

THE UNIVERSITY OF CHICAGO

THE POWER OF RANKINGS: QUANTIFYING THE EFFECT OF RANKINGS ON
ONLINE CONSUMER SEARCH AND PURCHASE DECISIONS

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE SOCIAL SCIENCES
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

DEPARTMENT OF ECONOMICS

BY

RALUCA MIHAELA URSU

CHICAGO, ILLINOIS

JUNE 2016

Copyright © 2016 by Raluca Mihaela Ursu
All Rights Reserved

Parintilor mei, Maria si Ioan Ursu, si surorii mele, Oana Ursu

TABLE OF CONTENTS

LIST OF FIGURES	vi
LIST OF TABLES	vii
ABSTRACT	viii
ACKNOWLEDGMENTS	ix
1 INTRODUCTION	1
2 RELATED LITERATURE	4
3 THE ONLINE TRAVEL AGENT INDUSTRY	7
4 DATA	9
4.1 Search Process	9
4.2 Data Description	11
4.3 Pros and Cons of the Data Set	12
5 REDUCED FORM EVIDENCE	15
5.1 Causal Effect of Rankings on Search and Purchase Decisions	15
5.2 Mechanism Driving the Effect of Rankings on Search	18
6 MODEL	21
6.1 Utility	21
6.2 Search Cost	21
6.3 Optimal Search	22
7 ESTIMATION	25
8 IDENTIFICATION	29
8.1 Model Parameters	29
8.2 Monte Carlo Simulation	32
9 RESULTS	34
10 COUNTERFACTUAL: UTILITY-BASED RANKING	37
11 LIMITATIONS AND FUTURE RESEARCH	40
12 CONCLUSION	41
REFERENCES	42
A DATA CLEANING	45

B	FURTHER EVIDENCE	46
B.1	The Experiment	46
B.2	Ranking Algorithm	49
B.3	The WCAI Data Set	50
B.4	Position Effect under the Random Ranking	53
B.5	Position Effect under Expedia’s Ranking	54
B.6	Illustration of Opaque Offers	55
B.7	Evidence on Consumer Click Order	56
B.8	Price Endogeneity Concerns	57
B.9	Summary Statistics for the Largest Destinations in the Data	58
B.10	Magnitude of Search Costs	59

LIST OF FIGURES

4.1	Consumer Search Process Illustration	10
5.1	Position Effect: Random Ranking	16
5.2	Position Effect: Expedia Ranking	17
B.1	Search Impressions with the Random Ranking	47
B.2	Position Effect for Destinations with Few Hotels	53
B.3	Position Effect for Search Impressions Longer than the Median	53
B.4	Position Effect for Search Impressions Ending in a Transaction	53
B.5	Illustration of Opaque Offers	55

LIST OF TABLES

4.1	Hotel and Search Impression Summary Statistics	12
5.1	Estimates of Clicks and Transactions (OLS): Random Ranking	17
5.2	Estimates of Click on the Fifth Displayed Hotel (OLS)	20
8.1	Monte Carlo Simulation Results	33
9.1	Main Estimation Results	35
10.1	Counterfactual Results	38
B.1	Comparison of Consumer Observables: Expedia versus Random Ranking (T-test)	48
B.2	Estimates of Position by Ranking (Rank Ordered Logit)	49
B.3	Summary Statistics Comparison with the WCAI Data Set (Mean)	52
B.4	Estimates of Clicks and Transactions (OLS): Expedia Ranking	54
B.5	Evidence on Consumer Click Order from the WCAI Data Set	56
B.6	Estimates of Hotel Prices (OLS)	57
B.7	Hotel and Search Impression Summary Statistics: Largest Destinations (Mean)	58
B.8	Estimation Results: Search Impressions with at least Two Clicks	60
B.9	WCAI Data set: Number of Clicks at Search Impression and Consumer Level .	61
B.10	WCAI Data Set: Estimation Results	61

ABSTRACT

Online search intermediaries, such as Amazon or Expedia, use rankings (ordered lists) to present third party sellers' products to consumers. These rankings decrease consumer search costs and increase the probability of a match with a seller, ultimately increasing consumer welfare. Constructing relevant rankings requires understanding their causal effect on consumer choices. However, this is challenging since rankings are endogenous: highly ranked products are also the most relevant ones for consumers. In this dissertation, I use the first data set with experimental variation in the ranking from a field experiment at Expedia to identify the causal effect of rankings. Using this data set, I make three contributions. First, I show that rankings affect what consumers search, but conditional on search, do not affect purchases. I also exploit a feature of the data set (opaque offers), to show that rankings lower search costs, instead of affecting consumer expectations or utility. Second, I quantify the effect of rankings using a sequential search model and find an average position effect of \$2.64, lower than previous estimates in the literature obtained without experimental variation. Finally, I show that a utility-based ranking built on this model's estimates leads to an almost twofold increase in consumer welfare, while also increasing transactions and revenue for the search intermediary.

ACKNOWLEDGMENTS

My dissertation was possible only because of the support and encouragement of many people. First, I would like to thank my advisors and thesis committee members, Professor Ali Hortaçsu, Professor Pradeep Chintagunta, Professor Hugo Sonnenschein, and Professor Richard Van Weelden, for their guidance and advice throughout my PhD. I had the honor of working with Professor Ali Hortaçsu, who has always encouraged me to work on big ideas and to strive for improvement. His vision, dedication and drive have led me on the trail to this point. I met Professor Pradeep Chintagunta during my third year in the program as I was exploring possible research questions for my dissertation. From early on, his constant guidance at every step of the way has taught me a lot of what I know about research. I am indebted to him for always believing in me and my work, even during times when I lost confidence in myself. Professor Hugo Sonnenschein has been a mentor for me from the first day I arrived at the University of Chicago, helping me navigate through graduate student life. When I went on the job market, he advised me to be “honest, transparent, and grateful”, lesson which I will carry with me all my life. I have also had the good fortune of working on the greater part of this dissertation with Professor Richard Van Weelden. He was instrumental in helping me develop the theoretical tools I used in my dissertation.

I greatly benefited from lengthy discussions about my dissertation with Professor Paulo Albuquerque, Professor Matthew Backus, Naiqing Gu, Professor Elisabeth Honka, Sergei Koulayev, Chris Nosko, Professor Anita Rao, Professor Stephan Seiler, Professor Bradley Shapiro, as well as the faculty members in the University of Chicago Department of Economics and the University of Chicago Booth School of Business Marketing Department. Throughout my years at the University of Chicago, I had the pleasure to interact with and learn from my classmates, especially Anushree Subramaniam and Sarah Komisarow. With them, I shared my concerns, struggles and hopes. For this, I thank them.

I also thank the participants at the Searle Conference on Internet Search and Innovation (June 2015, Chicago), the Workshop on Search and Switching Costs (May 2015, Groningen),

the IIOC conference (April 2015, Boston), the University of Chicago IO and Marketing working groups for all their feedback, as well as all the people I met during the job market for their questions and comments that helped improve my working paper that eventually became this dissertation.

I gratefully acknowledge the data sets I obtained from Expedia, Kaggle and the Wharton Customer Analytics Initiative and the financial support received from the Social Sciences Fellowship at the University of Chicago, the Theodore W. and Esther Schultz Economics Fellowship at the University of Chicago, and the Marketing Science Institute Alden G. Clayton Dissertation Competition.

At a personal level, I thank my parents, Maria and Ioan Ursu, my sister, Oana Ursu, and my friends, Liviu and Cristina Iftode. Thank you for your constant reminder of where I come from and where I am going.

CHAPTER 1

INTRODUCTION

The Internet has led to an explosion of product options facing consumers. Since evaluating these options is costly for consumers, in many industries, online search intermediaries have emerged. These intermediaries, such as Amazon or Expedia, return rankings (ordered lists) of third party sellers' products to consumers in response to their queries using algorithms that rank the most relevant products at the top. Intermediaries' rankings help consumers find a matching product more quickly, thereby decreasing search costs and increasing consumer welfare. Because of these benefits, consumers increasingly use search intermediaries in their search and purchase decisions. As consumers shift to using mobile devices to access Internet content, they become more impatient in their search, and hence the importance of intermediary rankings will further increase.

Constructing relevant rankings requires understanding their causal effect on consumer choices. However, this is challenging since rankings are endogenous: consumers pay more attention to higher ranked products both because of their position in the ranking and because they are more relevant for consumers' query. Therefore, identifying the causal effect of rankings requires determining to what extent consumers' choices depend on the position of a product in the ranking rather than its characteristics. This problem has been previously identified in the literature and existing approaches control for the endogeneity of position using a number of methods, including a control function approach (De los Santos and Koulayev, 2014), regression discontinuity design (Narayanan and Kalyanam, 2014), simultaneous equation model (Ghose et al., 2013) or latent instrumental variables (Rutz, Bucklin and Sonnier, 2012). However, without experimental variation, measuring the causal effect of rankings on consumer choices is not conclusive.

In this dissertation, I employ the first data set with experimental variation in the ranking to study the causal effect of rankings on consumer choices. This data set comes from a field experiment ran at Expedia, the world's largest online travel agent, where consumers

looking for hotels were randomized into either (i) seeing Expedia’s ranking, where hotels were ordered by relevance for consumers, or (ii) seeing a Random ranking, where hotels were listed in random order. The data includes 166,036 consumer queries for hotels along with their choices (search and purchases) over an eight-month period. This corresponds to 4.5 million observations on the hotels displayed on Expedia.

Using this data set, I make three contributions. First, using the Random ranking, I show that the position of a product in the ranking has a causal effect on what consumers search, but conditional on search, do not affect purchase decisions. This finding suggests that intermediaries’ rankings can influence consumer purchases only through their search decisions, emphasizing the importance of optimizing consumers’ search. Under Expedia’s ranking, however, both search and conditional purchase decisions are affected by rankings. To my best knowledge, this difference empirically demonstrates for the first time the endogeneity bias in the ranking, namely that the position effect is overestimated in the absence of experimental variation.

The causal effect of rankings cannot be fully understood without identifying the mechanism through which rankings affect consumer choices. The literature has suggested three possible mechanisms: rankings affect consumer search costs (Ghose et al. 2012b; Chen and Yao, 2014), expected utility (Varian, 2007; Athey and Ellison, 2011) or realized utility (De los Santos and Koulayev, 2015). In this dissertation, I use the same data set with the Random ranking to empirically distinguish between these mechanisms. As rankings do not affect purchases conditional on search, I conclude that they cannot affect consumer choices through realized utility. To further disentangle the two remaining mechanisms, I design a test based on a feature of the data set called opaque offers. Opaque offers are ads to purchase an unknown hotel at a discount. They are frequently displayed in certain positions in rankings, which creates variation in the position of a hotel that is unrelated to consumer expected utility. Comparing consumer choices with and without opaque offers, I find that consumers are less likely to search a hotel when it is displayed lower in the ranking due to

an opaque offer than when the offer is not present. This allows me to state that rankings affect consumer choices through search costs, rather than realized utility and independently of expected utility.

Second, to quantify the causal effect of rankings, I use a sequential search model following Weitzman (1979). This model captures consumer behavior in an ordered environment and explains their search and purchase decisions jointly. I apply the model to the Random ranking data to ensure that estimates are not biased by endogeneity. I find an average position effect of \$2.64, which is typically lower than previous results in the literature without experimental variation in the ranking.

Third, I investigate the welfare effects of a utility-based ranking, obtained by reordering hotels based on their estimated expected utility. I find that it leads to an almost twofold increase in consumer welfare, equaling on average \$28.69 (20.72% of the purchase price). A breakdown of this benefit shows that at least 40% of the value is due to lower search costs, confirming the importance of optimizing consumers' search decisions. Finally, I show that the utility-based ranking, not only increases welfare, but also benefits the search intermediary, with transactions increasing by at least 2.55% and revenues by as much as 4.59%.

The rest of the dissertation is organized as follows. In the next two chapters, I review related work and describe the institutional details of the online travel agent industry relevant for this dissertation. In Chapter 4, I present the data, and in Chapter 5, I describe reduced form evidence of the effect of position on consumer choices. In Chapter 6, I introduce the sequential search model, followed by a discussion of the estimation approach and identification. In Chapter 9, I present my results, while in Chapter 10, I investigate the welfare effects of a utility-based ranking through a counterfactual. Chapter 11 provides a discussion of limitations and future research, and the last chapter concludes.

CHAPTER 2

RELATED LITERATURE

In this chapter, I describe relevant previous work relating to the literature on consumer search, in particular to work examining the effect of rankings on consumer choices. Papers such as Chen and Yao (2014), De los Santos and Koulayev (2014), Koulayev (2014), and Ghose et al. (2012a, 2012b, 2013), consider the effect of rankings on consumer online choices in the hotel industry and find a broad range of position effect estimates, as high as \$35 in De los Santos and Koulayev (2014). De los Santos and Koulayev (2014) and Ghose et al. (2013) also address the endogeneity problem of the ranking using a control function approach and a simultaneous equation model, respectively. In contrast, I find lower position effects by using the Random ranking to eliminate endogeneity.

