h}]$ is:

$$L = \prod_{m=1}^{n} \frac{1}{c_1} e^{-(\frac{\tilde{d_{i,h}^{m}} + c_1 + c_{i,h}^{m} \cdot \alpha_{i,h}}{c_1})^2} \Longrightarrow$$

$$l = \sum_{m=1}^{n} ln(\frac{1}{c_1} e^{-(\frac{\tilde{d_{i,h}^{m}} + c_1 + c_{i,h}^{m} \cdot \alpha_{i,h}}{c_1})^2}) \Longrightarrow$$

$$l = -n \cdot ln(c_1) - \frac{1}{c_1^2} \cdot \sum_{m=1}^{n} (\tilde{d_{i,h}^{m}} + c_1 + c_{i,h}^{m} \cdot \alpha_{i,h})^2$$

I calculate the α that maximizes the log-likelihood:

$$\frac{\partial l}{\partial \alpha_{i,h}} |_{\tilde{\alpha_{i,h}}} = 0 \Longrightarrow$$

$$-\frac{2}{c_1} \cdot \sum_{m=1}^{n} (\tilde{d_{i,h}^{m}} + c_1 + c_{i,h}^{m} \cdot \tilde{\alpha_{i,h}}) \cdot c_{i,h}^{m} = 0 \Longrightarrow$$

$$\tilde{\alpha_{i,h}} = -\frac{\sum_{m=1}^{n}(\tilde{d_{i,h}^{m}} \cdot c_{i,h}^{m}) + c_1 \cdot \sum_{m=1}^{n} c_{i,h}^{m}}{\sum_{m=1}^{n}(c_{i,h}^{m})^2}$$

Finally, I adjust for the added term c_1 and get the estimator:

$$\hat{\alpha_{i,h}} = \tilde{\alpha_{i,h}} - \frac{c_1 \cdot \sum_{m=1}^{n} c_{i,h}^{m}}{\sum_{m=1}^{n}(c_{i,h}^{m})^2} = -\frac{\sum_{m=1}^{n}(\tilde{d_{i,h}^{m}} \cdot c_{i,h}^{m})}{\sum_{m=1}^{n}(c_{i,h}^{m})^2}$$

To prove that the estimator is consistent, I use the equation for $\tilde{d_{i,h}^{m}}$:

$$\hat{\alpha_{i,h}} = -\frac{\sum_{m=1}^{n}((-c_{i,h}^{m} \cdot \alpha_{i,h} + \tilde{s_{i,h}^{m}}) \cdot c_{i,h}^{m})}{\sum_{m=1}^{n}(c_{i,h}^{m})^2} = \alpha_{i,h} - \frac{\frac{1}{n}\sum_{m=1}^{n} \tilde{s_{i,h}^{m}} \cdot c_{i,h}^{m}}{\frac{1}{n}\sum_{m=1}^{n}(c_{i,h}^{m})^2}$$

Then, by using WLLN:

$$\lim_{n\to\infty} \hat{\alpha_{i,h}} = \alpha_{i,h} - \frac{\mathbb{E}[\tilde{s_{i,h}^{m}} \cdot c_{i,h}^{m}]}{\mathbb{E}[(c_{i,h}^{m})^2]}$$

The random variable $\tilde{s_{i,h}^{m}} = s_{i,h}^{m} - c_1$ is drawn from $W[1, c_1, -c_1]$ independently of the auction environment - including $c_{i,h}^{m}$ - hence:

$$\mathbb{E}[\tilde{s_{i,h}^{m}} \cdot c_{i,h}^{m}] = \mathbb{E}[\tilde{s_{i,h}^{m}} | c_{i,h}^{m}] \cdot \mathbb{E}[c_{i,h}^{m}] = 0$$

since $s_{i,h}^{m}$ follows $W[1, c_1]$, therefore $\mathbb{E}[s_{i,h}^{m}] = c_1$

Thus,

$$\lim_{n\to\infty} \hat{\alpha_{i,h}} = \alpha_{i,h}$$

making the estimator $\hat{\alpha_{i,h}} = -\frac{\sum_{m=1}^{n} \tilde{d_{i,h}^{m}} \cdot c_{i,h}^{m}}{\sum_{m=1}^{n}(c_{i,h}^{m})^2}$ consistent.

Given that $\tilde{s^m_{i,h}}$ is the error term in the equation that is used to estimate $\alpha_{i,h}$ and that $\mathbb{E}[\tilde{s^m_{i,h}}] = 0$, the residuals of the estimator $\hat{\alpha_{i,h}}$ follow:

$$\hat{\epsilon}^m_i = \frac{1}{M}\sum_{m=1}^{M} \tilde{d^m_{i,h}} + \hat{\alpha_{i,h}} \cdot c^m_{i,h}$$

Proposition 6 *The estimator* $\hat{\alpha_{i,h}}$ is unbiased, i.e.

$$\mathbb{E}[\hat{\alpha_{i,h}}|C_{i,h}] = \alpha_{i,h}$$

Proof: The calculation is straightforward:

$$\mathbb{E}[\hat{\alpha_{i,h}}|C_{i,h}] = \alpha_{i,h} - \frac{\sum_{m=1}^{n} \mathbb{E}[\tilde{s^m_{i,h}}|C_{i,h}] \cdot c^m_{i,h}}{\sum_{m=1}^{n}(c^m_{i,h})^2} = \alpha_{i,h}$$

since

$$\mathbb{E}[\tilde{s^m_{i,h}}|C_{i,h}] = 0, \quad \forall m \in \{1, 2, ...n\}$$

Proposition 7 $var(\hat{\alpha_{i,h}}|C_{i,h}) = \frac{c_1^2}{\sum_{m=1}^{n}(c^m_{i,h})^2}$

Proof: Given proposition 6, it follows that:

$$var(\hat{\alpha_{i,h}}|C_{i,h}) = \mathbb{E}[(\hat{\alpha_{i,h}} - \alpha_{i,h}) \cdot \hat{\alpha_{i,h}} - \alpha_{i,h})'|C_{i,h}] =$$

$$= \mathbb{E}[(C'_{i,h}C_{i,h})^{-1} \cdot C'_{i,h}\tilde{S}_{i,h} \cdot \tilde{S}'_{i,h}C_{i,h} \cdot (C'_{i,h}C_{i,h})^{-1}|C_{i,h}] =$$

$$= (C'_{i,h}C_{i,h})^{-1} \cdot C'_{i,h} \cdot \mathbb{E}[\tilde{S}_{i,h}\tilde{S}'_{i,h}|C_{i,h}] \cdot C_{i,h} \cdot (C'_{i,h}C_{i,h})^{-1}$$

Since $s_{i,h}^{\tilde{m}} \, W[1, c_1, -c_1]$, it holds that $\mathbb{E}[\tilde{S}_{i,h}\tilde{S}'_{i,h}|C_{i,h}] = c_1^2 \cdot (\Gamma(3) - \Gamma(2)^2) = c_1^2$, hence:

$$var(\hat{\alpha_{i,h}}|C_{i,h}) = c_1^2 \cdot (C'_{i,h}C_{i,h})^{-1} \cdot C'_{i,h} \cdot C_{i,h} \cdot (C'_{i,h}C_{i,h})^{-1} =$$

$$= c_1^2 \cdot (C'_{i,h}C_{i,h})^{-1} = \frac{c_1^2}{\sum_{m=1}^{n}(c_{i,h}^m)^2}$$

Proposition 8 *The estimator for* mar_i^l

$$\hat{mar}_i^l = -\frac{\sum_{h \in \mathbb{H}_0} \sum_{m=1}^{n_h} ratio_h \cdot d_{i,h}^{\tilde{m}} \cdot c_{i,h}^m}{\sum_{h \in \mathbb{H}_0} \sum_{m=1}^{n_h} (ratio_h \cdot c_{i,h}^m)^2}$$

is consistent as the number of auctions per hotel increases to infinity for all hotels, i.e.

$$\lim_{\substack{n_h \to \infty, \\ \forall h \in \mathbb{H}_0}} \hat{mar}_i^l = mar_i^l$$

Proof: The derivation of the estimator follows the same steps with the estimator in Proposition

5. Knowing that for the hotels used in this estimator, $opt_{i,h} = 1 \ \forall i$, we get for equation (1.3):

$$\alpha_{i,h} = ratio_h \cdot mar_i^l \overset{(1.9)}{\Longrightarrow} \tilde{d_{i,h}^m} = -ratio_h \cdot mar_i^l \cdot c_{i,h}^m + \tilde{s_{i,h}^m}$$

Therefore:

$$\hat{mar_i^l} = -\frac{\sum_{h\in\mathbb{H}_0}\sum_{m=1}^{n_h} ratio_h \cdot \tilde{d_{i,h}^m} \cdot c_{i,h}^m}{\sum_{h\in\mathbb{H}_0}\sum_{m=1}^{n_h}(ratio_h \cdot c_{i,h}^m)^2} =$$

$$= \frac{mar_i^l \cdot \sum_{h\in\mathbb{H}_0}\sum_{m=1}^{n_h}(ratio_h \cdot c_{i,h}^m)^2}{\sum_{h\in\mathbb{H}_0}\sum_{m=1}^{n_h}(ratio_h \cdot c_{i,h}^m)^2} - \frac{\sum_{h\in\mathbb{H}_0}\sum_{m=1}^{n_h} ratio_h \cdot \tilde{s_{i,h}^m} \cdot c_{i,h}^m}{\sum_{h\in\mathbb{H}_0}\sum_{m=1}^{n_h}(ratio_h \cdot c_{i,h}^m)^2} =$$

$$= mar_i^l - \frac{\frac{1}{n_1}\sum_{m=1}^{n_1} ratio_1 \cdot \tilde{s_{i,1}^m} \cdot c_{i,1}^m}{\frac{1}{n_1}(\sum_{m=1}^{n_1}(ratio_1 \cdot c_{i,1}^m)^2 + \ldots + \sum_{m=1}^{n_H}(ratio_H \cdot c_{i,H}^m)^2)} - \ldots$$

$$\ldots - \frac{\frac{1}{n_H}\sum_{m=1}^{n_H} ratio_H \cdot \tilde{s_{i,H}^m} \cdot c_{i,H}^m}{\frac{1}{n_H}(\sum_{m=1}^{n_1}(ratio_1 \cdot c_{i,1}^m)^2 + \ldots + \sum_{m=1}^{n_H}(ratio_H \cdot c_{i,H}^m)^2)} =$$

$$= mar_i^l - \frac{\frac{1}{n_1}\sum_{m=1}^{n_1} ratio_1 \cdot \tilde{s_{i,1}^m} \cdot c_{i,1}^m}{\frac{1}{n_1}\sum_{m=1}^{n_1}(ratio_1 \cdot c_{i,1}^m)^2 + \ldots + \frac{n_H}{n_1} \cdot \frac{1}{n_H}\sum_{m=1}^{n_H}(ratio_H \cdot c_{i,H}^m)^2} - \ldots$$

$$\ldots - \frac{\frac{1}{n_H}\sum_{m=1}^{n_H} ratio_H \cdot \tilde{s_{i,H}^m} \cdot c_{i,H}^m}{\frac{n_1}{n_H} \cdot \frac{1}{n_1}\sum_{m=1}^{n_1}(ratio_1 \cdot c_{i,1}^m)^2 + \ldots + \frac{1}{n_H}\sum_{m=1}^{n_H}(ratio_H \cdot c_{i,H}^m)^2}$$

As $n_1, n_2, \ldots, n_H$ increase to infinity at the same rate (e.g. linear), note that $\lim_{(n_1,n_2,\ldots,n_H)\to\infty} \frac{n_i}{n_j} = 1 \ \forall i,j \in \{1,\ldots,H\}$. Then by using WLLN:

$$\lim_{\substack{n_h\to\infty,\\ \forall h\in\mathbb{H}_0}} \hat{mar_i^l} = mar_i^l - \frac{\mathbb{E}[\tilde{s_{i,1}^m} \cdot c_{i,1}^m \cdot ratio_1]}{\mathbb{E}[(c_{i,1}^m \cdot ratio_1)^2] + \ldots + \mathbb{E}[(c_{i,H}^m \cdot ratio_H)^2]} - \ldots$$

$$\ldots - \frac{\mathbb{E}[\tilde{s_{i,H}^m} \cdot c_{i,H}^m \cdot ratio_H]}{\mathbb{E}[(c_{i,1}^m \cdot ratio_1)^2] + \ldots + \mathbb{E}[(c_{i,H}^m \cdot ratio_H)^2]}$$

The random variable $s_{i,h}^{\tilde{m}} = s_{i,h}^m - c_1$ is drawn from $W[1, c_1, -c_1]$, $\forall h$, independently of the auction environment - including $c_{i,h}^m$ - hence:

$$\mathbb{E}[s_{i,h}^{\tilde{m}} \cdot c_{i,h}^m \cdot ratio_h] = ratio_h \cdot \mathbb{E}[s_{i,h}^{\tilde{m}} | c_{i,h}^m] \cdot \mathbb{E}[c_{i,h}^m] = 0$$

Thus,

$$\lim_{n \to \infty} \hat{mar}_i^l = mar_i^l$$

making the estimator $\hat{mar}_i^l = -\frac{\sum_{h \in \mathbb{H}_0} \sum_{m=1}^{n_h} ratio_h \cdot d_{i,h}^{\tilde{m}} \cdot c_{i,h}^m}{\sum_{h \in \mathbb{H}_0} \sum_{m=1}^{n_h} (ratio_h \cdot c_{i,h}^m)^2}$ consistent.

Proposition 9 *The estimator for* $opt_{i,h}$

$$\hat{opt}_{i,h} = -\frac{\sum_{m=1}^{n} \hat{mar}_i^l \cdot ratio_h \cdot d_{i,h}^{\tilde{m}} \cdot c_{i,h}^m}{\sum_{m=1}^{n} (\hat{mar}_i^l \cdot ratio_h \cdot c_{i,h}^m)^2}$$

is consistent, i.e.

$$\lim_{n \to \infty} \hat{opt}_{i,h} = opt_{i,h}$$

Proof: We get from equation (1.3):

$$\alpha_{i,h} = ratio_h \cdot opt_{i,h} \cdot mar_i^l \overset{(1.9)}{\Longrightarrow} d_{i,h}^{\tilde{m}} = -ratio_h \cdot mar_i^l \cdot opt_{i,h} \cdot c_{i,h}^m + s_{i,h}^{\tilde{m}}$$

Therefore: Using the expression for $\tilde{d_{i,h}^{m}}$:

$$\hat{opt}_{i,h} = -\frac{\sum_{m=1}^{n} \tilde{d_{i,h}^{m}} \cdot c_{i,h}^{m}}{\hat{mar}_i^l \cdot ratio_h \cdot \sum_{m=1}^{n}(\cdot c_{i,h}^{m})^2} =$$

$$= mar_i^l \cdot ratio_h \cdot \frac{\sum_{m=1}^{n} opt_{i,h} \cdot (c_{i,h}^{m})^2}{\hat{mar}_i^l \cdot ratio_h \cdot \sum_{m=1}^{n}(\cdot c_{i,h}^{m})^2} - \frac{1}{\hat{mar}_i^l \cdot ratio_h} \cdot \frac{\sum_{m=1}^{n} \tilde{s_{i,h}^{m}} \cdot c_{i,h}^{m}}{\sum_{m=1}^{n}(c_{i,h}^{m})^2} =$$

$$= \frac{mar_i^l}{\hat{mar}_i^l} \cdot opt_{i,h} - \frac{1}{\hat{mar}_i^l \cdot ratio_h} \cdot \frac{\frac{1}{n}\sum_{m=1}^{n} \tilde{s_{i,h}^{m}} \cdot c_{i,h}^{m}}{\frac{1}{n}\sum_{m=1}^{n}(c_{i,h}^{m})^2}$$

The term $\frac{mar_i^l}{\hat{mar}_i^l}$ converges to 1 as $n \to \infty$, because the estimator for the marketing effect is consistent. As shown before, the second term converges to 0 as $n \to \infty$. Thus,

$$\lim_{n\to\infty} \hat{opt}_{i,h} = opt_{i,h}$$

making the estimator $\hat{opt}_{i,h} = -\frac{\sum_{m=1}^{n} \hat{mar}_i^l \cdot ratio_h \cdot \tilde{d_{i,h}^{m}} \cdot c_{i,h}^{m}}{\sum_{m=1}^{n}(\hat{mar}_i^l \cdot ratio_h \cdot c_{i,h}^{m})^2}$ consistent.

A.5 Identical Private Value Test

A crucial assumption of the model is that bidders have equal per click valuations for a given auction. In this section, I provide evidence in favor of this assumption, or rather that there is not sufficient evidence against it. The fact that per click valuations are impossible for the researcher to know makes this test difficult to implement. However, the data provide enough information to calculate a reliable proxy for the bidders' per click valuations. Denote the unknown per click valuation of a bidder in a specific auction with v. I calculate a proxy for v: $\tilde{v}$ as follows:

$$\tilde{v} = margin \cdot p \cdot \mathbb{P}(conversion) \cdot (1 - \mathbb{P}(cancellation))$$

where:

- $margin$: The percentage of the transaction value that goes to the bidder from the hotelier.
- p: The booking price that the user sees on the platform.[3]
- $\mathbb{P}(conversion)$: The probability that a user books, given that the user has clicked on the bidder's ad.
- $\mathbb{P}(cancellation)$: The probability that a user cancels her reservation, given a booking.[4]

The data allow for the calculation of the all these variables. However, the real bidder valuations can be affected by unobserved factors, not known to the researcher. I assume these unobserved factors to be random and multiplicative. Therefore, the proxy variable I calculate is a function of the unobserved real valuation, which is deterministic, and the unobserved factor which is random. Thus, the proxy variable is a random variable and its calculated values follow a probability distribution. Given the calculations, the distribution that appears to fit the best is the exponential distribution. The test I perform is testing whether the distribution of proxy valuations for every bidder is likely to be the same.

The steps of the test have as follows:

1. For every bidder i, calculate the set of proxy valuations $\tilde{V}_i$.
2. Derive the MLE estimator, $\hat{\lambda}_i$, for every set $\tilde{V}_i$, assuming that it follows the exponential distribution.
3. Derive the asymptotic distribution of λ_i: $\mathbb{N}(\frac{1}{\hat{\lambda}_i}, \frac{1}{n \cdot \hat{\lambda}_i^2})$.

[3]The price consists of the per night price times the number of nights.

[4]This means that the user (partially) receives her payment back and the bidder is not paid by the hotelier.

bidder	C.I. lower bound	C.I. upper bound
1	0.004303	0.005627
2	0.005026	0.00612
3	0.005406	0.007923
4	0.005186	0.006951
5	0.005502	0.007464
6	0.004861	0.006164
7	0.005262	0.007104

Table A.2: 95% confidence level for the parameter λ

4. For each estimation, $\hat{\lambda}_i$, derive the interval that corresponds to the 95% confidence level.

The results for the dataset for the 7 bidders that appear most often, and discussed in later sections, are displayed in table A.2. It is straightforward to see that the intersection of all intervals is the set $[0.005502, 0.005627]$ and not the empty set. Therefore, the hypothesis that $\lambda_1 = \lambda_2 = \ldots = \lambda_7 = c$, for any $c \in [0.005502, 0.005627]$ would not be rejected at the 95% confidence level. This result adds a level of plausibility to the identical private value assumption.