Recently, several papers have used a sequential search model following Weitzman (1979) to quantify consumers' preference and search cost parameters. Kim et al. (2010) is one of the earliest papers to show how to use Weitzman (1979) to model consumer search decisions in an empirical setting, while Honka and Chintagunta (2014) extend it to also model purchase decisions. I use the same model to quantify the causal effect of rankings on both consumer search and purchase decisions. Furthermore, Chen and Yao (2014) and Ghose et al. (2012b) use a maximum likelihood approach that imposes restrictions on the parameters to be estimated to ensure consistency with optimal sequential search. My approach has the benefit of modeling the likelihood of observing consumer search and purchase decisions jointly, thereby capturing a more complete set of parameter restrictions consistent with Weitzman's (1979) search rules. Finally, Kim et al. (2014) use the search rules in Weitzman (1979) in a probit model of sequential search. Their approach is an alternative to the one used in this dissertation that lowers the computational burden of estimation by providing semi-closed form expressions for the probability of choice.

Exploring how to improve rankings, in particular by using a utility-based method, is closely related to work in online recommender systems (see Ansari et al. 2000, Ansari

and Mela 2003), and has been the subject of several recent papers. Ghose et al. (2012a) was one of the earliest papers to use a utility-based ranking, which they show through lab experiments, is superior to several baseline rankings. Furthermore, Ghose et al. (2013) show that a utility-based ranking outperforms other rankings in terms of revenues, while De los Santos and Koulayev (2014) show that it can increase click through rates almost twofold. Finally, Chen and Yao (2014) find that a utility-based ranking increases consumer utility by 2.3%. Compared to the literature, I find greater benefits of the utility-based ranking in terms of consumer welfare, likely due to the elimination of endogeneity bias. Also, I decompose this welfare and find that a large fraction of the improvement is due to lower search costs.

Understanding how rankings affect consumer search is a broader question which is also present in the online sponsored search literature (Ghose and Yang, 2009; Yang and Ghose, 2010; Agarwal et al. 2011; Jerath et al. 2011; Yao and Mela, 2011; Athey and Ellison, 2011; Jeziorski and Segal, 2012; Blake et al. 2014; Baye et al., 2014; Jeziorski and Moorthy, 2014; Chan and Park, 2015; Narayanan and Kalyanam, 2014; Athey and Imbens, 2015), as well as in the theoretical search literature (Hagiu and Jullien, 2011; Berman and Katona, 2013; De Corniere and Taylor, 2014) and the emerging work on mobile search (Ghose et al. 2012c). In this literature, three papers are most relevant for my work. Baye et al. (2014) study search results at Google and Bing to measure the importance of name prominence and position on consumers' clicks. They also find that failing to account for the endogeneity in position inflates the position effect and minimizes the effect of name prominence. Second, Narayanan and Kalyanam (2014) show how to use advertisers' quality scores in a regression discontinuity design framework to address the endogeneity of ad position, which in my setting, I eliminate using the Random ranking. Most recently, Athey and Imbens (2015) focus on estimating heterogeneity in causal effects of rankings on clicks using data from an experiment demoting the best matched search result to the third position. In contrast, I employ a data set where rankings were fully randomized and I also investigate the effect of rankings on purchases.

In sum, this dissertation draws on the vast literature emphasizing the importance of

rankings in consumer choices, including work on the online hotel industry and work on online sponsored search ads. It adds to this literature by measuring the causal effect of rankings, by employing a sequential search model to quantify this effect, and by investigating the welfare benefits of a utility-based ranking using data with experimental variation in the ranking.

CHAPTER 3

THE ONLINE TRAVEL AGENT INDUSTRY

In this chapter, I describe the institutional details of the online travel agent industry that are relevant for this dissertation. In 2013, almost 80% of bookings made online were made through online travel agents (OTA), which had combined revenues of \$157 billion in the U.S. and \$278 billion world wide.¹ In the U.S., four OTAs, Expedia, Booking, Orbitz and Travelocity, account for 95% of bookings, with Expedia being the largest. OTAs revenue is derived under both the agency and the merchant model. Under the agency model, the OTAs receive a commission (e.g. ranging from 10% to 25% for hotels) from third-party sellers for a purchase. Under the merchant model, the OTA purchases the seller's product, which it marks up and sells to consumers.²

To compete for consumers, OTAs rank third-party sellers' products, such as flights, hotels and rental cars. OTAs invest in constructing relevant rankings for consumers on the basis of machine learning techniques. In this chapter, I provide a general overview of one such ranking algorithm, while in Appendix B.2, I describe the technical details behind "learning to rank" algorithms. The position of a product in the ranking is a function of its historical purchases and clicks, its characteristics (e.g. price, quality) and its match with an individual consumer (e.g. consumer's past purchases). The algorithm learns a function of these components (i.e. a score) that best predicts the probability of a purchase or click and those products with a higher score are ranked at the top. In this case, sellers cannot directly affect their position in the ranking. In addition, OTAs reserve specific positions in the ranking for sponsored

1. All figures reported come from three sources: 1. www.economist.com/news/business/21604598-market-booking-travel-online-rapidly-consolidating-sun-sea-and-surfing; 2. www.forbes.com/sites/greatspeculations/2014/04/08/competitive-landscape-of-the-u-s-online-travel-market-is-transforming/; and 3. www.wsj.com/articles/amazons-new-travel-service-enters-lucrative-online-travel-market-1429623993.

2. In the case of Expedia, in the first quarter of 2013 (the relevant period for this analysis), 70% of its global revenues came from the sale of hotel rooms and most of its bookings were done under the agency model (54%). See Expedia's Earnings Release for the first quarter of 2013 for details: ir.expediainc.com/common/download/download.cfm?companyid=EXPEfileid=657057filekey=F8EB6718-117C-444C-A1B8-7C5FD4826A54filename=EXPE-Q1_2013_Earnings_Release.pdf.

ads.³ More precisely, a seller enters a pay-per-click auction for the sponsored ad slots by bidding to target consumer queries for particular destinations and travel dates.⁴ The auction is adjudicated by evaluating the bids submitted by all sellers and their scores, and the winner is displayed in the same position on all result pages.

The institutional details of the OTA market motivate and affect my analysis. In particular, the algorithm used makes the position of a hotel endogenous. As a result, using observational data on consumer queries for hotels given the default ranking will not suffice in determining the causal effect of the ranking on consumer choices. Thus, I rely on estimates using the Random ranking that eliminate the endogeneity bias. Without properly separating the effect of position on choices from other hotel characteristics, the ability to improve the current ranking is limited.

3. In the case of Expedia, ads may appear at the top of the ranking and in the last two positions. Details on the auction it uses for sponsored ads can be found at searchsolutions.expedia.com/how-it-works /searchsolutions.expedia.com/how-it-works/.

4. Search impressions with sponsored ads are included in the data, but are not flagged as such. To identify them, I use the fact that ads are more likely to appear in popular destinations with a large number of hotels.

CHAPTER 4

DATA

This chapter describes the data set employed in this dissertation. The data set comes from consumers' queries for hotels on Expedia. It includes results from 166,036 queries along with consumers' clicks (interpreted as their search decisions) and purchases over an eight-month period between November 1, 2012 and June 30, 2013.¹ There are 4.5 million observations on 54,877 hotels located in 55 countries and 788 different destinations. The Expedia data is provided at the level of a search impression. A search impression is an ordered list of hotels and their characteristics seen by consumers in response to a query describing their trip.

The main feature of this data set is that in only two thirds of the search impressions consumers saw Expedia's ranking, where hotels were order by relevance, while in one third of search impressions hotels were ordered randomly (Random ranking). In Appendix B.1, I show that consumers were randomly assigned to these two rankings. I also show that the position of the hotel was randomly generated under the Random ranking. To my best knowledge, I am the first to use a data set with a Random ranking to investigate the causal effect of rankings on consumer search and purchase decisions.

4.1 Search Process

To understand the characteristics of this data set, in this section I explain the three steps of consumers' search process for hotels on Expedia. This process is illustrated in Figure 4.1. First, the consumer begins her search query on Expedia by specifying details of her trip, such as the destination (city, country), the travel dates and the number of travelers and rooms requested. In addition, based on her query, the number of days before the beginning of the trip is recorded (booking window). Second, in response to her query, the consumer gets a search impression of all the hotels that match her request, distributed over multiple pages.

1. Appendix A contains details about data cleaning.

From this search impression, I observe the first page of results displayed to consumers, which includes the hotel ID, its position in the ranking and its characteristics (price, number of stars and reviews, location, a chain and a promotion indicator).² On Expedia, consumers can sort or filter results, however, the data set only contains those search impressions where consumers made choices from the ranking displayed.³ After observing the list of hotels, the consumer can click on a particular one to observe more information (third step). In this case, she navigates to a sub-page reserved for that hotel, where she can see additional pictures, previous customers’ reviews, etc. Then, she can either return to the previous screen to click on another hotel, leave the site without purchasing, or she can purchase. I observe both clicks and purchases that consumers make. However, I do not observe the additional information that consumers see on the hotel’s page.

Figure 4.1: Consumer Search Process Illustration



2. In a companion data set from Wharton Customer Analytics Initiative (WCAI) on consumer queries for hotels in Manhattan on a similar online travel agent, I find that in 67% of search impressions, consumers only consider the first page of results. See Appendix B.3 for a description of this data set.

3. The WCAI companion data set shows that few queries actually contain sorted/filtered results: only 32% of all search impressions and 34% consumers sort/filter. For more details, see Appendix B.3.

4.2 Data Description

In this section, I provide a detailed description of the data. Table 4.1 provides summary statistics at the hotel and the search impression level. Hotels charge on average \$160 per night, have more than 3 stars and a median review score of 4. Most hotels belong to a chain and one quarter of them display a promotion. Location attractiveness is summarized by a score ranging from 0 to 7 designed by Expedia to measure how central a hotel is located, what amenities surround it, etc, and the average hotel has a location score of 3.17. In a search impression, consumers see on average 27 hotels displayed. On average, consumers look for trips lasting two days that begin one month from their query. The median search impression is for a trip for one hotel room and two adults traveling with no children. There are a total of 186,171 clicks, with 58,501 clicks under the Random ranking. There is at least one click per search impression, with 7% of search impressions including two or more clicks. Two thirds of all search impressions end in a transaction for a total of 108,903 transactions overall and 3,930 under the Random ranking. Approximately 7,500 search impressions have historical information on the average star rating and price of hotels previously purchased by a consumer. However, this is not enough information to link consumers who are making repeated queries over time, thus I interpret each search impression as being made by a different consumer.⁴

The data is anonymized, so identifying the destination to which a consumer wishes to travel is not possible. However, there exists suggestive evidence that the largest country is the U.S.⁵ In the data, 80% of queries to this country are made by consumers also located there, suggesting that it has a large territory with a large fraction of domestic travel. This evidence is also consistent with the fact that 73% of Expedia’s traffic comes from U.S. visitors.⁶

The data set I use resulted from an effort by Expedia to elicit the involvement of data

4. In the WCAI data set, I find that more than 40% of consumers make only one query.

5. The largest country in the data is labeled 219.

6. Information retrieved in May 2015 from Alexa.com: www.alexa.com/siteinfo/expedia.com.

Table 4.1: Hotel and Search Impression Summary Statistics

	Observations	Mean	Median	SD	Min	Max
<i>Hotel level</i>						
Price	4,503,043	159.71	132.00	102.43	10	1000
Stars	4,418,366	3.32	3.00	0.87	1	5
Review Score	4,498,652	3.89	4.00	0.86	0	5
Chain	4,503,043	0.66	1.00	0.47	0	1
Location Score	4,503,043	3.17	3.14	1.51	0	7
Promotion	4,503,043	0.25	0.00	0.44	0	1
<i>Search impression level</i>						
Number of Hotels Displayed	166,036	27.12	31.00	8.10	5	38
Trip Length (days)	166,036	2.42	2.00	1.98	1	40
Booking Window (days)	166,036	39.26	18.00	53.89	0	498
Saturday Night (percent)	166,036	0.50	1.00	0.50	0	1
Adults	166,036	2.00	2.00	0.90	1	9
Children	166,036	0.39	0.00	0.79	0	9
Rooms	166,036	1.12	1.00	0.44	1	8
Total Clicks	166,036	1.12	1.00	0.61	1	25
Two or More Clicks (percent)	166,036	0.07	0.00	0.25	0	1
Total Transactions	166,036	0.66	1.00	0.48	0	1
Random Ranking (percent)	166,036	0.31	0.00	0.46	0	1

miners in the improvement of its ranking algorithm. Through the International Conference on Data Mining (ICDM 2013) and Kaggle.com (an online platform facilitating participation of data miners in competitions posted by companies), Expedia released two data sets on consumers' hotel queries: one from its default ranking and another where the position of hotels was randomly generated (Random ranking data).⁷

4.3 Pros and Cons of the Data Set

This section considers both the pros and the cons of using this Expedia data set. The main advantage of using it comes from the Random ranking, which provides the unique opportunity to study the causal effect of rankings on consumer choices. However, there are two features of the data that Expedia released that limit the analysis in this study: (i) search impressions included in the data contain at least one click, and (ii) search impressions

⁷. This data is available at www.kaggle.com/c/expedia-personalized-sort/data.

leading to a transaction were oversampled.^{8,9} These features have three implications for the analysis. First, the fact that only search impressions with at least one click are observed means that consumers who did not click are missing from the sample, and hence the results may not generalize to all consumers who saw the Random ranking. Second, the fact that only search impressions with at least one click are observed for the two rankings means that they may correspond to distinct fractions of the consumer populations who saw the respective rankings, and hence a quantitative comparison of consumers' choices under the two rankings may be problematic. Even though I show that consumers were randomly assigned to seeing either ranking and exhibit similar observable characteristics in the sample (see Appendix B.1), a quantitative comparison of the two rankings is still difficult since consumers may differ on unobservables (for example, consumers who click under the Random ranking may have lower search costs than those who click under Expedia's ranking). Third, the fact that converting search impressions were oversampled means that the total number of purchases cannot be compared across the two rankings. In what follows, I show how I was able to overcome these limitations to study the causal effect of rankings.

For the reduced form chapter of the dissertation, I use the Random ranking to compare the fraction of clicks and purchases across positions. The sampling done by Expedia may affect the absolute number of clicks and purchases by position, but the relative effect of position on choices is unchanged. Since this relative effect recovers the causal effect of rankings, the sampling done does not affect the results. If search impressions contain at least one click and converting search impressions were oversampled, another concern is that consumers have identified a hotel to purchase on previous visits to the site, which would minimize the observed effect of rankings on their choices. However, this story is contradicted by the fact that consumers click more often on higher ranked hotels under the Random ranking, which are

8. Also, a higher fraction of converting search impressions appear under Expedia's ranking than the Random ranking, possibly to hide its true conversion rate.

9. Other data sets in the literature have similar properties. For example, Chen and Yao (2014), restrict their attention to search impressions that end in a transaction, thereby also reducing their data set to one that has at least one click and one transaction per search impression.

unlikely to display their previously identified hotel (see result in Section 5.1). For this part of the dissertation, I use the Expedia ranking only to point out qualitatively the direction of the bias that would arise in the absence of experimental variation, without quantifying differences between the two rankings.

For the sequential search model, I use the Random ranking to quantify the effect of rankings on consumer choices. In this case, removing search impressions without clicks may translate into an underestimation of search costs. However, the oversampling of transactions does not affect preference and search cost parameters estimates: preference parameters can be consistently estimated since I observe both consumers who purchase and those who do not, while search costs are not identified from purchase decisions and are thus not affected by the sampling done (see Chapter 8 for more details).

For the counterfactual, I compare the utility-based ranking built on the model's estimates with simulated choices under Expedia's ranking (instead of actual choices). Simulated choices are computed using the preference and search cost parameters estimated using the Random ranking and Weitzman's (1979) search rules. By doing so, I avoid the direct comparison of consumer choices under the two rankings, which may be biased since their underlying parameters may be different due to the sampling done.

Despite its limitations, I conclude that the causal effect of rankings on choices can be studied with this data set that contains experimental variation in the ranking. In the following chapters, I use this data to recover the causal effect of rankings, as well as quantify this effect and show how these results can be used to improve the ranking algorithm.

CHAPTER 5

REDUCED FORM EVIDENCE

In this chapter, I present reduced form results of the causal effect of rankings on consumer search and purchase decisions, as well as determine the mechanism through which rankings affect these choices.

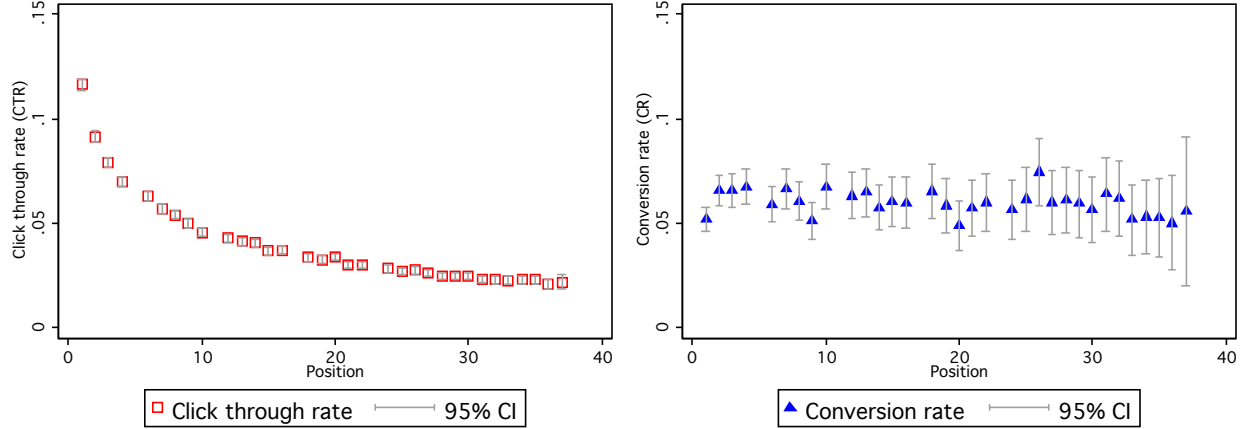
5.1 Causal Effect of Rankings on Search and Purchase Decisions

To explore the causal effect of rankings, in this section I look at the effect of position on consumer search and purchase decisions under the Random ranking. In the left panel of Figure 5.1, I illustrate the click through rate of a position, i.e. the fraction of times a position was clicked out of all the times it was displayed.¹ I find that the higher hotels are ranked the more clicks they receive, which happens despite the fact that under the Random ranking the quality of hotels is unrelated to their position in ranking. This proves that rankings have a causal effect on consumer search.² In the right panel of Figure 5.1, I plot the conversion rate of a position, i.e. the percent of clicks that end in a purchase. I find that the conversion rate of the Random ranking is approximately constant across positions. This means that rankings have no causal effect on consumers' purchases conditional on search. Table 5.1 shows that the same pattern as in Figure 5.1 holds in a regression controlling for hotel and destination characteristics. In particular, hotels ranked lower (higher position) have a lower probability of being clicked (columns 1-3) and purchased (column 4), whereas conditional on a click, the position of the hotel has no effect on purchases (columns 5-7). This result holds both when considering the linear effect of position (columns 1, 2 and 4-6) and the non-linear fixed effect of being among a group of five hotels (columns 3 and 7). Note

1. I restrict attention to search impressions that do not include a hotel in positions 5, 11, 17, 23, which are reserved for opaque offers by Expedia. Under the Random ranking this represents 93% of search impressions.

2. See Appendix B.4 for robustness checks, controlling for search impressions that more likely contain sponsored ads, that display few hotels, or that lead to a transaction. The same pattern as in Figure 5.1 holds.

Figure 5.1: Position Effect: Random Ranking



that unconditional on a click, higher ranked hotels lead to more purchases. However, the fact that conditional on a click, consumers who clicked at the top and those who clicked lower in the ranking are equally likely to purchase, means that intermediaries’ rankings can influence consumer purchases only through their search decisions.

To put these results into context, in Figure 5.2, I plot the click through rate and the conversion rate of Expedia’s ranking. As described in Section 4.3, I cannot quantitatively compare the two rankings.³ However, I can point out qualitatively what the pattern in click through rate and conversion rate is in the absence of experimental variation in the ranking. In contrast to the results for the Random ranking, Figure 5.2 shows that under Expedia’s ranking higher ranked hotels lead to both more clicks and more purchases conditional on a click, hence the effect of position is increased (Table B.4 in Appendix B.4 shows regression results mirroring the pattern in Figure 5.2). This difference empirically demonstrates, for the first time to my best knowledge, the endogeneity bias in the ranking, namely that the position effect is overestimated in the absence of experimental variation in the ranking.

3. Note that the conversion rate under Expedia’s ranking is higher than that of the Random ranking.

Table 5.1: Estimates of Clicks and Transactions (OLS): Random Ranking


	Click		Transaction	Transaction conditional on click			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
<i>Position effect</i>							
Position	-0.00192*** (0.00002)	-0.00186*** (0.00002)		-0.00012*** (0.00000)	-0.00009 (0.00010)	-0.00001 (0.00010)	
Positions 1-5			0.06168*** (0.00068)				0.00009 (0.00282)
Positions 6-10			0.02697*** (0.00053)				-0.00120 (0.00302)
Positions 11-15			0.01381*** (0.00054)				-0.00141 (0.00356)
Positions 16-20			0.00788*** (0.00052)				-0.00202 (0.00378)
<i>Hotel characteristics</i>							
Price		-0.00013*** (0.00000)	-0.00012*** (0.00000)	-0.00001*** (0.00000)		-0.00010*** (0.00001)	-0.00010*** (0.00001)
Stars		0.01608*** (0.00030)	0.01610*** (0.00030)	0.00110*** (0.00008)		0.00094 (0.00169)	0.00095 (0.00169)
Review Score		0.00124*** (0.00020)	0.00126*** (0.00020)	0.00021*** (0.00004)		0.00615*** (0.00114)	0.00616*** (0.00114)
Chain		0.00221*** (0.00046)	0.00231*** (0.00046)	0.00026* (0.00011)		0.00466 (0.00242)	0.00466 (0.00242)
Location Score		0.00439*** (0.00017)	0.00443*** (0.00017)	0.00053*** (0.00004)		0.00717*** (0.00085)	0.00717*** (0.00085)
Promotion		0.01162*** (0.00050)	0.01152*** (0.00050)	0.00125*** (0.00013)		0.01036*** (0.00243)	0.01036*** (0.00243)
Query characteristics	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Destination FE	No	Yes	Yes	Yes	No	Yes	Yes
Adjusted R^2	0.010	0.015	0.017	0.003	-0.000	0.033	0.033
Observations	1,245,455	1,220,917	1,220,917	1,220,917	54,614	53,921	53,921

Standard errors in parentheses

Notes: In columns (3) and (7), the omitted category is [Positions 21-end]. Standard errors clustered at the search impression level. Restrict sample to search impressions with opaque offers. Query characteristics included in all regressions: trip length, booking window, number of adults and children traveling, number of rooms, and Saturday night dummy.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Figure 5.2: Position Effect: Expedia Ranking



Note that the scales of the y-axis of the two figures are different.

5.2 Mechanism Driving the Effect of Rankings on Search

The last section established the causal effect of rankings: positions affect clicks, but not subsequent purchases. However, this effect cannot be fully understood without identifying the mechanism through which rankings affect consumer search. Why do consumers click more at the top of the ranking? The literature identifies three possible mechanisms that could explain this effect. Either consumers incur a cost for exploring lower ranked hotels, so rankings affect consumer search costs (Ghose et al. 2012b; Chen and Yao, 2014). Or, consumers click more at the top, because they expect those hotels to be of higher quality, so rankings affect consumers' expected utility (Varian, 2007; Athey and Ellison, 2011). Or, consumers click more at the top, because they derive utility from higher ranked hotels (De los Santos and Koulayev, 2015).⁴ Unfortunately, there has been little empirical evidence on the exact mechanism through which rankings affect choices. This is without a doubt due to the difficulty of controlling one channel and exploring how rankings affect another.

I use the same data set with the Random ranking to empirically distinguish between these mechanisms. First, since rankings do not affect purchases conditional on a click, they cannot affect consumers' realized utility. Thus, the evidence from Section 5.1 eliminates the third mechanism. Second, to further disentangle the two remaining mechanisms, I use the fact that, in most search impressions, Expedia reserves positions 5, 11, 17 and 23 in the ranking for opaque offers (see Appendix B.6 for an illustration of opaque offers). Opaque offers are ads to purchase an unknown hotel at a discount. They occupy a position in the ranking, but are not one of the ranked hotels. Thus, when an opaque offer is displayed, consumers' expectations of the ranked hotels should not change, since their order in the search impression has not changed (the n th best hotel according to Expedia's algorithm is still the n th displayed hotel, although it may not appear in position n). However, the opaque offer demotes a hotel to a lower position, increasing its search cost. The variation

4. Murthi and Rao (2012) investigate a related question, looking at the extent to which consumer choices of offline grocery products are affected by their price expectations or the actual posted prices.

in the position of hotels due to opaque offers is unrelated to individual consumer queries (exogenous), since offers are displayed based on availability.

To perform the test, I run a regression of the probability of clicking on the 5th displayed hotel in the ranking as a function of its position, hotel and query characteristics for the Random ranking.⁵ In Table 5.2, I show my results. In column 1, I find that when the 5th displayed hotel is shown in position 6 rather than position 5 (because it is demoted by an opaque offer), it receives 1.8% fewer clicks. I find a similar result after controlling for hotel and query characteristics (second column), as well as after restricting attention to search impressions with clicks lower in the ranking (third column), and to the largest destination in the data in order to include hotel fixed effects (fourth column). These results allow me to state that rankings affect consumer choices through search costs, rather than realized utility and independently of expected utility.

One concern with this result is that consumers who clicked less on the 5th displayed hotel when it was shown in position 6 actually clicked on the opaque offer. Even though I cannot completely alleviate this concern, as I do not observe clicks on the opaque offer, it is improbable that such behavior is predominant in the data since it involves navigating to a different screen. Thus, if a click on the opaque offer occurs, it is unlikely that the consumer will return to search lower ranked hotels. To support this claim, in the third column of Table 5.2, I show that even in the sample of search impressions where consumers click lower in the ranking than the position of the opaque offer (these are search impressions that are least likely to contain a click on the opaque offer), a similar result holds as in the full sample. Therefore, I expect clicks on opaque offers to be more common in search impressions without any clicks on the displayed hotels and thus not detract from the main findings of Table 5.2.

In this chapter, I provided evidence of the causal effect of rankings on consumer search and purchase decisions, as well as identified the mechanism through which rankings affect

5. The scarcity of data prevents me from providing conclusive evidence from similar tests in the other positions reserved for opaque offers, i.e. 11, 17 and 23.