A.6 Figures

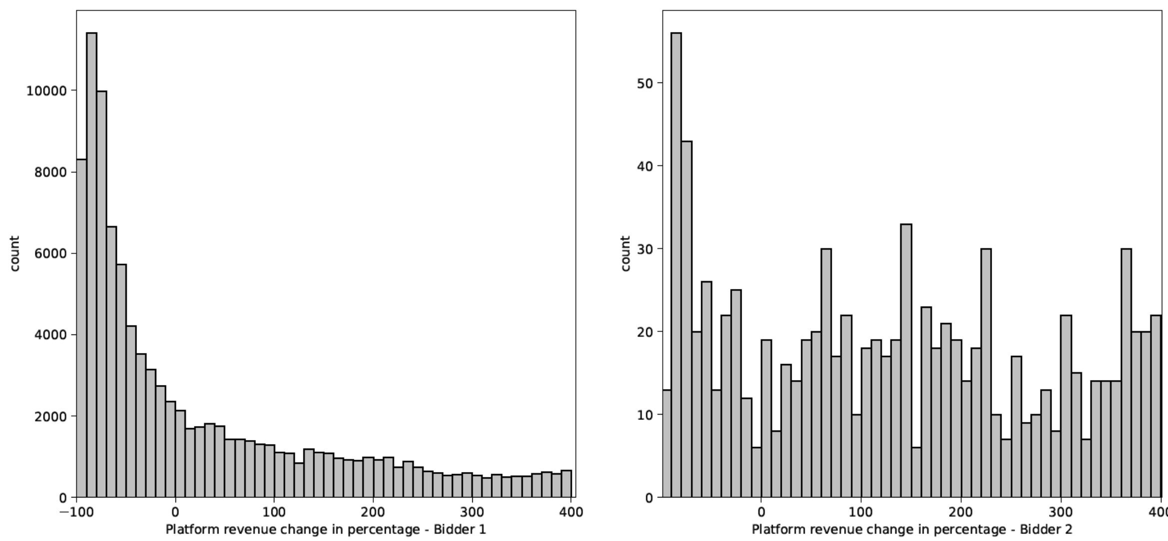

Figure A.2: Platform revenue percentage change for bidders 1-2

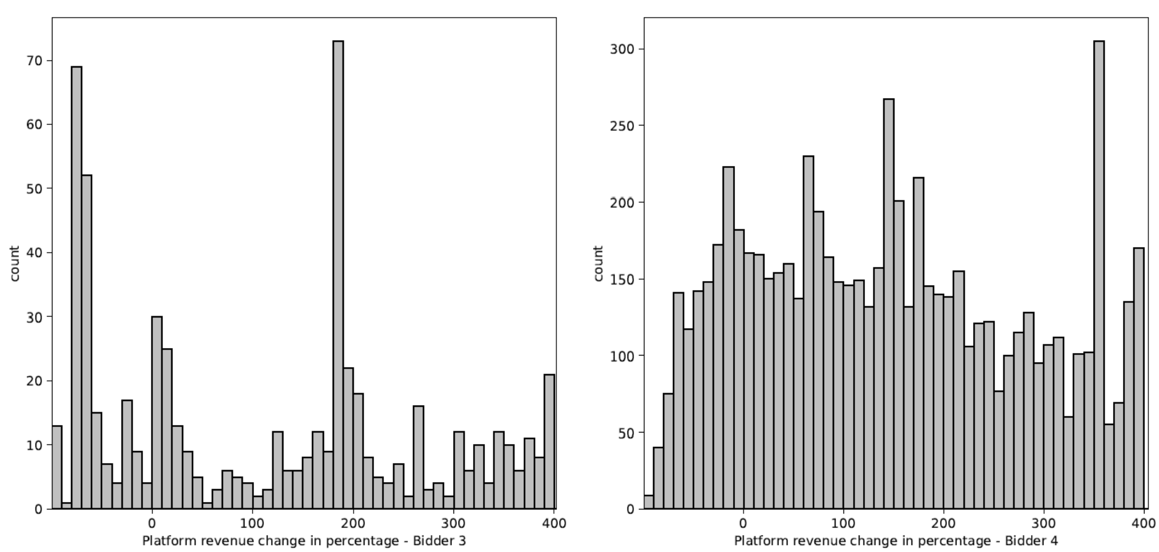

Figure A.3: Platform revenue percentage change for bidders 3-4

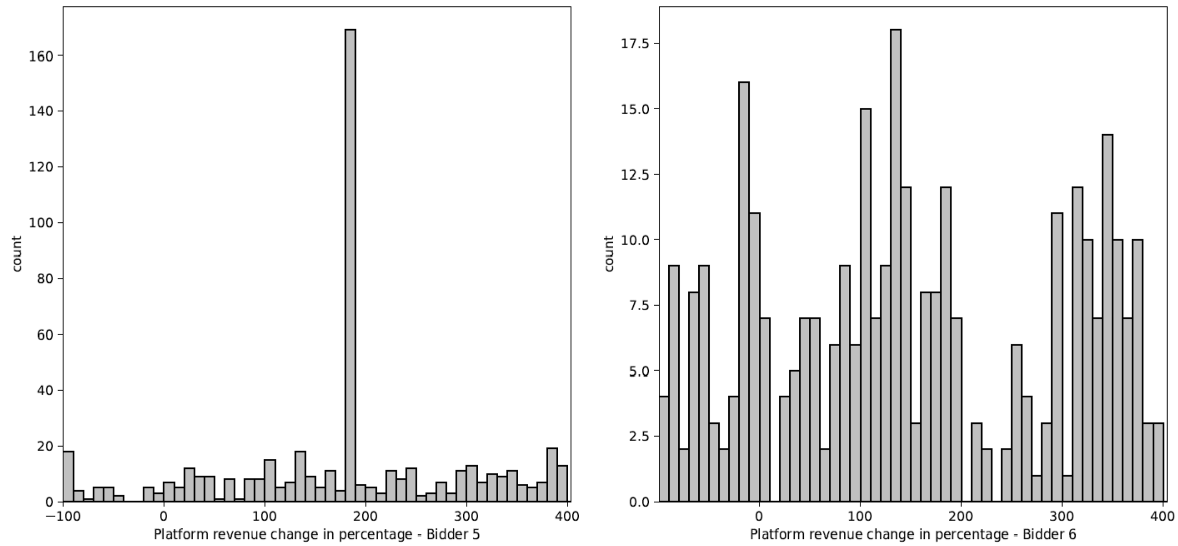

Figure A.4: Platform revenue percentage change for bidders 5-6

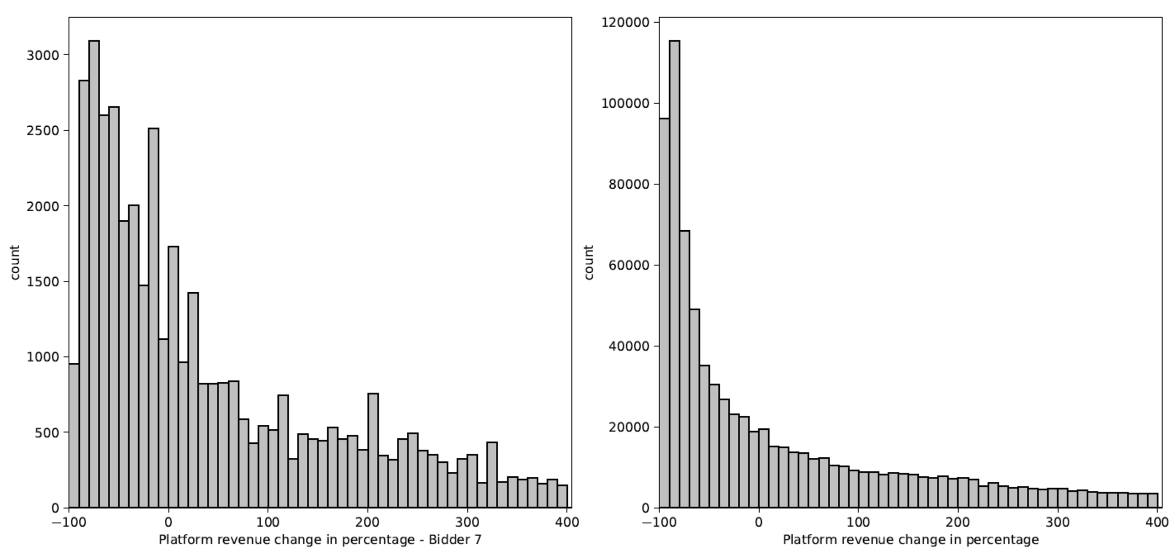

Figure A.5: Platform revenue percentage change for bidder 7 - Aggregate for all bidders

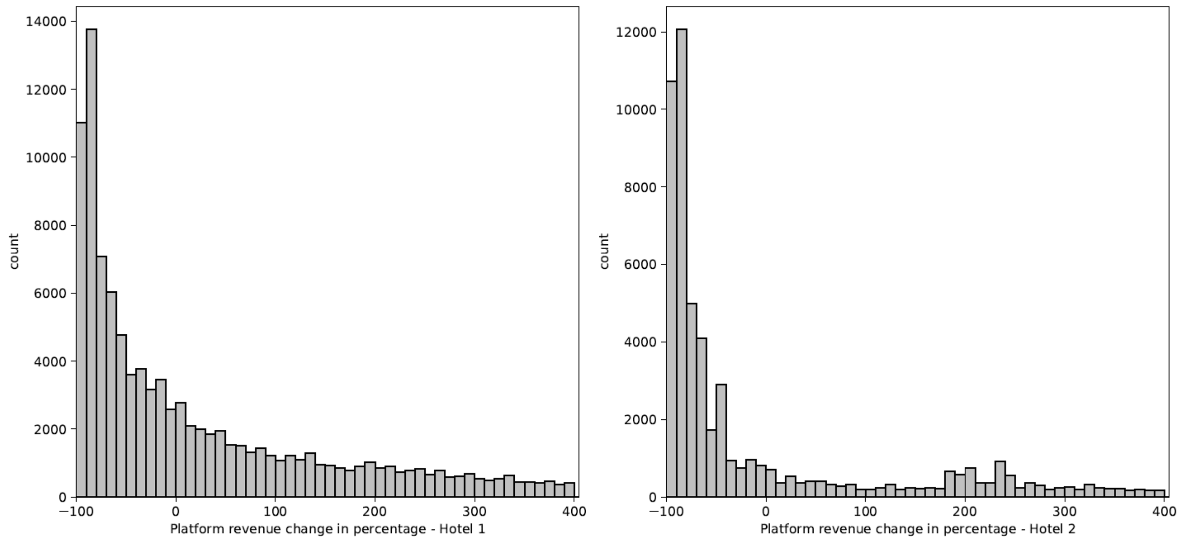

Figure A.6: Platform revenue percentage change for hotels 1-2

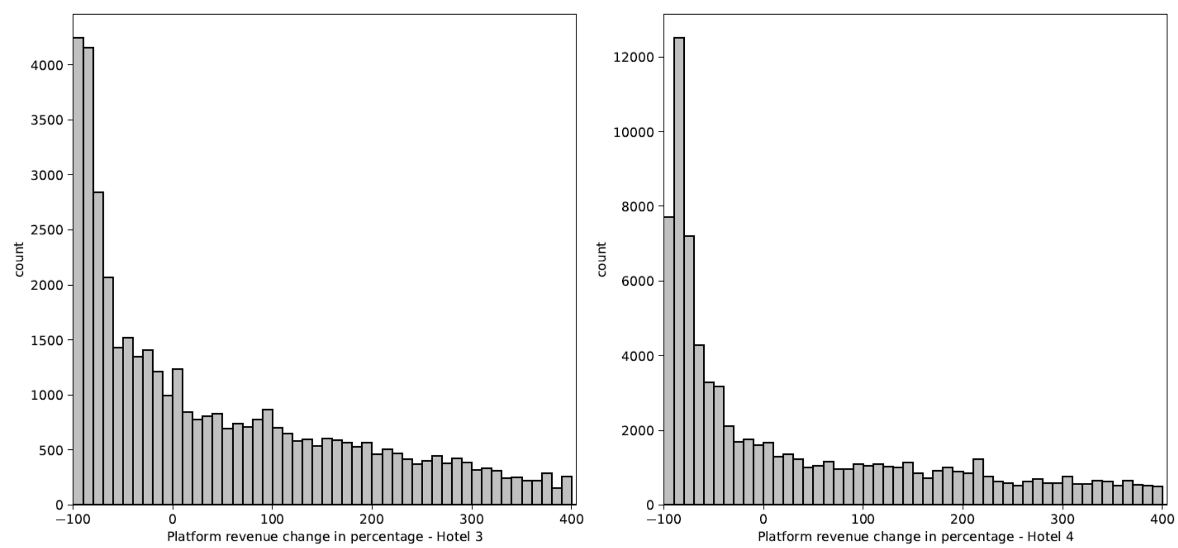

Figure A.7: Platform revenue percentage change for hotels 3-4

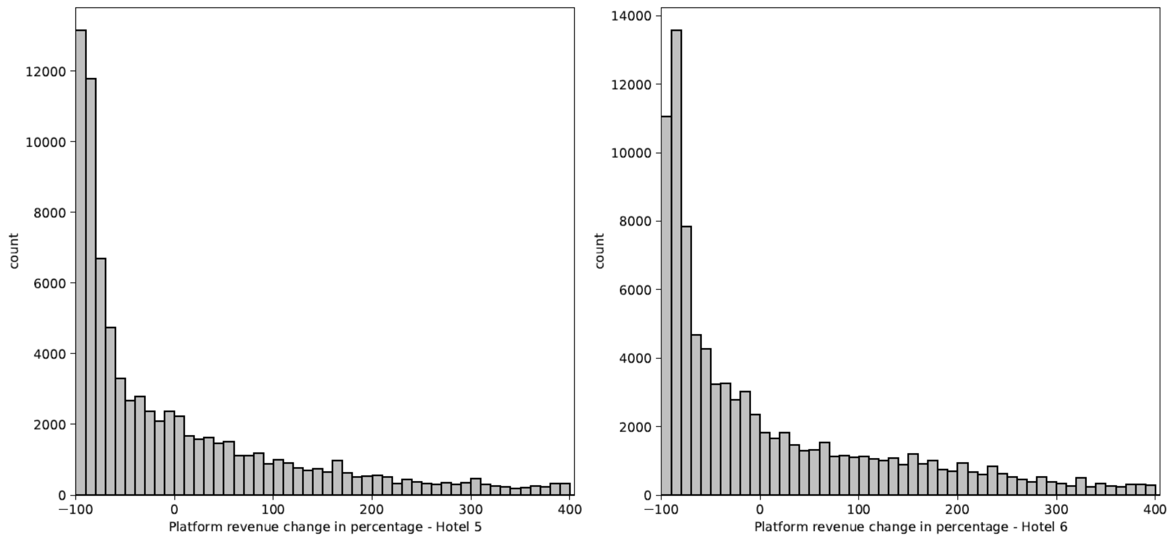

Figure A.8: Platform revenue percentage change for hotels 5-6

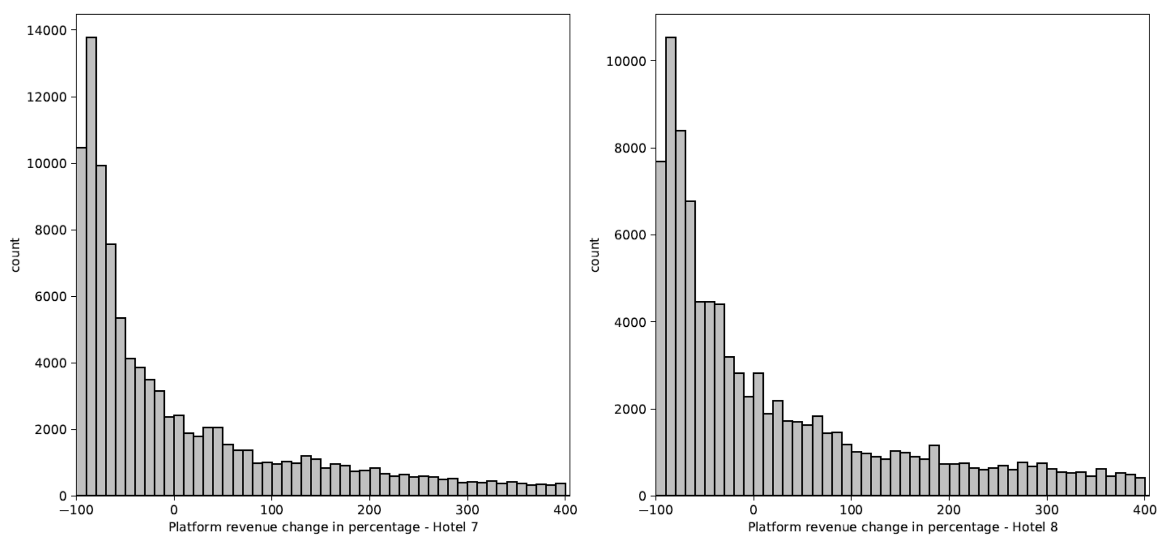

Figure A.9: Platform revenue percentage change for hotels 7-8

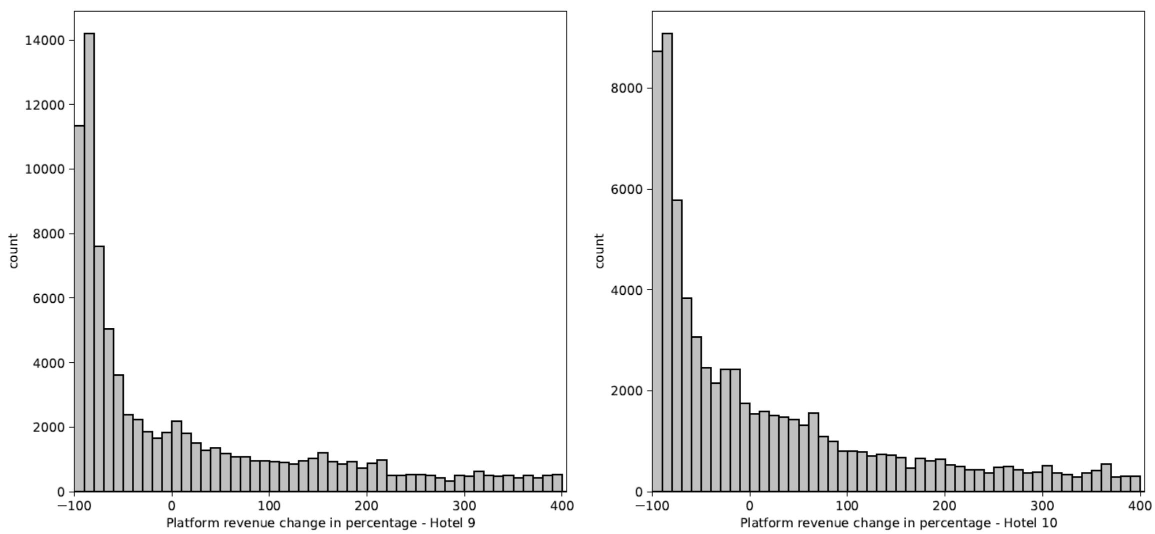

Figure A.10: Platform revenue percentage change for hotels 9-10

Appendix B: The Tragedy of Commons in the Taxi Industry: A Case for a Central Dispatcher

B.1 Proof of Lemma 1

First, I show that the neural network estimator is unbiased at every t. Then, I show that the estimator remains unbiased when transiting from $t+1$ to t.

The neural network solves for the coefficients w in order to solve the following problem:

$$min_w \frac{1}{|S|} \sum_{i=1}^{|S|} [V_t(S_i) - \hat{V}_t(S_i, w)]^2$$

where:

- $|S|$: The number of all possible allocations
- $\hat{V}_t$: The output of the neural network
- V_t : The true value function

Denote the squared errors of the latter problem as $g(w, S_i) = [V_t(S_i) - \hat{V}_t(S_i, w)]^2$

The neural network converges to the solution by using the gradient descent algorithm. The linearity of the neural network output with respect to S and the fact that f is increasing ensures the convexity of g. This method iterates until the result converges, that is:

$$w^k = w^{k-1} - \tau_k \cdot \nabla\left(\tfrac{1}{|S|}\textstyle\sum_{i=1}^{|S|} g(w^{k-1}, S_i)\right)$$

until $w^k \approx w^{k-1} \Rightarrow \nabla\left(\tfrac{1}{|S|}\sum_{i=1}^{|S|} g(w^{k-1}, S_i)\right) \approx 0$

The latter equation is not implementable in computational terms because of the large number of total allocations, $|S|$. Instead of considering all allocations, I randomly sample K allocations, and then I calculate:

$$w^k = w^{k-1} - \tau_k \cdot \nabla g(w^{k-1}, S_k)$$

for $k = 1, 2, ..., K$. Note that K is big enough to ensure convergence, i.e. I draw samples until convergence is achieved. The samples are drawn randomly, thus:

$$E_S[\nabla g(w^{k-1}, S_k)] = \tfrac{1}{|S|}\textstyle\sum_{i=1}^{|S|} \nabla g(w^{k-1}, S_i) = \nabla\left(\tfrac{1}{|S|}\textstyle\sum_{i=1}^{|S|} g(w^{k-1}, S_i)\right)$$

The latter suggests that the expression $\nabla g(w^{k-1}, S_k)$ is equal to the gradient of the squared errors in expectation, hence:

$$E_S[V_t(S) - \hat{V}_t(S, \hat{w})] = 0 \qquad (I)$$

Then, I show by induction that the estimator remains unbiased when transiting from $t+1$ to t.

For $t = T$ it holds that $E_S[V_T(S) - \hat{V}_T(S, \hat{w})] = 0 \qquad (II)$ as equation (I) suggests.

For $t < T$, assume that $E_S[V_{t+1}(S) - \hat{V}_{t+1}(S)] = 0$. I show that the same holds for t. The function $V_t(S_k, x_k) = TW_t(S_k, x_k) + V_{t+1}(S'_k)$ cannot be estimated using the neural network, as the function V_{t+1} is not known. Instead, I estimate the function $\tilde{V}_t(S_k, x_k) = TW_t(S_k, x_k) + \hat{V}_{t+1}(S'_k)$ from which I can draw samples. By (I):

$$E_S[\tilde{V}_t(S) - \hat{V}_t(S)] = 0 \Rightarrow$$

$$E_S[TW_t(S)] + E_S[\hat{V}_{t+1}(S')] - E_S[\hat{V}_t(S)] = 0 \overset{\text{(II)}}{\Longrightarrow}$$