Table 5.2: Estimates of Click on the Fifth Displayed Hotel (OLS)

	Baseline model		Searches with click after position 4	Largest destination
	(1)	(2)	(3)	(4)
<i>Position effect</i>				
Position (5 or 6)	-0.01801 ⁺ (0.00990)	-0.02061* (0.00999)	-0.03290* (0.01437)	-0.02812 (0.07727)
<i>Hotel characteristics</i>				
Price		-0.00016*** (0.00001)	-0.00022*** (0.00002)	-0.00031* (0.00013)
Stars		0.02177*** (0.00175)	0.03070*** (0.00249)	
Review Score		-0.00007 (0.00134)	-0.00077 (0.00190)	
Chain		0.00781** (0.00262)	0.01188** (0.00369)	
Location Score		0.00634*** (0.00102)	0.00878*** (0.00145)	
Promotion		0.01086*** (0.00268)	0.01684*** (0.00378)	
Query characteristics	Yes	Yes	Yes	Yes
Destination FE	No	Yes	Yes	Yes
Hotel FE	No	No	No	Yes
Adjusted R^2	0.000	0.008	0.013	0.018
Observations	50,947	50,238	35,103	1,353

Standard errors in parentheses

Notes: Search impressions with opaque offers will display the 5th hotel in position 6, while those without will display it in position 5. The largest destination in the data is labeled 8192. Query characteristics included in all regressions: trip length, booking window, number of adults and children traveling, number of rooms, and Saturday night dummy.

+ $p < 0.10$ * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

search. In the next chapter, I present a sequential search model that can describe consumer choices from an ordered list.

CHAPTER 6

MODEL

In this chapter, I develop a sequential search model following Weitzman (1979), which will be used in the following chapters to quantify the effect of rankings on consumer search and purchase decisions and to measure the welfare effects of a utility-based ranking.

6.1 Utility

Suppose consumer i 's utility from purchasing hotel $j \in \{1, \dots, J\}$ is given by

$$\begin{aligned} u_{ij} &= v_{ij} + \epsilon_{ij} \\ &= x_j \beta - \alpha p_{ij} + \epsilon_{ij} \end{aligned} \tag{6.1}$$

where v_{ij} contains hotel characteristics observable to the consumer in a search impression and ϵ_{ij} captures those characteristics that she searches for. More precisely, v_{ij} (the expected utility of the hotel) has two components: x_j , which contains the hotel's number of stars, review score, location, a chain and a promotion indicator, and p_{ij} , which gives the price of the hotel at the time of i 's query. The consumer observes ϵ_{ij} by clicking (searching) on j , but this information is not observed in the data. I thus model it as following a standard normal distribution, consistent with previous literature (Kim et al., 2010; Chen and Yao, 2014). Since the data is provided at the level of a search impression (instead of the consumer level), consumer heterogeneity is only captured by ϵ_{ij} . The consumer has an outside option that is not observed in the data, which I model as an idiosyncratic shock to utility, $u_{i0} = \epsilon_{i0}$.

6.2 Search Cost

Consumer i observes characteristics v_{ij} for each hotel j displayed in a search impression for free. To discover ϵ_{ij} , the consumer has to pay a search cost for clicking, which I model as

$$c_{ij} = \exp(k + \delta \rho_{ij}) \tag{6.2}$$

where k gives the baseline level of search costs and ρ_{ij} gives the position of hotel j in the ranking at the time of i 's search query. The exponential function assumption of search costs is consistent with prior literature (Kim et al. 2010, Ghose et al. 2012b, Chen and Yao 2014) and it ensures that search costs are positive. I assume the position of the hotel affects search costs, instead of expected or realized utility, consistent with findings in Section 5.2.

6.3 Optimal Search

To study consumers' optimal search strategy, consider the following scenario. Suppose the consumer has searched a number of hotels (as well as the outside option, which is always searched). She must then decide whether to continue searching, and if so, which hotel to search. If she decides to stop searching, she must determine whether to purchase, and if so, which of the searched hotels to choose. To model this behavior, I rely on Weitzman (1979), which characterizes consumers' optimal search strategy using the following search rules.

1. **Selection Rule:** If a search is to be made, the hotel with the highest reservation utility should be searched next.

The reservation utility z_{ij} of a hotel j is defined as the level of utility that would make consumer i indifferent between searching it or not. It is computed by equating the marginal gains from searching j with the marginal cost as in

$$c_{ij} = \int_{z_{ij}}^{\infty} (u_{ij} - z_{ij})f(u_{ij})du_{ij} \quad (6.3)$$

2. **Stopping Rule:** Search should terminate when the maximum utility observed exceeds the reservation utility of any unsearched hotel.
3. **Choice Rule:** Once the consumer stops searching, she will choose the hotel with the highest utility among those searched.

Kim et al. (2010) show that equation (6.3) can be rewritten by taking advantage of the distributional assumptions made. More precisely, using ϵ_{ij} 's normality assumption and the expression for the expectation of the truncation of normally distributed random variables, equation (6.3) becomes

$$\begin{aligned} c_{ij} &= (1 - \Phi(m_{ij}))(\lambda(m_{ij}) - m_{ij}) \\ &= B(m_{ij}) \end{aligned} \tag{6.4}$$

where $\lambda(\cdot) = \frac{\phi(\cdot)}{1 - \Phi(\cdot)}$ is the hazard function and where $m_{ij} = z_{ij} - v_j$. The result in equation (6.4) provides a straightforward way of computing the reservation utility z_{ij} . More precisely, it says that given search cost c_{ij} , one can invert $B(\cdot)$ and solve for m_{ij} .¹ Then, using the definition of m_{ij} , the reservation utility of hotel j equals $z_{ij} = v_j + m_{ij}$, i.e. it equals the expected utility of the hotel plus a function of the search costs for clicking on the hotel.

Why is this model based on Weitzman's (1979) search rules preferable to other possible approaches? The model assumes that consumers observe all hotels in a search impression at no cost, and decide which ones to click and whether to purchase based on their reservation and realized utilities. In reality, there is a cost to observe a hotel, which is associated with scrolling through the ranking until it is reached. Thus, an alternative model could be one that includes consumers' pre-click decision to observe a hotel. Such a model would require solving two dynamic programming problems (computing two reservation utilities): one to decide how many hotels to observe and another to decide when to stop searching and what to purchase (Chan and Park, 2015). However, a way to simplify this model is to realize that consumers can only click on hotels they have observed. Thus, when they click, they must pay both a clicking and an observation cost. In my model, I capture both of these costs by defining the search cost as having two components: the baseline level of search costs can be

1. Kim et al. (2010) show that function $B(\cdot)$ is monotonic and decreasing in its argument and that a unique solution to $c_{ij} = B(m_{ij})$ exists. Thus, this inversion is possible. To speed up computation, I follow Kim et al. (2010) and construct a look-up table for $c_{ij} = B(m_{ij})$ outside the estimation loop. During estimation, for a particular value of search costs, I use the table to look up the value of m_{ij} and construct the reservation utility.

interpreted as the clicking cost and the position effect as the observation cost. Thus, in the absence of data on consumers' scrolling behavior, I believe that the present model is a good approximation to one that accounts for both decisions.

Another variation of this model could be one where the order in which consumers search is exogenously given by the ranking instead of Weitzman's (1979) search rules. However, this approach is rejected by the data. This alternative model suggests that a consumer's n th click would occur in position n . Instead, in the data the difference between clicks in search impressions with at least two clicks is on average 12 positions (median is 10). This suggests that consumers' click order is not only determined by the ranking, but also by observed characteristics, which the present model accounts for by modeling click order as being determined by reservation utilities following Weitzman (1979).

Having introduced the model, the following chapter discusses its estimation.

CHAPTER 7

ESTIMATION

In this chapter, I show how the parameters of the model are estimated. Suppose consumer i searched h of the total J hotels displayed in a search impression and that she chose j (including the outside option of not purchasing). Denote by $R_i(n)$ the identity of the hotel with the n th largest reservation utility for consumer i . Order these hotels by their reservation utilities and let $R_i = [R_i(1), \dots, R_i(h)]$ denote the set of searched hotels and the order in which they were searched by consumer i . The outside option is always searched (denoted for simplicity as either $j = 0$ or $R_i(0)$).

In this setting, Weitzman's (1979) optimal search strategy translates into the following restrictions on the preference and search cost parameters of the consumer. The selection rule requires that, if the consumer makes an n th search, her reservation utility from that hotel exceed her reservation utility from all hotels searched next and all those not searched. Formally, it must be that

$$z_{iR_i(n)} \geq \max_{k=n+1}^J z_{iR_i(k)}, \quad \forall n \in \{1, \dots, J-1\} \quad (7.1)$$

otherwise the consumer would have searched a hotel with a higher reservation utility instead.

The stopping rule imposes two restrictions on the preference and search cost parameters. First, if the consumer makes an n th search, then her reservation utility from that hotel must exceed her utility from all hotels searched so far, including the outside option. Otherwise the consumer would have stopped searching. Formally,

$$z_{iR_i(n)} \geq \max_{k=0}^{n-1} u_{iR_i(k)}, \quad \forall n \in \{1, \dots, h\} \quad (7.2)$$

Second, all unsearched hotels must have a lower reservation utility than the maximum utility of the searched alternatives, including the outside option,

$$z_{iR_i(m)} \leq \max_{k=0}^h u_{iR_i(k)}, \quad \forall m \in \{h+1, \dots, J\} \quad (7.3)$$

otherwise the consumer should have continued searching.

Finally, consistent with the choice rule, if the consumer chooses j (including the outside option), then her utility from this choice must exceed the utilities of all the hotels searched and the outside option. Formally,

$$u_{ij} \geq \max_{k=0}^h u_{iR_i(k)}, \quad \forall j \in R_i \cup \{0\} \quad (7.4)$$

If consumers search sequentially, then they make search and purchase decisions jointly. This means that the probability of observing a certain outcome in the data is characterized by the joint probability of equations (7.1-7.4) holding. However, the data set does not contain information on the order in which consumers click. Thus, even though consumers decide in which order to click, as captured by the selection rule in equation (7.1), only observing clicks and not click order does not allow this decision to inform the parameter estimates.¹ However, in this data, most consumers only click once, while for the rest of the 7% of search impressions with at least two clicks, I assume that consumers click first on hotels displayed closer to the top of the ranking.² This assumption is supported by data from WCAI where consumers' search order is observed and where the position of the hotel explains most of the click order (see Appendix B.7), allowing me to use equations (7.2-7.4) in estimation. Alternatively, the estimation method proposed by Honka and Chintagunta (2014) could be used, which accounts for consumers' click order in the absence of data showing the order in which they searched.

The probability P_{ijR_i} that i searches in the order R_i and chooses j (including the outside option) is given by³

1. If data on consumers' click order were available, one way to incorporate the selection rule in estimation would be to model consumers' search costs as including an unobserved component.

2. This assumption may lead to an overestimation of the position effect if the expected utility of the lower ranked hotel was greater than that of the higher ranked hotel. However, this is unlikely since hotels are ranked randomly, leading to significant variation in the characteristics displayed.

3. In equation (7.5), I suppress the domain of possible values. It appears in equations (7.2-7.4).

$$\begin{aligned}
P_{ijR_i} &= Pr(z_{iR_i(n)} \geq \max_{k=0}^{n-1} u_{iR_i(k)} \cap z_{iR_i(m)} \leq \max_{k=0}^h u_{iR_i(k)} \cap u_{ij} \geq \max_{k=0}^h u_{iR_i(k)}) \\
&= \int I(cond) \phi(\epsilon_i) d\epsilon_i
\end{aligned} \tag{7.5}$$

where $I(cond)$ is an indicator for whether conditions (7.2-7.4) hold. The log-likelihood is

$$LL = \sum_i \sum_{R_i} \sum_j d_{ijR_i} \log P_{ijR_i} \tag{7.6}$$

where $d_{ijR_i} = 1$ if i chose search order R_i and chose j (including outside option). Since search and purchase decisions are made jointly, the integral in equation (7.5) does not have a closed form solution. Thus, I replace P_{ijR_i} with the simulated choice probability \hat{P}_{ijR_i} . This results in the following simulated log-likelihood

$$SLL = \sum_i \sum_{R_i} \sum_j d_{ijR_i} \log \hat{P}_{ijR_i} \tag{7.7}$$

The choice probability \hat{P}_{ijR_i} can be simulated in a number of ways. A straightforward simulator is accept-reject (AR), which was originally proposed by Manski and Lerman (1981) for probit models. This simulator approximates P_{ijR_i} by the proportion of draws from the appropriate distribution that satisfy the conditions in (7.5). However, using the AR simulator in maximizing the SLL can be problematic for two reasons. First, any finite number of draws can result in a reject, so that \hat{P}_{ijR_i} is zero, which is especially likely if the data contains very few choices, as is the case in this dissertation. The second difficulty comes from the fact that the choice probabilities are not twice differentiable, so the simulated probabilities will not be smooth. Thus, finding a maximum by optimizing the SLL using first and second derivatives will not be effective. Even though there is a way to circumvent this problem by using an approximation of the gradient to the SLL instead, in practice AR is difficult to use (Train, 2009). An alternative simulator can be obtained by replacing the indicator function in the AR simulator with a smooth and increasing function that has defined first and second derivatives. As suggested by McFadden (1989), I choose the logit function that satisfies these conditions and is convenient to use. This method is known as the logit-smoothed AR

simulator. It has been used to estimate probit models, as well as to estimate search models (see Honka, 2014 and Honka and Chintagunta, 2014).

Simulating \hat{P}_{ijR_i} using the logit-smoothed AR simulator involves the following steps:

1. Draw $d = \{1, \dots, D\}$ samples of ϵ_{ij}^d for each consumer and each hotel.
2. Use ϵ_{ij}^d to form utility u_{ij}^d .
3. Use the relation $c_{ij}^d = B(m_{ij}^d)$ to compute m_{ij}^d and form reservation utilities z_{ij}^d .
4. Define the following expressions for each draw d
 - (a) $\nu_1^d = z_{iR_i(n)}^d - \max_{k=0}^{n-1} u_{iR_i(k)}^d$
 - (b) $\nu_2^d = \max_{k=0}^h u_{iR_i(k)}^d - z_{iR_i(m)}^d$
 - (c) $\nu_3^d = u_{ij}^d - \max_{k=0}^h u_{iR_i(k)}^d$
5. Compute S^d for each draw d using the expressions above

$$S^d = \frac{1}{1 + \sum_{n=1}^3 e^{-\frac{\nu_n^d}{\lambda}}} \quad (7.8)$$

where $\lambda > 0$ is a scaling parameter.

6. The average of S^d over D draws of the error terms gives the simulated choice probability

$$\hat{P}_{ijR_i} = \frac{1}{D} \sum_d S^d \quad (7.9)$$

There is little guidance in choosing the scaling parameter λ . As $\lambda \rightarrow 0$, the simulator approaches the AR simulator and is thus unbiased. So, the researcher should use a small enough λ , but not too small to reintroduce the numerical problems one faces when optimizing with a non-smooth function. For this data set, I determine the appropriate scaling parameter using Monte Carlo simulations, which are described in Section 8.2.

This chapter described the estimation of the model, while in the next chapter I discuss its identification.

CHAPTER 8

IDENTIFICATION

In this chapter, I show how the model parameters are identified, as well as provide Monte Carlo simulation results that validate the estimation approach.

8.1 Model Parameters

This section demonstrates the identification of preference and search cost parameters. Weitzman's (1970) stopping and choice rules in equations (7.2)-(7.4) help identify the preference parameters of characteristics that vary by hotel. Characteristics in the utility function that do not vary by hotel cannot be identified, as they shift both reservation and realized utilities by the same amount. Therefore, I do not include them in the specification of the utility function. Stopping rules in equations (7.2)-(7.3), impose an upper and a lower bound on search costs, respectively, that must have made it optimal for the consumer to conduct a certain number of searches. Importantly, the choice rule does not affect the search cost estimate. The stopping rules, however, only recover a range of search costs. The level of search costs (parameter k) is pinned down by the functional form and the distribution of the utility function through the optimal search relation in equation (6.4). The position effect is identified from variation in where in the ranking consumers click. Other factors that may affect consumer search costs that do not vary across hotels but only vary across consumers, such as the booking window or the trip length, cannot be recovered without data showing multiple search impressions per consumer. Note that I had previously normalized the systematic component of the outside option utility to zero. One reason for this is that this utility is not separately identified from the baseline search cost parameter, k . Further, most consumers only click once in the data set and I cannot link queries made by the same consumer over time. For these reasons, the implied level of search costs is large and difficult to interpret.¹

1. See Appendix B.10 for more details.

An obstacle to identifying model parameters is that the position of a hotel (in a curated ranking) as well as its price may be endogenous. In what follows, I describe how the Random ranking eliminates the endogeneity bias in position and show evidence to alleviate price endogeneity.

Position Endogeneity: The position of a hotel in a curated ranking is endogenous since consumer choices depend on the hotel’s current position and the hotel’s position depends partially on its past performance (clicks and purchases) and thus, on consumer preferences. When unaccounted for, position endogeneity leads to biased preference and search cost estimates. To illustrate this problem, decompose the probability that consumer i chooses hotel j (click or purchase) into the conditional choice probability given j ’s position and the marginal probability of observing j in a given position.² More formally,

$$Pr(\text{choice}_{ij}|\theta_{ij}) = Pr(\text{choice}_{ij}|x_{ij}, \text{position}_j, \theta_{ij})Pr(\text{position}_j|x_{ij}, \theta_{ij}) \quad (8.1)$$

where θ_{ij} denotes consumer i ’s preference parameter for hotel j and x_{ij} denotes the hotel characteristics revealed at the time of i ’s query. Because the position of the hotel in the ranking is determined by an algorithm reflecting the probability that the consumer chooses j , the marginal probability $Pr(\text{position}_j|x_{ij}, \theta_{ij})$ depends not only on hotel characteristics, but also on consumer preferences. Thus, in maximizing $Pr(\text{choice}_{ij}|\theta_{ij})$ to estimate θ_{ij} (equivalent to maximizing the likelihood function in equation 7.6), one cannot omit the marginal probability from the likelihood, because it would lead to biased estimates (similar to Manchanda et al. 2004). At the same time, modeling the marginal probability is challenging since the exact ranking algorithm used by the search intermediary is proprietary. However, when the ranking is randomly generated, the marginal probability of observing j in a given position is independent of θ_{ij} , making maximization of the conditional likelihood sufficient for unbiased estimates, which is the approach taken in this dissertation.

Price Endogeneity: The price of the hotel may be endogenous for two reasons. First,

2. I adapt this explanation from Nair et al. (2014).

an unobserved quality shock may affect both consumer choices and hotel prices. Second, consumer specific choice probabilities may affect what prices hotels set. The most common method to alleviate price endogeneity is using instrumental variables. One possibility is using Hausman (1996) style instruments to approximate marginal costs, such as the average price of the same hotel or of same star hotels across destinations. Neither instrument is available in this study: the former is not available since in the data I do not observed the same hotel in different destinations, while the latter is not available since information on destinations is anonymized, thus the average price of same star hotels across different destinations (for example, cities on different continents) may not recover the marginal costs of the focal hotel. Other possible instruments are lagged prices of the same hotel, which may not be valid if the unobserved quality of the hotel is correlated over time. Finally, another set of instruments are the average price of other hotels for the same trip and the focal hotel's non-price characteristics. These instruments have been used by Chen and Yao (2014) in the hotel industry, and by Hortaçsu and Syverson (2004) and Berry, Levinsohn and Pakes (1995) in different settings. They capture the position in characteristics space of the focal hotel relative to others, assuming that hotels' characteristics are predetermined. However, the last assumption may not be tenable in my data.

Even though price instruments in the hotel industry are difficult to obtain, concerns about endogeneity may be partially alleviated by the observation that prices, set by the hotel's revenue management system, are not set in response to individual consumers' preferences (Cross et al. 2009; Mauri, 2013; Koulayev 2014). In Table B.6 (Appendix B.8), using the Random ranking data, I show that observable hotel and query characteristics explain most of the variation in prices. More precisely, in the first column, I regress price on hotel and trip date fixed effects in the largest destination in the data and obtain an adjusted R^2 of 0.746. Therefore, specific dates command different prices, but all consumers searching for the same trip date see the same price for a particular hotel. In the next column, I add additional query characteristics, and include information about the average prices of similar

hotels for the same trip and obtain a larger adjusted R^2 of 0.815. In the last three columns of Table B.6, I show that a similar pattern holds across different destinations. This analysis suggest that observable characteristics explain most of the variation in price of a hotel, with the trip date explaining the majority of it. From discussions with an employee at a large hotel chain and from previous literature, I learned that the remaining price variation may be due either to (i) different suppliers selling the particular hotel, or to (ii) experimental price variation (Einav et al., 2015; Koulayev, 2014). Since both of these explanations are not demand-related, I conclude that, conditional on the query, the price variation observed is unlikely to be correlated with the utility error term and thus does not need an instrument.

8.2 Monte Carlo Simulation

In this section, I describe simulation results to show that Simulated Maximum Likelihood using the logit-smoothed AR simulator can be used to recover preference and search costs parameters in this model.³ To this end, I generate a data set of 1,000 consumers, each searching among five firms. Hotel characteristics are assumed to be drawn from a normal distribution with mean and standard deviation equal to those found in the data (see Table 4.1). The true values of the parameters are assumed to be consistent with those from a preliminary estimation of the model. For estimation, I use 50 draws from the distribution of the utility error terms for each consumer and hotel combination and I repeat the estimation 50 times. I have performed simulations with (inverse) scaling factor $1/\lambda$ ranging from 1 to 7 and found that $1/\lambda = 3$ works best in recovering the true parameters in the simulation sample, which motivated me to use the same scaling parameter in estimation.⁴ The simulation results are given in Table 8.1. The first column shows the true parameters and the second column shows the estimated parameters. I find that this method works well in recovering the parameters of interest.

3. I thank Elisabeth Honka for her hints on running this simulation.

4. Results with the remaining values of the scaling factor are similar and are available upon request.

Table 8.1: Monte Carlo Simulation Results

	True values	Estimated values
<i>Preferences (u)</i>		
Price	-2	-2.0109*** (0.0548)
Stars	0.2	0.2015*** (0.0259)
Review Score	-0.2	-0.1605*** (0.0231)
Location Score	0.2	0.2082*** (0.0246)
Chain	-0.2	-0.1856*** (0.0495)
Promotion	0.2	0.2135*** (0.0575)
<i>Search cost (c)</i>		
Position	2	1.7819*** (0.0036)
Constant k	-11	-9.5531*** (0.0030)
Log-likelihood		-1,545
Observations		5,000

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Having introduced the model and discussed its estimation and identification, in the next chapter, I present estimation results using the Random ranking data set.

CHAPTER 9

RESULTS

In this chapter, I present the main estimation results of the dissertation. To recover the preference and the search cost parameters, I estimate the sequential search model using Simulated Maximum Likelihood with the logit-smoothed AR simulator. I use 50 draws for each consumer-hotel combination of utility error terms and a scaling factor of $\lambda = 1/3$. In order to make the estimation feasible, I focus on the four largest destinations in the data.¹ This approach controls for destination specific differences, as well as shows that the results are not limited to a particular destination. Furthermore, I restrict attention to search impressions with opaque offers that have the same length (25, i.e. the average for the four destinations), to ensure that position effects are estimated consistently. Finally, I use the Random ranking data to eliminate the endogeneity bias in the ranking. The resulting estimation sample has a total of 67,075 observations and 2,683 search impressions. Each search impression has at least one click and there are 161 (6%) transactions.

Table 9.1 shows the main estimation results. In Panel A, I show the estimated coefficients, while in Panel B, I derive the magnitude of the position effect (the focus of this dissertation). In general, preference and search cost estimates are economically meaningful and significant. For example, search costs are higher for clicks occurring lower in the ranking (higher position) and consumers derive higher utility from cheaper hotels, with more stars and those that run a promotion. I find mixed evidence for the effect of the location score on choices (depending on the destination), as well as a negative effect of the chain dummy (insignificant for most destinations) and of review score (possibly due to nonlinearity of the effect).

Using the Random ranking data for estimation eliminates the endogeneity bias in the ranking and recovers the true position effect. To measure it, I compute the dollar equivalent of a change in search costs resulting from an increase in position by one. In Panel B, I find an average position effect of \$2.64, which ranges from \$0.84 to \$5.09 across the four

1. The four destinations are 8192, 4562, 8347 and 9402, described in Table B.7 in Appendix B.9.

Table 9.1: Main Estimation Results

Destination	(1)	(2)	(3)	(4)
<i>Panel A: Coefficients</i>				
<i>Preferences (u)</i>				
Price (\$100)	-0.2747*** (0.0355)	-0.1926*** (0.0301)	-0.1668*** (0.0380)	-0.2521*** (0.0468)
Stars	0.3554*** (0.0169)	0.0333 (0.0315)	0.1269*** (0.0305)	0.2071*** (0.0378)
Review Score	-0.2590*** (0.0422)	-0.0726*** (0.0171)	-0.0800** (0.0261)	-0.1156*** (0.0334)
Location Score	-0.1356*** (0.0366)	0.0882*** (0.0139)	-0.0230 (0.0225)	0.0481** (0.0183)
Chain	-0.0184 (0.0418)	-0.0359 (0.0435)	-0.1180* (0.0561)	-0.0265 (0.0711)
Promotion	0.1392*** (0.0404)	0.0028 (0.0418)	-0.0030 (0.0580)	0.0504 (0.0599)
<i>Search Cost (c)</i>				
Position	0.0065* (0.0025)	0.0225*** (0.0023)	0.0096** (0.0037)	0.0107*** (0.0032)
Constant	-1.1136*** (0.0032)	-1.1220*** (0.0009)	-0.9800*** (0.0055)	-0.7962*** (0.0031)
Observations	26,325	19,350	10,500	10,900
Log-likelihood	-3,404	-2,520	-1,381	-1,419
<i>Panel B: Equivalent Change \$</i>				
Position	\$0.84	\$5.09	\$2.45	\$2.19

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

destinations. As expected, these results are typically smaller than those found in the hotel industry literature, which does not have experimental variation in the ranking. For example, De los Santos and Koulayev (2014) find position effects ranging from \$7.76 to \$35.32, Ghose et al. (2012b) find a position effect of \$6.24, and Koulayev (2014) finds position effects that range from \$2.93 to \$18.78. The exception is Chen and Yao (2014) who find a position effect of approximately 25 cents. The difference likely comes from their ability to estimate the model at the consumer level, thereby aggregating different clicks made by the same consumer across search impressions, leading to a lower estimated search cost.