$$E_S[TW_t(S)] + E_S[V_{t+1}(S')] - E_S[\hat{V}_t(S)] = 0 \Rightarrow$$

$$E_S[TW_t(S) + V_{t+1}(S') - \hat{V}_t(S)] = 0 \Rightarrow$$

$$E_S[V_t(S) - \hat{V}_t(S)] = 0 \qquad \blacksquare$$

B.2 Figures

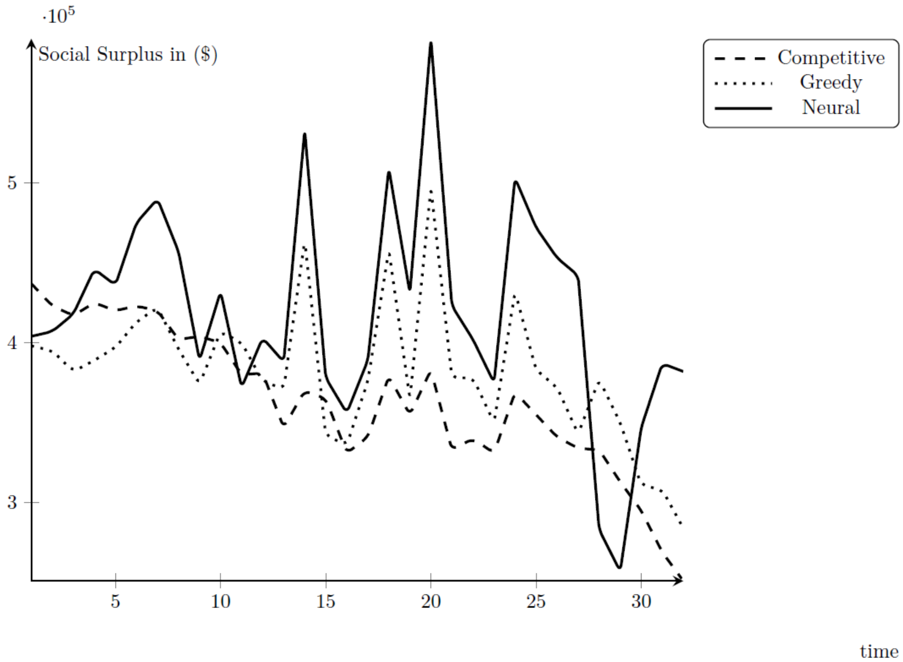

Figure B.1: Social Surplus

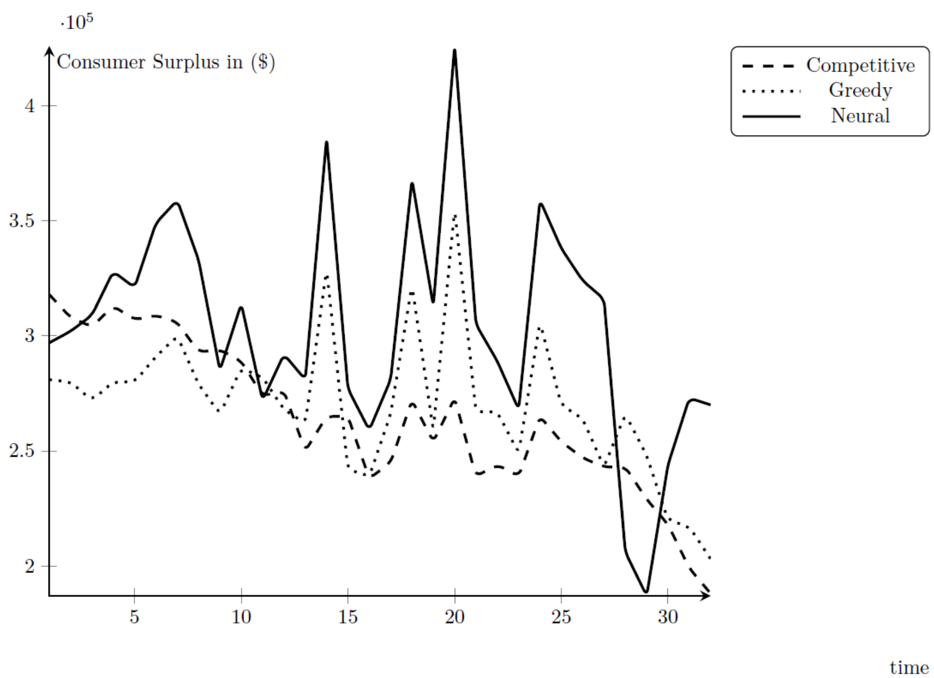

Figure B.2: Consumer Surplus

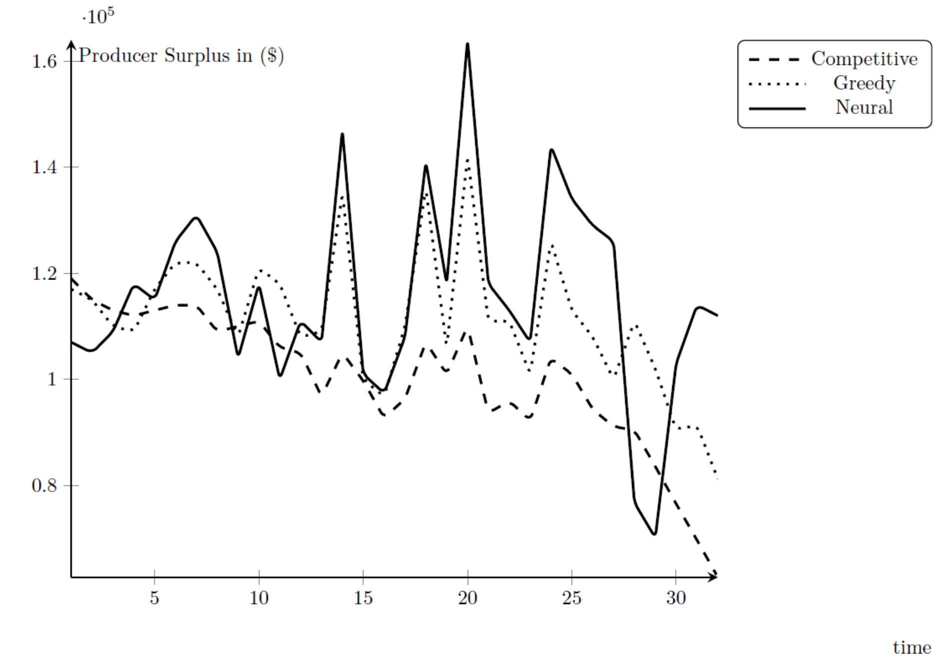

Figure B.3: Producer Surplus

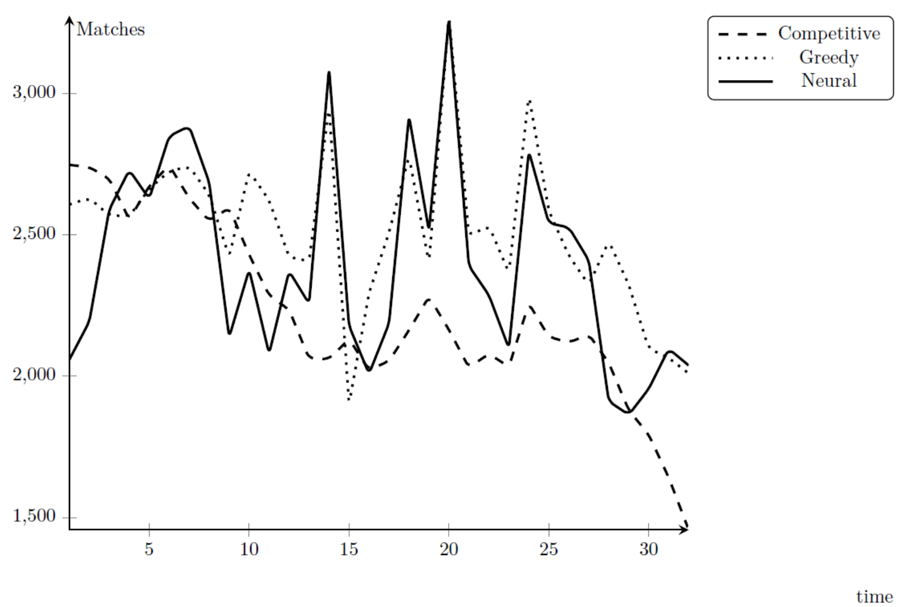

Figure B.4: Matching, Lower Manhattan

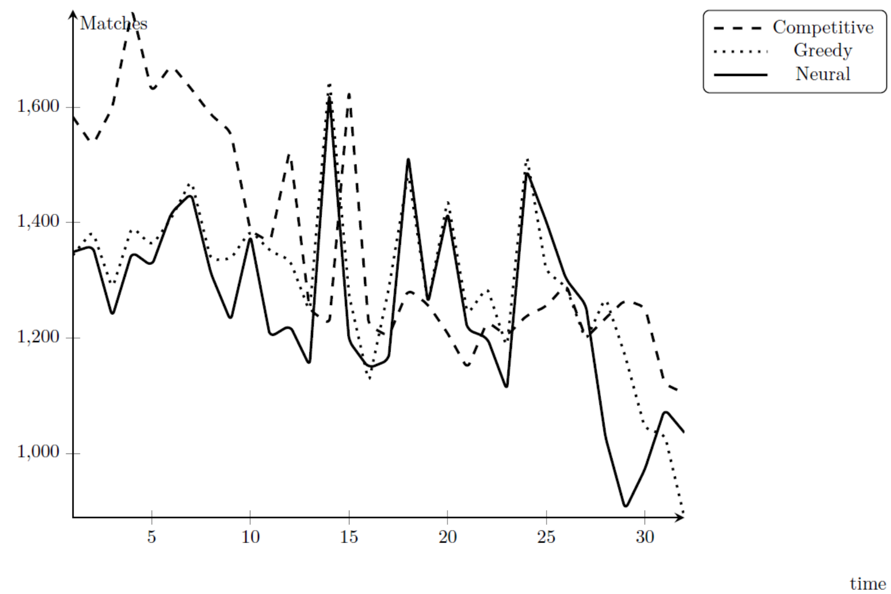

Figure B.5: Matching, Mid Manhattan

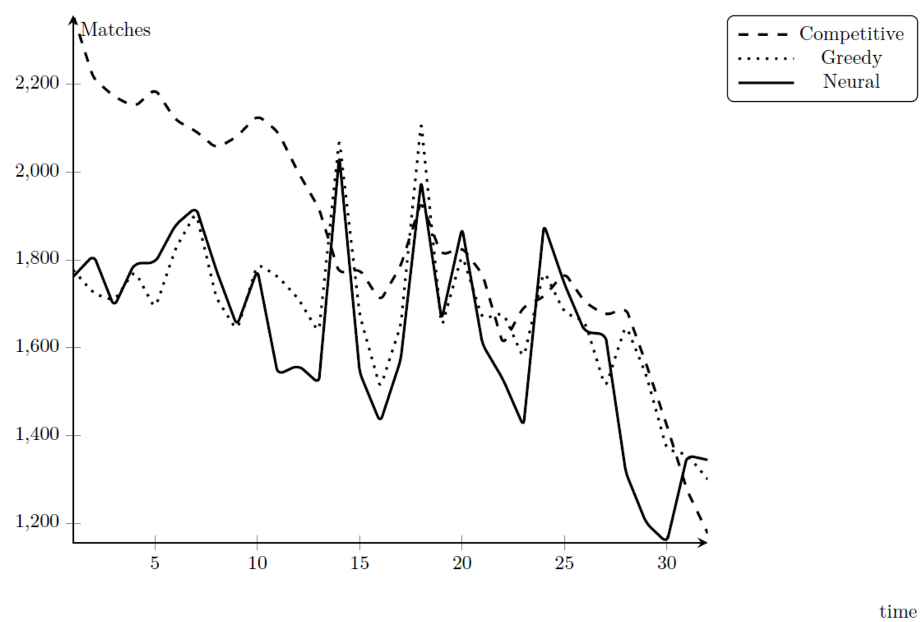

Figure B.6: Matching, Upper Manhattan

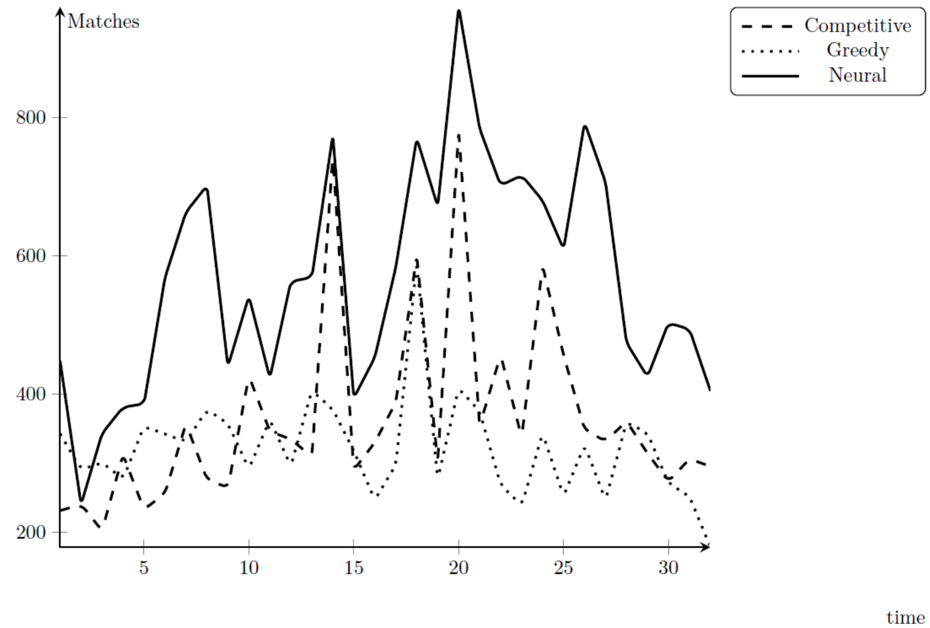

Figure B.7: Matching, Boroughs

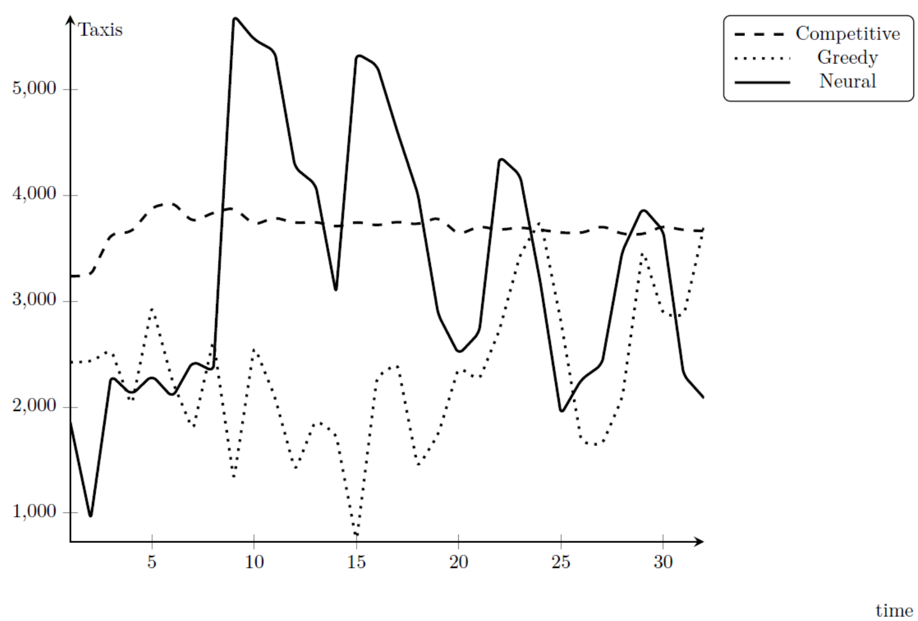

Figure B.8: Vacant taxis, Lower Manhattan

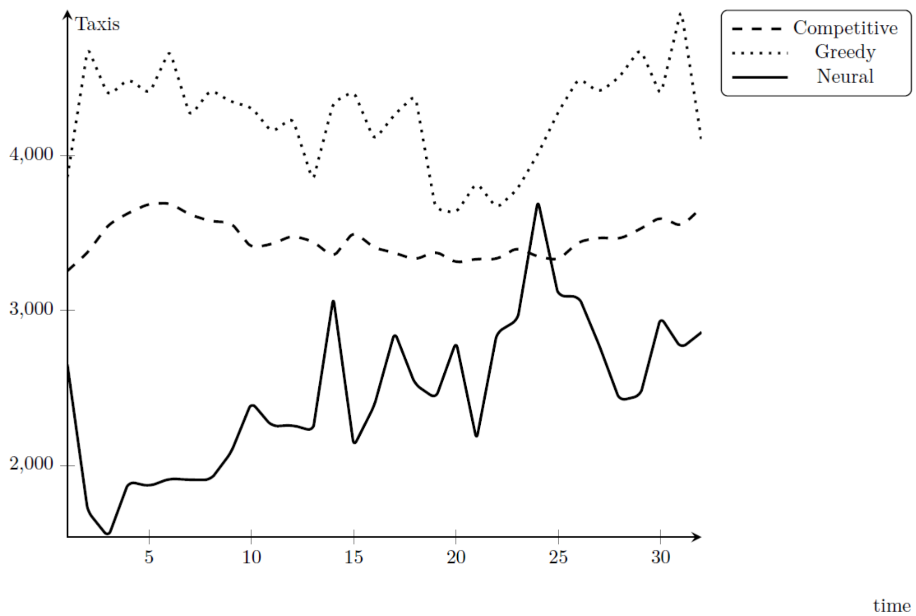

Figure B.9: Vacant taxis, Mid Manhattan

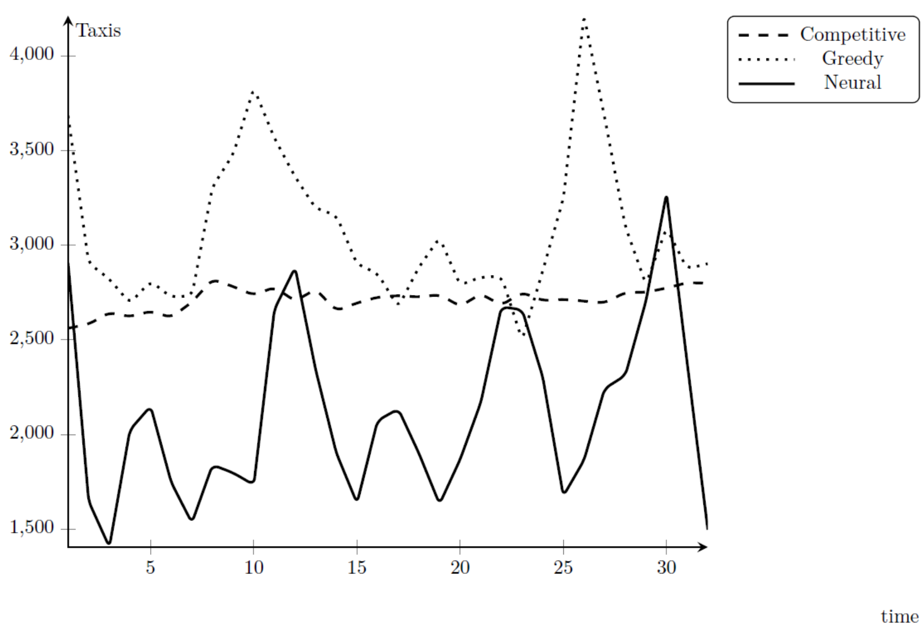

Figure B.10: Vacant taxis, Upper Manhattan

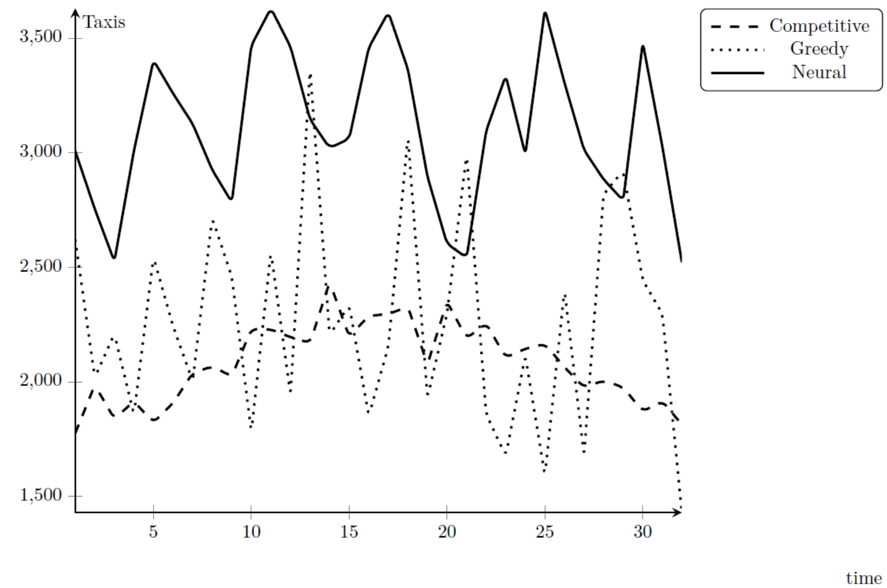

Figure B.11: Vacant taxis, Boroughs

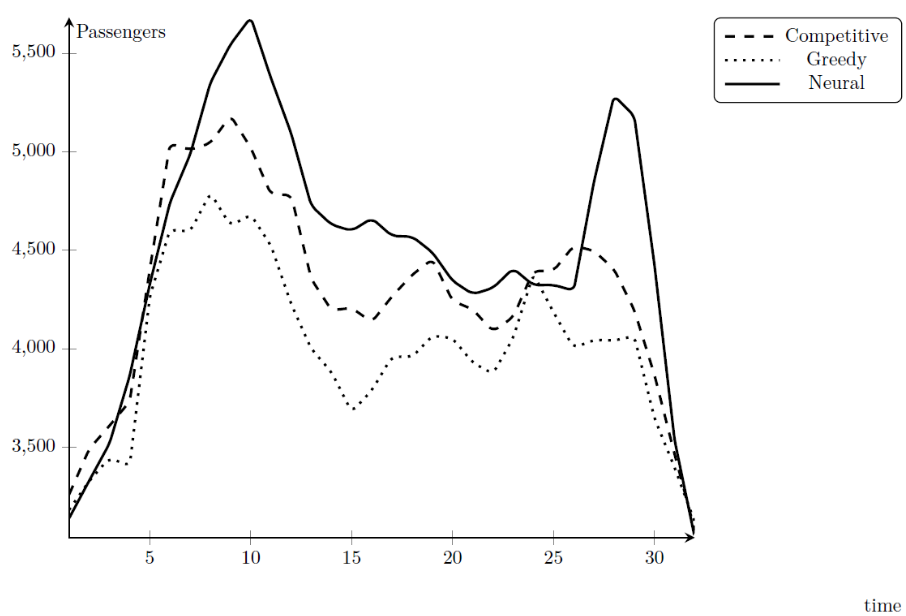

Figure B.12: Passengers, Lower Manhattan

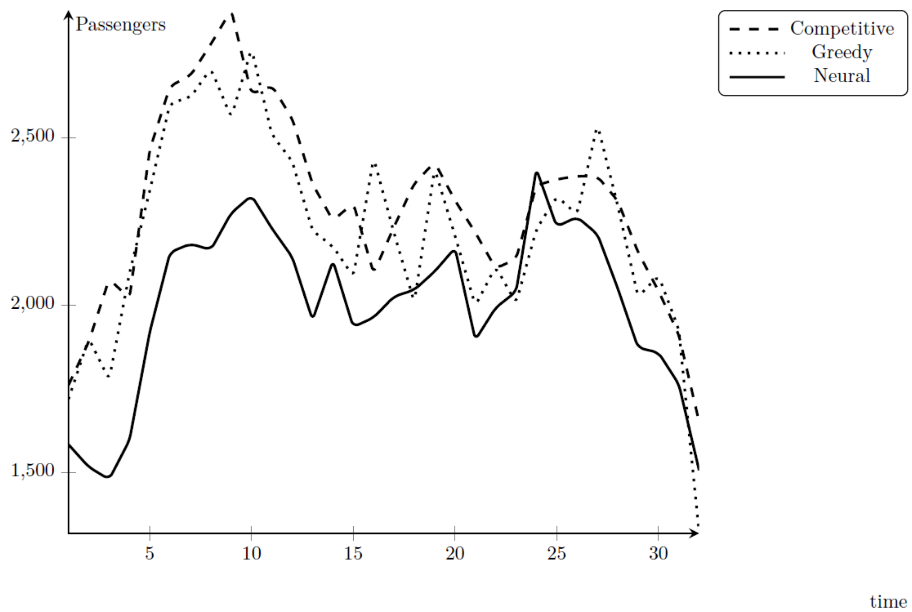

Figure B.13: Passengers, Mid Manhattan

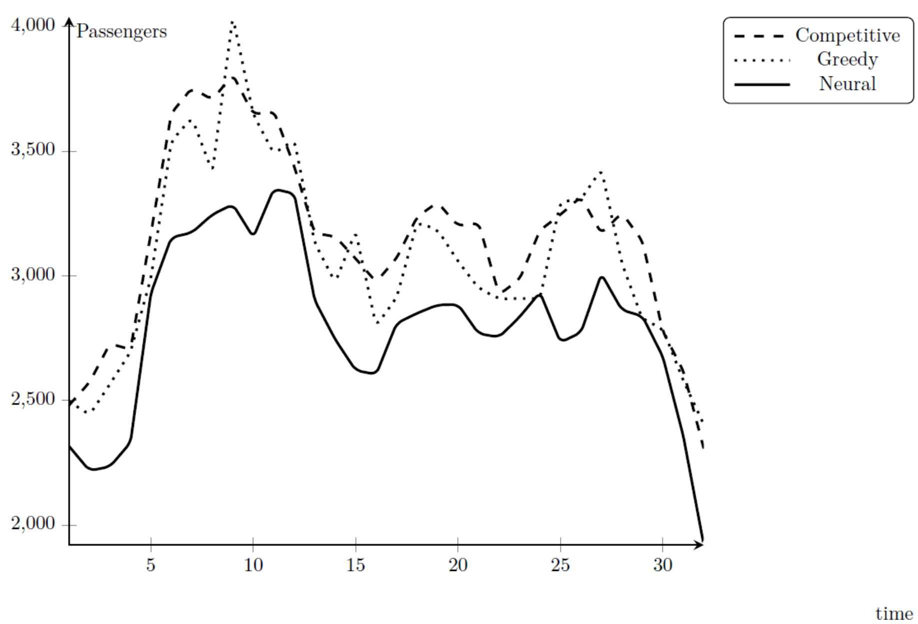

Figure B.14: Passengers, Upper Manhattan

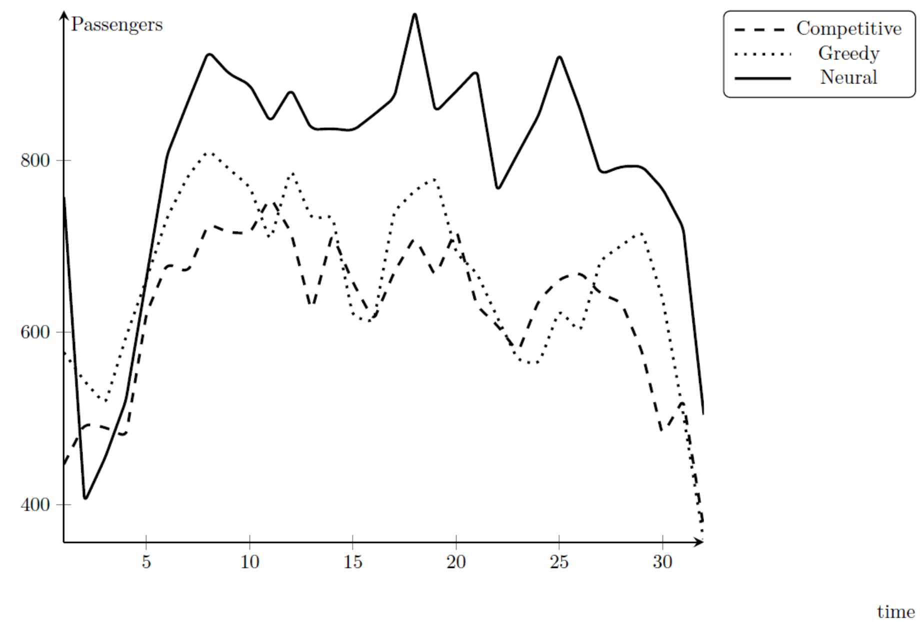

Figure B.15: Passengers, Boroughs Manhattan

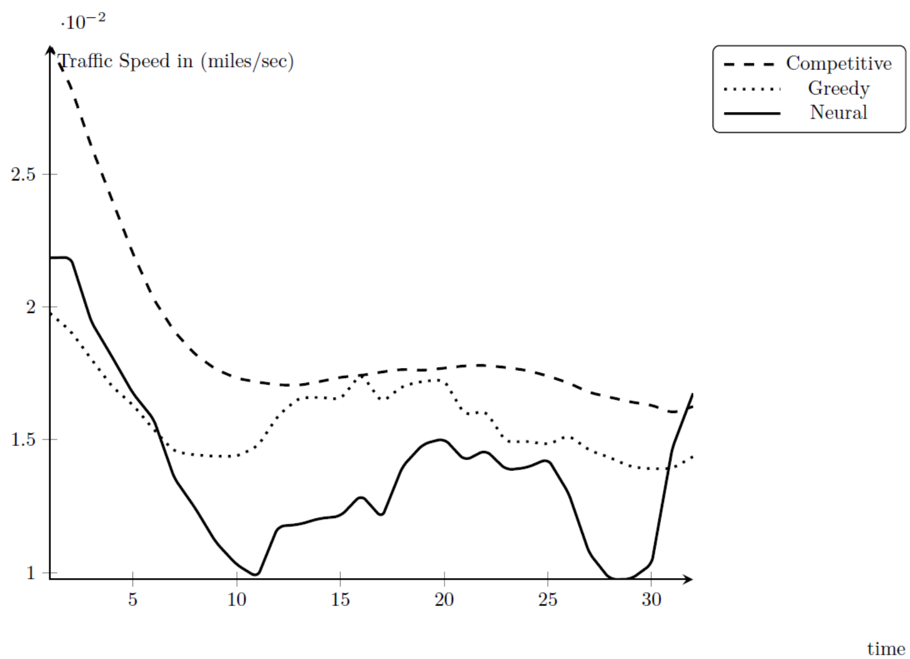

Figure B.16: Traffic Speed, Lower Manhattan

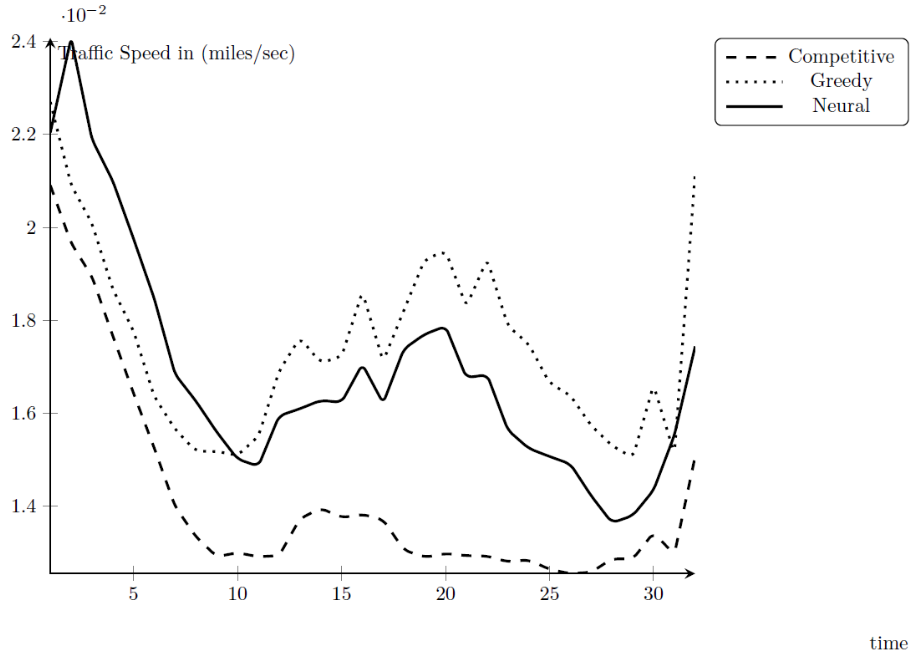

Figure B.17: Traffic Speed, Mid Manhattan

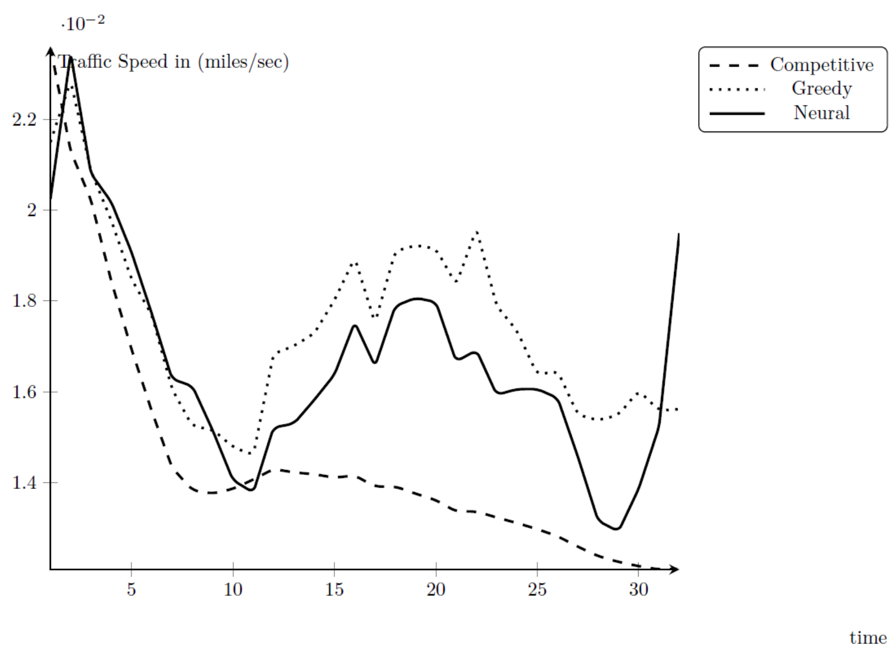

Figure B.18: Traffic Speed, Upper Manhattan

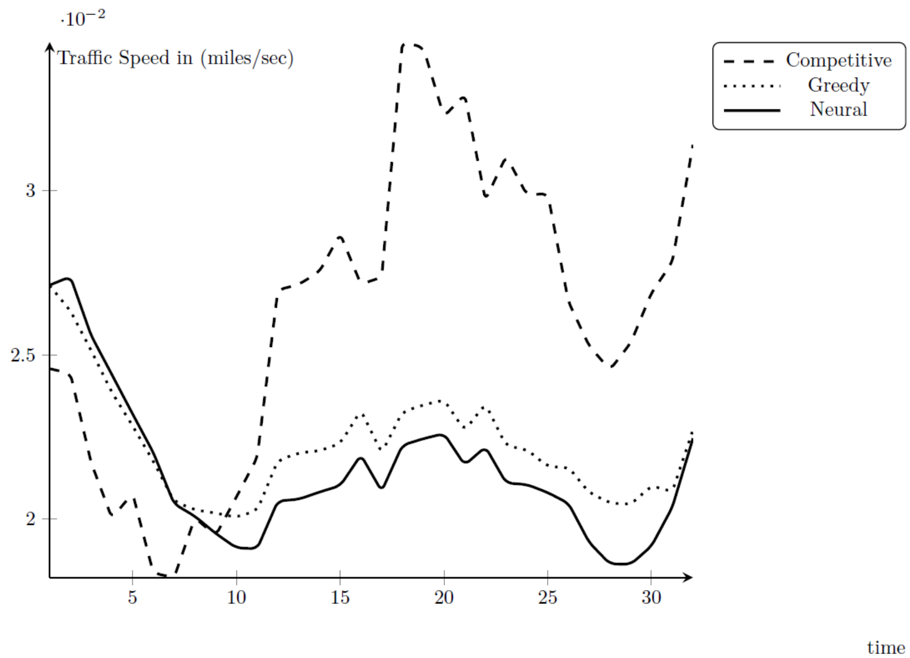

Figure B.19: Traffic Speed, Boroughs

Bibliography

[1] Armstrong, Mark (2000), ”Optimal Multi-Object Auctions,” Review of Economic Studies, 67(3), pp. 455–481.

[2] Asdemir, K. (2006): “Bidding Patterns in Search Engine Auctions.”, Paper presented at the Second Workshop on Sponsored Search Auctions, Ann Arbor, MI.

[3] Athey, Susan, Peter Cramton, and Allan Ingraham (2002), ”Setting the Upset Price in British Columbia Timber Auctions,” Working paper.

[4] Athey, Susan, and Glenn Ellison (2011), ”Position Auctions with Consumer Search,” Quarterly Journal of Economics, 126(3), pp. 1213–1270.