Regarding the magnitude of the baseline search cost level (k), as mentioned in Chapter 8 on identification, I expect that the lack of information on consumers' outside option, together with the fact that most search impressions contain only one click, and the inability to link different queries made by the same consumer will result in large search cost estimates that are difficult to interpret. Indeed, I find baseline search costs exceeding \$127. However, the coefficient of main interest, the position effect (δ), is correctly estimated from variation in where in the search impression the consumer clicks, independently of the total number of clicks (which identifies the baseline level of search costs). To support this claim, I show that when I restrict attention to search impressions with more than one click, search cost estimates decrease by as much as 89%, without significant changes in the position effect (see Appendix B.10). This result is comparable to the search cost of \$25 estimated by Chen and Yao (2014) and to the range of search costs from \$10 to \$50 estimated by De los Santos and Koulayev (2014).

In sum, eliminating the endogeneity bias in the ranking leads to a lower position effect, which I document. This result is consistent with the reduced form evidence presented in Section 5.1: rankings only affect consumer search, not their purchase decisions conditional on search, but in the absence of experimental variation, this effect is increased. Understanding the causal effect of rankings can help to improve rankings, as I show in the next chapter.

CHAPTER 10

COUNTERFACTUAL: UTILITY-BASED RANKING

In this chapter, I investigate the welfare effects of a utility-based ranking through a counterfactual. To perform the counterfactual, I compare the utility-based ranking built on the estimates of the search model with simulated choices under Expedia's ranking. This involves two steps. First, I reorder hotels displayed under Expedia's ranking by their estimated expected utility to construct the utility-based ranking, and then simulate consumer choices and measure their welfare under the resulting ranking. Second, I simulate consumer choices using the list of hotels displayed under Expedia's ranking and measure consumers' welfare. Note that in comparing the utility-based ranking with simulated choices under Expedia's ranking (instead of actual choices), I avoid the bias of a direct comparison of the two rankings, as discussed in Section 4.3. An alternative to this approach would be to approximate Expedia's ranking using data on which hotels it prioritizes and to simulate consumer choices and measure their welfare under the resulting ranking. However, this may introduce a bias, since approximating Expedia's ranking algorithm is challenging.

Consumer welfare from a ranking is defined as consumers' utility of the choice made under that ranking, net of their total search costs. To obtain consumers' choices under a ranking, I simulate their click and purchase decisions using the preference and search costs parameters estimated from the Random ranking in Chapter 9 and Weitzman's (1979) optimal search rules. To integrate over the unobserved component in consumers' utility function, I repeat the simulation 50 times and I report the average result. To further minimize the impact of the unobserved component, I restrict attention to consumers who purchase since their utility is less dependent on it.

In this dissertation, a utility-based ranking coincides with a ranking based on consumers' reservation utility. The reservation utility of a hotel equals its expected utility plus a function of search costs. Hence, the hotel with the highest expected utility assigned the top position in the ranking also has the highest reservation utility. As a result, under this utility-based

ranking, consumers find it optimal to search in order of position.

My results can be found in Table 10.1. The counterfactual results represent short-run effects of changing the structural parameters. I find that the utility-based ranking increases consumer welfare on average by \$28.69 (20.72% of the purchase price), ranging from an increase of \$9.89 (11.03% of the purchase price) to an increase of \$52.74 (28.41% of the purchase price). This corresponds to as much as a twofold increase in consumer welfare (for example, in the second destination I find an increase of 96.97%). A breakdown of this welfare shows that consumers benefit from better matches, as well as lower search costs. In particular, more than 40% of the increase in welfare comes from lower search costs¹, with the utility-based ranking decreasing the average position in which a transaction happens by approximately 4 positions.² This result confirms the importance of optimizing consumers' search decisions in improving rankings. Therefore, I conclude that ranking hotels by their expected utility allows consumers to find better matching products more quickly.

Table 10.1: Counterfactual Results

Destination	(1) U-E	(2) U-E	(3) U-E	(4) U-E
<i>Consumer</i>				
Change in Consumer Valuation	\$9.89	\$52.74	\$30.42	\$21.70
(% Transaction Price)	(11.03%)	(28.41%)	(29.63%)	(13.81%)
(% Change)	(82.28%)	(96.97%)	(56.39%)	(54.59%)
Match	\$6.56	\$23.22	\$7.61	\$6.57
Price	\$1.79	-\$8.60	-\$9.47	-\$4.17
Total Search Costs	-\$5.11	-\$20.93	-\$13.34	-\$10.96
Change in Transaction Position	-4.32	-3.74	-3.94	-4.39
<i>Search intermediary</i>				
% Change in Transactions	2.55%	7.00%	4.58%	4.30%
% Change in Revenue	4.59%	2.04%	-5.07%	1.53%

U=Utility-based ranking; E=Expedia ranking.

Note: Due to rounding, the change in consumer valuation may differ from the sum of match and search costs net of price.

1. For example, in the first destination, total search costs decrease by \$5.11, accounting for 51% of the increase consumer valuation of \$9.89.

2. Note that the change in total search costs is influenced both by the change in the position of a click, as well as by the change in the total number of clicks that consumers make.

At the same time, the utility-based ranking also benefits the search intermediary, with transactions increasing by at least 2.55%, while revenues increase by as much as 4.59%. This is expected, since using historical data to predict click and purchase probabilities may not always provide a good approximation to consumers' utility, and thus may not maximize purchase opportunities. In addition, contaminating this prediction with endogeneity bias by using estimates from consumer choices under a curated ranking, can further deteriorate the ranking. As a result, I conclude that a utility-based ranking provides a viable option for search intermediaries to improve their rankings.

In sum, I find that a utility-based ranking benefits both consumers and the search intermediary. This result is generally in line with the benefits of the utility-based ranking reported in the literature (Ghose et al., 2013; De los Santos and Koulayev, 2014; Chen and Yao, 2014). However, in terms of magnitude, I find larger benefits for consumers than previously perceived, likely due to the elimination of endogeneity bias, which results in precise estimates of consumers' search and preference parameters. This result further cements the value of the Random ranking in providing clean estimates that help to improve the ranking.

CHAPTER 11

LIMITATIONS AND FUTURE RESEARCH

In this chapter, I describe the limitations of this dissertation, as well as discuss future research directions. The current research could be improved in at least three directions, if the data limitations would be addressed. First, this data set does not allow comparing the two rankings in quantitative terms. If this data were available, one could quantify the endogeneity bias of the ranking, as well as evaluate different methods to eliminate this bias. This analysis would result in a method for resolving the endogeneity bias inherent in rankings without the need to conduct experiments. Second, the data set does not provide enough information to link different queries made by the same consumer. With this type of data, one could analyze the effect of rankings on consumer learning across search impressions, as well as provide a better measurement of the magnitude of consumers' search costs. Finally, consumers' click order is not observed in this data set. Since consumers choose the order in which they search, this type of data would allow one to account for click order in estimation and thus improve preference and search cost parameter estimates.

I also see two avenues for future research. First, as consumers transition to using mobile devices to access Internet content, intermediaries' rankings have to adjust to consumers' different behavior on these devices. For example, if the smaller screen size of mobile devices makes consumers focus even more on higher ranked products than on a computer (by increasing their search costs), then the need for a relevant ranking will be even greater on mobile devices. Further research should focus on the relation between consumer search costs and the impact of the ranking. Another avenue for future research is to investigate the consequences of search intermediaries ranking only a subset of products, such as independent hotels in the case of online travel agents. Even though such rankings decrease the diversity of the products displayed, which apparently may hurt consumers, they also substantially speed up consumers' search (making these rankings especially relevant for mobile devices). Thus, exploring this tradeoff can provide a new method to improve consumers' search experience.

CHAPTER 12

CONCLUSION

In this dissertation, I study the causal effect of rankings on consumer search and purchase decisions. This is challenging since rankings are endogenous. To resolve this problem, I use the first data set with experimental variation in the ranking from a field experiment at Expedia. Using this data set, I make three contributions. First, I show that rankings affect what consumers search, but conditional on search, do not affect purchases. I also exploit a feature of the data set (opaque offers), to show that rankings lower search costs, rather than consumer expectations or utility. Second, I quantify the effect of rankings using a sequential search model following Weitzman (1979) and find position effects lower than previous estimates in the literature obtained without experimental variation. Finally, I use the model's estimates to measure the welfare effects of a utility-based ranking, and find that it improves matches and lowers consumer search costs, as well as benefits the search intermediary.

REFERENCES

- [1] Agarwal A., K. Hosanagar, M. Smith. 2011. Location, Location, Location: An Analysis of Profitability of Position in Online Advertising Markets. *J. Marketing Res.* **48**(6) 1057-1073.
- [2] Ansari, A., C. Mela. 2003. E-customization. *J. Marketing Res.* **40**(2) 131-145.
- [3] Ansari, A., S. Essegai, R. Kohli. 2000. Internet Recommendation Systems. *J. Marketing Res.* **37**(3) 363-375.
- [4] Athey, S., G. Ellison. 2011. Position Auctions with Consumer Search. *Quart. J. Econ.* **126**(3) 1213-1270.
- [5] Athey, S., G. Imbens. 2015. Machine Learning Methods for Estimating Heterogeneous Causal Effects. Working paper.
- [6] Baye, M., B. De los Santos, M. Wildenbeest. 2014. What's in a Name? Measuring Prominence, and Its Impact on Organic Traffic from Search Engines. Working paper.
- [7] Berman, R., Z. Katona. 2013. The Role of Search Engine Optimization in Search Marketing. *Marketing Sci.* **32**(4) 644-651.
- [8] Berry, S., J. Levinsohn, A. Pakes. 1995. Automobile Prices in Market Equilibrium. *Econometrica.* **63**(4) 841-890.
- [9] Blake, T., C. Nosko, S. Tadelis. 2014. Consumer Heterogeneity and Paid Search Effectiveness: A Large Scale Field Experiment. Working paper.
- [10] Chan, T., Y. Park. 2015. Consumer Search Activities and the Value of Ad Positions in Sponsored Search Advertising. *Marketing Sci.* **34**(4) 606-623.
- [11] Chen, Y., S. Yao. 2014. Sequential Search with Refinement: Model and Application with Click-stream Data. Working paper.
- [12] Cross, R., J. Higbie, D. Cross. 2009. Revenue Management's Renaissance: A Rebirth of the Art and Science of Profitable Revenue Generation. *Cornell Hospitality Quart.* **50**(1) 56-81.
- [13] De los Santos, B., S. Koulayev. 2014. Optimizing Click-through in Online Rankings for Partially Anonymous Consumers. Working paper.
- [14] De Corniere, A., G. Taylor. 2014. Quality Provision in the Presence of a Biased Intermediary. Working paper.
- [15] Einav, L., T. Kuchler, J. Levin, N. Sundaresan. 2015. Assessing Sale Strategies in Online Markets Using Matched Listings. *Amer. Econ. J.: Microecon.* **7**(2) 215-247.

- [16] Ghose, A., P. Ipeirotis, B. Li. 2012a. Designing Ranking Systems for Hotels on Travel Search Engines by Mining User-Generated and Crowd-Sourced Content. *Marketing Sci.* **31**(3) 492-520.
- [17] Ghose, A., P. Ipeirotis, B. Li. 2012b. Surviving Social Media Overload: Predicting Consumer Footprints on Product Search Engines. Working paper.
- [18] Ghose, A., A. Goldfarb, S.P. Han. 2012c. How Is the Mobile Internet Different? Search Costs and Local Activities. *Info. System Res.* **24**(3) 613-631.
- [19] Ghose, A., P. Ipeirotis, B. Li. 2013. Examining the Impact of Ranking on Consumer Behavior and Search Engine Revenue. Working paper.
- [20] Ghose, A., S. Yang. 2009. An Empirical Analysis of Search Engine Advertising: Sponsored Search in Electronic Markets. *Manage. Sci.* **55**(10) 1605-1622.
- [21] Hagiu, A., B. Jullien. 2011. Why do Intermediaries Divert Search? *RAND J. Econ.* **42**(2) 337-362.
- [22] Hausman, J. 1996. Valuation of New Goods Under Perfect and Imperfect Competition. In T. Bresnahan, and R. Gordon, ed., *The Economics of New Goods, Studies in Income and Wealth* 207-248.
- [23] Yoganarasimhan, H. 2014. Search Personalization. Working paper.
- [24] Honka, E. 2014. Quantifying search and switching costs in the U.S. auto insurance industry. *RAND J. Econ.* **45**(4) 847-884.
- [25] Honka, E., P. Chintagunta. 2014. Simultaneous or Sequential? Search Strategies in the U.S. Auto Insurance Industry. Working paper.
- [26] Hortacısu, A., C. Syverson. 2004. Product Differentiation, Search Costs, and Competition in the Mutual Fund Industry: A Case Study of S&P 500 Index Funds. *Quart. J. Econ.* **119**(2) 403-456 .
- [27] Jerath, K., L. Ma, Y. Park, K. Srinivasan. 2011. A Position Paradox in Sponsored Search Auctions. *Marketing Sci.* **30**(4) 612-627.
- [28] Jeziorski, P., S. Moorthy. 2014. Advertiser Prominence Effects in Search Advertising. Working paper.
- [29] Jeziorski, P., I. Segal. 2012. What Makes them Click: Empirical Analysis of Consumer Demand for Search Advertising. Working paper.
- [30] Kim, J. B., P. Albuquerque, B. J. Bronnenberg. 2010. Online Demand under Limited Consumer Search. *Marketing Sci.* **29**(6) 1001-1023.
- [31] Kim, J. B., P. Albuquerque, B. J. Bronnenberg. 2014. The Probit Choice Model under Sequential Search with an Application to Online Retailing. Working paper.

- [32] Koulayev, S. 2014. Search for Differentiated Products: Identification and Estimation. *RAND J. Econ.* **45**(3) 553-575.
- [33] Manchanda, P., P. Rossi, P. Chintagunta. 2004. Response Modeling with Nonrandom Marketing-Mix Variables. *J. Marketing Res.* **41**(4) 467-478.
- [34] Manski, C., S. Lerman. 1981. On the Use of Simulated Frequencies to Approximate Choice Probabilities. In C. Manski and D. McFadden, ed., *Structural Analysis of Discrete Data with Econometric Applications* MIT Press, Cambridge, MA.
- [35] Mauri, A.G. 2013. Hotel Revenue Management: Principles and Practices. *Pearson*.
- [36] McFadden, D. 1989. A Method of Simulating Moments for Estimation of Discrete Response Models Without Numerical Integration. *Econometrica.* **57**(5) 995-1026.
- [37] Murthi, B., R. Rao. 2012. Price Awareness and Consumers' Use of Deals in Brand Choice. *J. Retailing.* **88**(1) 34-46.
- [38] Nair, H. S., S. Misra, W. J. Hornbuckle IV., R. Mishra, A. Acharya. 2014. Big Data and Marketing Analytics in Gaming: Combining Empirical Models and Field Experimentation. Working paper.
- [39] Narayanan, S., K. Kalyanam. 2014. Position Effects in Search Advertising: A Regression Discontinuity Approach. Working paper.
- [40] Rutz, O., R. Bucklin, G. Sonnier. 2012. A Latent Instrumental Variables Approach to Modeling Keyword Conversion in Paid Search Advertising. *J. Marketing Res.* **49**(3) 306-319.
- [41] Train, K. 2009. Discrete Choice Methods with Simulation. *Cambridge University Press*.
- [42] Varian, H. 2007. Position Auctions. *Int. J. Ind. Organ.* **25** 1163-1178.
- [43] Weitzman, M. 1979. Optimal search for the best alternative. *Econometrica.* **47**(3) 641-654.
- [44] Yang, S., A. Ghose. 2010. Analyzing the Relationship Between Organic and Sponsored Search Advertising: Positive, Negative, or Zero Interdependence? *Marketing Sci.* **29**(4) 602-623.
- [45] Yao, S., C. Mela. 2011. A Dynamic Model of Sponsored Search Advertising. *Marketing Sci.* **30**(3) 447-468.

APPENDIX A

DATA CLEANING

In this appendix, I describe the steps I took to clean the data set I used in this dissertation. The data set made available contains 9,917,530 observations. For my analysis, I filtered out the following three categories of observations, which reduced the data set to 4,503,043.

First, the data set contains some errors in the way price information was stored. For example, both very high (more than \$19 million per night) and very low (\$0.01 per night) prices appear in the data set. I corrected these errors for those consumers who made purchases by removing search impressions that contain at least one observation where the total amount spent exceeded the price paid multiplied by the length of the trip and the number of rooms booked plus taxes.¹

Second, I chose to focus on “typical” search impressions and removed those that include prices lower than \$10 or higher than \$1,000 per night. By eliminating these search impressions, I also mitigated the first problem above for queries not ending in a transaction.

Third, the original data set contains observations on more than 20,000 total destinations, with a median of 2 search impressions per destination. In order to observe enough variation in the position of a hotel under the Random ranking, I focused my attention on destinations with at least 50 search impressions.

1. In the US, hotel taxes range from 7% to almost 20% according to www.consumerreports.org/cro/news/2014/06/booking-a-hotel-these-cities-have-the-highest-hotel-taxes/index.htm, motivating my conservative approach of dropping observations where the implied tax is larger than 30%.

APPENDIX B

FURTHER EVIDENCE

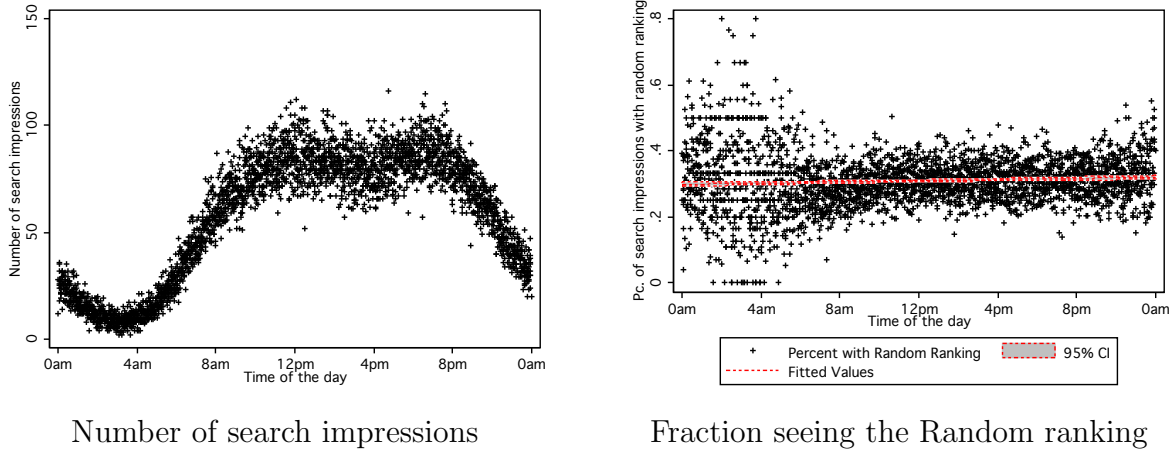
In this appendix, I provide further evidence for several chapters of this dissertation.

B.1 The Experiment

In this section, I show evidence for Chapter 4. In particular, I shows that (i) consumers were randomly assigned to the two types of rankings and that (ii) in constructing the Random ranking, the position of the hotel was randomly generated.

To show that consumers were randomly assigned to each type of ranking, I perform two tests. First, I test whether the time of arrival to Expedia’s website is related to the type of ranking the consumer saw. One concern may be that different types of consumers visit the website at different times of the day, which would bias the results if the probability of observing one type of ranking is different at different times of the day. For example, suppose business travelers search after 5pm and they have a higher probability to purchase. Then, if after 5pm the probability of observing Expedia’s ranking were higher (instead of the Random ranking), there would be a correlation between consumer choices and Expedia’s ranking in the data that is not due to the ranking observed, but rather to the way in which consumers were assigned to the two rankings. However, Figure B.1 shows that this is not a concern in the data. More precisely, the left panel plots the number of search impressions made by the time of the day and shows that more search impressions occur in the afternoon and evening than in the morning. The right panel plots the fraction of search impressions seeing the Random ranking every 30 seconds during the course of one day in the entire data set. It shows that the fraction of search impressions seeing the Random ranking is constant throughout the day, so that, even though more queries happen in the second part of the day, they have the same probability of seeing the Random ranking. Thus, these figures suggest that consumers were randomly assigned to seeing either type of ranking.

Figure B.1: Search Impressions with the Random Ranking



Second, I test whether consumer characteristics observed by Expedia prior to showing them a ranking are different between the two rankings. When the consumer arrives at Expedia’s website, she reveals details of her upcoming trip, such as her destination, the length of the trip, how long in advance she is searching for, the number of travelers and rooms requested, as well as whether her trip includes a Saturday night. For some consumers, Expedia also has historical information such as the average price and the number of stars of hotels previously purchased. One concern might be that this information affects consumers’ probability of seeing either ranking. Table B.1 shows that this is also not a concern. Comparing separately search impressions ending and not ending in a transaction across the two rankings by means of a t-test, I find that consumers seeing Expedia’s ranking have very similar characteristics to those seeing the Random ranking.¹ In particular, the difference in most characteristics is not statistically significant, while for booking window and trip length, which show a significant effect, the magnitude of the effect is small. For example, the difference in trip length is less than 0.2 days, while for booking window is less than 2 days (the average booking window in the sample is 39 days). Combined, these findings suggest that there are no systematic differences in consumer observables across the two rankings, confirming that consumers were randomly assigned either Expedia’s or the Random ranking.

1. I separate search impressions ending and not ending in a transaction to avoid biasing results since converting search impressions were oversampled.

Table B.1: Comparison of Consumer Observables: Expedia versus Random Ranking (T-test)

	(1)	(2)
	No Tran	Tran
Trip Length (days)	-0.1792*** (0.0275)	-0.0857*** (0.0255)
Booking Window (days)	-1.9284** (0.7103)	0.4653 (0.7476)
Adults	-0.0034 (0.0105)	-0.0475*** (0.0141)
Children	-0.0051 (0.0093)	-0.0205 (0.0125)
Rooms	-0.0034 (0.0052)	-0.0058 (0.0068)
Saturday Night	0.0126* (0.0056)	-0.0018 (0.0081)
Consumer Hist. Stars	0.0684 (0.0551)	0.0665 (0.0503)
Consumer Hist. Price	-9.3606 (7.7715)	-8.0906 (8.2443)
Observations	57,133	108,903

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

To show that hotels were randomly ordered under the Random ranking², I run a rank ordered logit regression of the position of the hotel in the ranking on past transactions and characteristics. My results can be found in Table B.2 below.³ I find that past transactions of the hotel or its characteristics do not determine its position within the Random ranking, while under Expedia’s ranking, there is a strong correlation between these characteristics and the hotel’s position, which is consistent with hotels being ranked randomly under the Random ranking. Moreover, this result provides insights into the construction of Expedia’s ranking: it favors non-chains that were purchased more often in the past, that are cheaper, of higher observed quality and are running a promotion.

2. See www.kaggle.com/c/expedia-personalized-sort/forums/t/5772/meaning-of-random-bool.

3. I restrict attention to the estimation sample to control for destination fixed effects (see Chapter 9 for details).

Table B.2: Estimates of Position by Ranking (Rank Ordered Logit)

Destination Ranking	(1) Random	(1) Expedia	(2) Random	(2) Expedia	(3) Random	(3) Expedia	(4) Random	(4) Expedia
Past Transactions	-0.5730 (0.4337)	7.8986*** (0.3350)	0.0637 (0.2902)	1.6914*** (0.1412)	-0.2174 (0.2784)	0.7765*** (0.1627)	0.3734 (0.3829)	2.7221*** (0.1970)
Price	-0.0004** (0.0001)	-0.0016*** (0.0001)	-0.0003** (0.0001)	-0.0019*** (0.0001)	-0.0003* (0.0001)	-0.0022*** (0.0002)	-0.0001 (0.0002)	-0.0034*** (0.0001)
Stars	0.0242 (0.0149)	0.3167*** (0.0155)	0.0188 (0.0137)	0.1942*** (0.0127)	0.0092 (0.0201)	0.1765*** (0.0201)	0.0248 (0.0200)	0.4838*** (0.0157)
Review Score	-0.0586* (0.0260)	-0.7605*** (0.0251)	-0.0123 (0.0093)	0.0409*** (0.0086)	0.0416** (0.0159)	0.2548*** (0.0226)	-0.0309 (0.0190)	0.1527*** (0.0180)
Chain	-0.0185 (0.0173)	-0.0277 (0.0158)	-0.0416* (0.0175)	-0.0674*** (0.0128)	-0.0133 (0.0297)	-0.1442*** (0.0215)	-0.0028 (0.0251)	0.0614*** (0.0175)
Location Score	-0.0493* (0.0231)	0.0295 (0.0217)	0.0037 (0.0062)	0.2774*** (0.0091)	0.0063 (0.0120)	0.0465*** (0.0124)	-0.0181* (0.0078)	0.1762*** (0.0091)
Promotion	-0.0000 (0.0153)	0.4241*** (0.0140)	0.0548** (0.0192)	0.2657*** (0.0128)	0.0658* (0.0270)	0.4882*** (0.0236)	0.0073 (0.0284)	0.4316*** (0.0163)
Log-likelihood	-61,057	-83,331	-44,878	-80,806	-24,351	-29,872	-25,285	-54,242
Observations	26,325	37,200	19,350	35,250	10,500	13,075	10,900	23,925

Standard errors in parentheses

Notes: Rank ordered logit regression with dependent variable position. A positive coefficient means correlation with the top of the ranking. Restrict attention to the data set used in estimation. Past transaction is defined as the total number of past transactions of a hotel normalized by the number of times it was displayed.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

B.2 Ranking Algorithm

In this section, I describe a basic *learning to rank* algorithm introduced in Chapter 3. *Learning to rank* algorithms are an example of machine learning algorithms with the goal of ranking documents based on relevance. In the current application, a document is a hotel and its characteristics. The system (e.g. the search intermediary) maintains a collection of these hotels and when a consumer makes a search query, it proceeds to rank them. More precisely, the system uses a so-called training data set, which contains consumer queries, hotel characteristics x and associated clicks and purchases. This latter information is interpreted as an indication of how relevant a hotel is for consumer queries and is known as the relevance score s of a hotel. This relevance score is highest if a purchase occurred, and lowest if no click occurred. The purpose of the *learning to rank* algorithm is to use the data on hotel j 's characteristics x_{ij} for a query performed by i and i 's clicks/purchases s_{ij} for each j and learn a function

$$f(x_{ij}, \delta) = \hat{s}_{ij}$$

so that the ranking order of predicted scores \hat{s}_{ij} are exactly equal to that of observed s_{ij} . More concretely, this means that the goal is to find a function that will rank the hotel that

will be purchased by the consumer at the top (the most relevant), followed by hotels that the consumer will click and finally, by those that will not be considered by the consumer. Note that a ranking model $f(x_{ij}, \delta)$ can be constructed even without learning, by only taking into account the characteristics of the hotel, such as the price and the number of stars. In contrast, a *learning to rank* model exploits the availability of data on relevance scores of consumers, which makes the ranking of the search intermediary endogenous.