[5] Bajari, Patrick, and Ali Hortaçsu (2003), ”The Winner’s Curse, Reserve Prices, and Endogenous Entry: Empirical Insights from eBay Auctions,” RAND Journal of Economics, 34(2), pp. 329–355.

[6] Berry, S., Waldfogel, J. (1999): “Free Entry and Social Inefficiency in Radio Broadcasting”, RAND Journal of Economics, 30(3): 397-420.

[7] Bian, B. (2019): “Search Frictions, Network Effects and Spatial Competition: Taxis versus Uber”, Working Paper.

[8] Bie, Y., Seraj, M., Qiu, T. and Wang, X. (2016): “Implementation of Variable Speed Limits: Preliminary Test on Whitemud Drive, Edmonton, Canada”, Journal of Transportation Engineering Volume 142 Issue 12.

[9] Brancaccio, G., M. Kalouptsidi, and T. Papageorgiou (2019): “Geography, Transportation, and Endogenous Trade Costs”, Econometrica.

[10] Brancaccio, G., M. Kalouptsidi, T. Papageorgiou, and N.Rosaia (2020): ”Search Frictions and Efficiency in Decentralized Transport Markets”, NBER working paper 27300.

[11] Brennan, A. (2014): “IBIS World Industry Report 48533 Taxi & Limousine Services in the US”, Tech. rep., IBIS World.

[12] Brophy, T. (2013): “GPSmap: Routine for verifying and returning the attribute table of given decimal degree GPS coordinates”, in 2013 Stata Conference, Stata Users Group, 14.

[13] Buchholz, N. (2020): “Spatial Equilibrium, Search Frictions and Dynamic Efficiency in the Taxi Industry”, Working Paper.

[14] Cairns, R. D. and C. Liston-Heyes (1996): “Competition and regulation in the taxi industry”, Journal of Public Economics, 59, 1–15.

[15] Castillo, J. C., D. Knoepfle, and G. Weyl (2017): “Surge pricing solves the wild goose chase”, in Proceedings of the 2017 ACM Conference on Economics and Computation, ACM, 241–242.

[16] Dimitrellos, Panos (2020), ”A Field Experiment on Hybrid and Second Price Rules in Internet Advertising Auctions,” working paper.

[17] Dimitrellos, Panos (2021a), ”Information, Bias, and Revenues in Sponsored Search Auctions,” working paper.

[18] Dimitrellos, Panos (2021b): “Optimal Reserve Prices with Endogenous Demand: A Field Experiment”, working paper.

[19] Edelman, B. and Ostrovsky, M. (2006): “Strategic bidder behavior in sponsored search auctions”, Decision Support Systems 43 (2007) 192–198.

[20] Edelman, B., Ostrovsky, M. and Schwarz, M. (2007): “Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords”, American Economic Review vol. 97, no. 1, March 2007.

[21] Edelman, B. and Schwarz, M. (2010): “Optimal Auction Design and Equilibrium Selection in Sponsored Search Auctions”, American Economic Review vol. 100, no. 2, May 2010.

[22] Ericson, R. and A. Pakes (1995): “Markov-perfect industry dynamics: A framework for empirical work”, The Review of Economic Studies, 62, 53–82.

[23] Frechette, G. R., Lizzeri, A., and Salz, T. (2018): “Frictions in a competitive, regulated market: Evidence from taxis”, American Economic Review, 109, 2954–92.

[24] Haile, Philip A., and Elie Tamer (2003), ”Inference with an Incomplete Model of English Auctions,” Journal of Political Economy, 111(1), pp. 1–51.

[25] Hall, J., C. Kendrick, and C. Nosko (2015): “The effects of Uber’s surge pricing: A case study”, The University of Chicago Booth School of Business.

[26] Iyengar, G., Kumar, A. (2006): “Characterizing Optimal Adword Auctions”,Second Workshop on Sponsored Search Auctions, June 11, 2006 Ann Arbor, Michigan in conjunction with the ACM Conference on Electronic Commerce (EC’06).

[27] Jeziorskiy, Przemyslaw, and Ilya Segal (2015), ”What Makes Them Click: Empirical Analysis of Consumer Demand for Search Advertising,” American Economic Journal: Microeconomics, 7(3), pp. 24–53.

[28] Lahaie, S. (2006): “An Analysis of Alternative Slot Auction Designs for Sponsored Search”, EC’06, June 11–15, 2006, Ann Arbor, Michigan, USA.

[29] Lagos, R. (2000): “An Alternative Approach to Search Frictions”, The Journal of Political Economy, 108, 851–873.

[30] Lagos, R. (2003): “An analysis of the market for taxicab rides in New York City”, International Economic Review, 44, 423–434.

[31] Mangrum, D., Molnar, A. (2017): “The Marginal Congestion of a Taxi in New York City”, Working Paper. Shapiro, M. (2018): “Density of Demand and the Benefit of Uber”, Working Paper.

[32] Maskin, Eric, and John Riley (1984), ”Optimal Auctions with Risk Averse Buyers,” Econometrica, 52(6), pp. 1473–1518.

[33] McAfee, R. Preston, Daniel C. Quan, and Daniel R. Vincent (2002), ”How to Set Minimum Acceptable Bids, with Application to Real Estate Auctions,” Journal of Industrial Economics, 50(4), pp. 391–416.

[34] McAfee, R. Preston, and Daniel R. Vincent (1992), ”Updating the Reserve Price in Common Value Auctions,” American Economic Review Papers and Proceedings, 82(2), pp. 512–518.

[35] Myerson, Roger B. (1981), ”Optimal Auction Design,” Mathematics of Operations Research, 6(1), pp. 58–73.

[36] Ostrovsky, Michael and Michael Schwarz (2016), ”Reserve Prices in Internet Advertising Auctions: A Field Experiment,” revise and resubmit, Journal of Political Economy.

[37] Parkes, D. and Singh, S. (2004): “An MDP-Based Approach to Online Mechanism Design”, In Advances in neural information processing systems 16: Proceedings of the 2003 conference.

[38] Rafieian, Omid (2020), ”Revenue-Optimal Dynamic Auctions for Adaptive Ad Sequencing,” working paper.

[39] Reiley, David (2006), ”Field Experiments on the Effects of Reserve Prices in Auctions: More Magic on the Internet.” RAND Journal of Economics, 37(1), pp. 195–211.

[40] Riley, John G., and William F. Samuelson (1981), ”Optimal Auctions,” American Economic Review, 71(3), pp. 381–392.

[41] Roughgarden, Tim, and Mukund Sundararajan (2007), ”Is Efficiency Expensive?,” Proceedings of the 3rd Workshop on Sponsored Search Auctions, WWW2007.

[42] Said, M. (2012): “Auctions with dynamic populations: Efficiency and revenue maximization”, Journal of Economic Theory, Elsevier, vol. 147(6), pages 2419-2438.

[43] Sun, Yang, Yunhong Zhou, and Xiaotie Deng (2014), ”Optimal Reserve Prices in Weighted GSP Auctions,” Electronic Commerce Research and Applications, 13(3), pp. 178–187.

[44] Tang, Xun (2011), ”Bounds on Revenue Distributions in Counterfactual Auctions with Reserve Prices,” RAND Journal of Economics, 42(1), pp. 175–203.

[45] Topinsky, Valery (2014), Reserve Prices in Asymmetric Auctions, Ph.D. dissertation, Higher School of Economics, Moscow. Available at http://www.ccas.ru/avtorefe/0023d.pdf.

[46] Ulku, L. (2013): “Optimal combinatorial mechanism design”, Economic Theory, Springer; Society for the Advancement of Economic Theory (SAET), vol. 53(2), pages 473-498, June.

[47] Vickrey, William (1961), ”Counterspeculation, Auctions, and Competitive Sealed Tenders,” Journal of Finance, 16(1), pp. 8–37.

[48] Weintraub, G. Y., C. L. Benkard, P. Jeziorski, and B. Van Roy (2008a): “Nonstationary Oblivious Equilibrium”, Manuscript, Columbia University.

[49] Wong, J. (2019): “E-hail, Flexible Fare and Leasing: Improving the New York Taxicab Market”, Working Paper.

[50] Yan, H. (2019): “Position Auctions with Interdependent Values”, working paper