A commonly used *learning to rank* algorithm is known as LambdaMART. It is described in more detail by Yoganarasimhan (2014).

B.3 The WCAI Data Set

In this section, I describe the companion data set from the Wharton Customer Analytics Initiative (WCAI). It provides information on consumer queries for hotels on a popular online travel agent’s (OTA) website between October 1 and 9, 2009.⁴ This data set cannot be used to study the causal effect of rankings since it does not include search impressions performed under a random ranking. However, it has information on some of the consumer groups that are excluded from the Expedia data set that I use. For example, it contains queries where consumers do not click, where they sort/filter search results, as well as where they consider the additional result pages (beyond the first page of results). I thus use the WCAI data to explore the impact of the absence of these groups from the analysis. I find that the groups that are missing generally constitute a small fraction of search queries, allowing me to state that their absence will not affect the representativeness of the Expedia data.

The WCAI data set includes search impressions in Paris, Budapest, Cancun and Manhattan, but for the current analysis I focus on the Manhattan data set. This data contains 431,820 observations and 15,171 search impressions made by 4,752 consumers. Table B.3 compares three data sets: the Expedia data, the WCAI data, and the WCAI data set containing search impressions with *at least* one click (similar to the Expedia data set), which

4. I cannot disclose the name of the OTA that provided the data.

I call WCAI-Click. The hotels displayed in the three data sets have similar characteristics (except that those in the Expedia data are cheaper, most likely due to the fact that it contains a much greater pool of destinations (788) and hotels (54,877) than those found in Manhattan).

Comparing the three data sets, I draw several conclusions. First, only 23% (3,560 out of 15,171) of the queries in the WCAI data set have at least one click. However, the Expedia data set, which only contains search impressions with at least one click, has a similar average number of clicks as the WCAI-Click data set (1.12 compared to 1.80), indicating that for the group of queries that have at least one click, the Expedia data set is representative.

Second, in the Expedia data set, two thirds of search impressions end in a transaction, while in the WCAI (WCAI-Click) data only 1% (3%) lead to a transaction. However, as discussed in Section 4.3, search impressions ending and those not ending in a transaction in the Expedia data set were sampled randomly, thus not affecting the main analysis.

Third, in the Expedia data set, there are no queries in which consumers sort or filter. However, in the WCAI or WCAI-Click data sets, very few consumers actually make these choices: in the WCAI data set, out of all search impressions, only 4,860 contain filtered results (32%), while only 1,601 consumers filter (34%). Therefore, most consumers make choices from the ranking that is displayed, instead of sorting/filtering the hotels observed, making the Expedia data set representative.

Finally, in the Expedia data set, I only observe the first page of results that consumers saw in a search impression. However, the WCAI and the WCAI-Click data sets show that most consumers only consider the first page of results: out of all search impressions, 10,228 (67%) choose to only consider the first page of results, while out of all consumers, 3,027 (64%) only consider the first page of results. Thus, only observing consumer choices from the first page of results captures most consumers' behavior on an OTA.

The comparison of the Expedia and the WCAI data sets allows me to I state that the former data set is representative of most consumers searching on an OTA.

Table B.3: Summary Statistics Comparison with the WCAI Data Set (Mean)

	Expedia	WCAI	WCAI If Click
<u>Hotel level</u>			
Price	159.71	281.50	281.63
Stars	3.32	3.08	3.16
Review Score	3.89	4.01	4.03
Chain	0.66	0.35	0.37
Promotion	0.25	0.35	0.28
<i>Observations</i>	<i>4,503,043</i>	<i>431,820</i>	<i>108,066</i>
<u>Search impression level</u>			
Hotels Displayed	27.12	25.00	25.00
Total Clicks	1.12	0.42	1.80
Total Transactions	0.66	0.01	0.03
Sort/filter	0%	32%	33%
First Page	100%	67%	68%
Random ranking	31%	0%	0%
<i>Observations</i>	<i>166,036</i>	<i>15,171</i>	<i>3,560</i>
Number of hotels	54,877	301	300
Number of destinations	788	1	1

B.4 Position Effect under the Random Ranking

This section provides robustness checks for the causal effect of rankings identified in Section 5.1. It shows that the same patterns as in Figure 5.1 hold when restricting attention to:

- small destinations with less than the median number of hotels in the sample, to control for the fact that in large destinations sponsored ads are more likely: Figure B.2
- search impressions with more than 30 hotels displayed, to control for the fact that some search impressions have few hotels displayed: Figure B.3
- search impressions ending in a transaction: Figure B.4


Figure B.2: Position Effect for Destinations with Few Hotels



Figure B.3: Position Effect for Search Impressions Longer than the Median



Figure B.4: Position Effect for Search Impressions Ending in a Transaction



Note that the scales of the y-axis of the two figures are different.

B.5 Position Effect under Expedia's Ranking

This section provides regression results showing the effect of position under Expedia's ranking. More precisely, it shows that under Expedia's ranking, higher ranked hotels lead to more clicks and purchases, as well as more purchases conditional on a click. This result confirms the data patterns shown in Figure 5.2 in Section 5.1.

Table B.4: Estimates of Clicks and Transactions (OLS): Expedia Ranking

	Click		(3)	Transaction	Transaction conditional on click		
	(1)	(2)		(4)	(5)	(6)	(7)
<i>Position effect</i>							
Position	-0.00308*** (0.00001)	-0.00247*** (0.00003)		-0.00199*** (0.00003)	-0.00391*** (0.00016)	-0.00196*** (0.00037)	
Positions 1-5			0.08917*** (0.00137)				0.04825*** (0.01003)
Positions 6-10			0.03523*** (0.00095)				0.00581 (0.01111)
Positions 11-15			0.01781*** (0.00090)				0.01176 (0.01290)
Positions 16-20			0.00947*** (0.00081)				-0.02488 (0.01469)
<i>Hotel characteristics</i>							
Price		-0.00015*** (0.00000)	-0.00015*** (0.00000)	-0.00012*** (0.00000)		-0.00016** (0.00006)	-0.00016** (0.00006)
Stars		0.00622*** (0.00053)	0.00662*** (0.00052)	0.00469*** (0.00045)		-0.00529 (0.00668)	-0.00550 (0.00667)
Review Score		0.00058 (0.00055)	0.00050 (0.00055)	0.00070 (0.00048)		0.01008 (0.00704)	0.01003 (0.00703)
Chain		0.00290*** (0.00071)	0.00285*** (0.00071)	0.00248*** (0.00063)		0.00258 (0.00817)	0.00267 (0.00816)
Location Score		0.00367*** (0.00029)	0.00359*** (0.00029)	0.00367*** (0.00024)		0.02878*** (0.00441)	0.02888*** (0.00440)
Promotion		-0.00037 (0.00071)	-0.00140* (0.00071)	0.00057 (0.00062)		0.01894* (0.00784)	0.01793* (0.00783)
Query characteristics	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Destination FE	No	Yes	Yes	Yes	No	Yes	Yes
Adjusted R^2	0.026	0.024	0.029	0.020	0.008	0.036	0.037
Observations	2,522,288	377,224	377,224	377,224	104,874	14,690	14,690

Standard errors in parentheses

Notes: In columns (3) and (7), the omitted category is [Positions 21-end]. Standard errors clustered at the search impression level. Restrict sample to search impressions with opaque offers. To control for destination FE, I restrict attention to destinations with at least 10,000 observations. Query characteristics included in all regressions: trip length, booking window, number of adults and children traveling, number of rooms, and Saturday night dummy.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

B.6 Illustration of Opaque Offers

In this section, I provide an illustration of opaque offers that are analyzed in Section 5.2.

Figure B.5: Illustration of Opaque Offers

The figure illustrates the concept of opaque offers on a travel website. It shows two side-by-side screenshots of a search results page for hotels in Chicago. The left screenshot, labeled 'No opaque offer', shows a list of hotels including The Peninsula Chicago, Hyatt Place Chicago-South/University Medical Center, Chicago Lake Shore Hotel, La Quinta Inn & Suites Chicago - Lake Shore, and Amber Inn. The Amber Inn listing shows a price of \$100. The right screenshot, labeled 'With opaque offer', shows the same list but with a red box highlighting an offer: '4.5-Star Hotels from \$100 near Chicago (and vicinity)'. This offer is described as 'Expedia Unpublished Rate Hotels' with 'Deep discounts on quality hotels' and 'Get the Deal now - and the Hotel Name after you book'. A blue arrow points from the Amber Inn listing in the left screenshot to the opaque offer in the right screenshot, indicating that the offer is a hidden discount on a hotel that would otherwise be visible.

No opaque offer

With opaque offer

B.7 Evidence on Consumer Click Order

In this section, I show that the position of a hotel in the ranking is a good predictor of the order in which consumers click. To do so, I use the companion data set from WCAI, which contains information on consumers' click order in the form of time stamps associated with each click. I then ask what fraction of search impressions with at least two clicks had a click order that matched the position of the hotels clicked. I also compare this fraction with that ordered by price. Table B.5 shows my results. I find that in 35% of all search impressions with at least two clicks and in the majority of search impressions (65%) with exactly two clicks the position of the hotel exactly matches the click order of the consumer. In contrast, the price of the hotels clicked only matches the order of 20% of the clicks. This finding allows me to approximate consumers' click order with data on the position in which their click occurred.

Table B.5: Evidence on Consumer Click Order from the WCAI Data Set

	Percentage		
	Price	Position	Price or Position
Search impressions with at least two clicks	20	35	40
Search impressions with exactly two clicks	49	65	77

B.8 Price Endogeneity Concerns

This section provides regression results mentioned in Chapter 8, showing the effect of observable hotel and query characteristics on hotel prices.

Table B.6: Estimates of Hotel Prices (OLS)

Destination	(1)	(1)	(2)	(3)	(4)
Hotel and trip date FE	Yes	Yes	Yes	Yes	Yes
<i>Query characteristics</i>					
Trip Length (days)		0.7090*** (0.1989)	0.1863 (0.2381)	0.2468 (0.3751)	0.8082* (0.3548)
Adults		0.6894* (0.3451)	9.1458*** (0.6202)	6.3270*** (0.6864)	8.5790*** (1.1580)
Children		1.6298** (0.5502)	7.2412*** (0.8576)	4.4478*** (0.6353)	5.5602*** (1.0746)
Rooms		0.9244 (0.8875)	-2.7969* (1.2859)	-3.4720 (1.8805)	-9.3391*** (2.7853)
Saturday Night		12.9115*** (1.5119)	-5.9249** (1.9529)	8.5286** (2.7417)	-3.2758 (3.3133)
Booking Window (days)		-0.1054*** (0.0075)	0.0276* (0.0132)	-0.0252 (0.0182)	0.0550*** (0.0161)
9am-6pm		0.7679 (0.9302)	-1.5403 (1.4026)	0.1794 (1.7561)	-3.2830 (1.8978)
6pm-midnight		-0.0233 (0.9721)	-6.0739*** (1.6059)	-4.2705* (2.0372)	-1.6627 (2.0652)
Weekend		1.5421* (0.6882)	-2.9467* (1.1612)	1.5573 (1.6241)	-6.7465*** (1.6480)
<i>Competition</i>					
Avg. prices of similar hotels		-2.1410*** (0.0249)	-1.1359*** (0.0181)	-1.0116*** (0.0261)	-1.0976*** (0.0246)
Promotion		-19.1166*** (0.7218)	-26.2001*** (1.2646)	-20.8092*** (2.0716)	-28.6557*** (1.8307)
Adjusted R^2	0.746	0.815	0.849	0.850	0.819
Observations	26,321	25,503	16,689	8,166	9,056

Standard errors in parentheses

Notes: OLS regression with dependent variable price. Time of day of the search is with respect to the left out variable: searches performed between midnight and 9am (local time). The average price of similar hotels is computed as the average price of hotels of same type (chain vs independent) with the same number of stars and reviews as the focal hotel for the same trip date (excluding the focal hotel). I restrict attention to hotels that are displayed at least 15 times to be able to include hotel fixed effects in all specifications above.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

B.9 Summary Statistics for the Largest Destinations in the Data

This section complements Chapter 9 by summarizing hotel and search impression characteristics for the four largest destinations in the data.

Table B.7: Hotel and Search Impression Summary Statistics: Largest Destinations (Mean)

Destination	(1)	(2)	(3)	(4)
<i>Hotel level</i>				
Price	142.85	250.46	203.38	144.34
Stars	4.01	3.62	3.45	3.43
Review Score	4.03	3.97	4.05	4.05
Chain	0.80	0.59	0.68	0.71
Location Score	4.05	4.68	4.16	2.80
Promotion	0.64	0.39	0.32	0.46
<i>Observations</i>	<i>63,535</i>	<i>54,600</i>	<i>34,835</i>	<i>23,575</i>
<i>Search impression level</i>				
Number of Hotels Displayed	25.00	25.00	25.00	25.00
Trip Length (days)	2.71	3.18	2.67	3.74
Booking Window (days)	51.43	48.31	46.71	47.12
Saturday Night (percent)	0.49	0.44	0.48	0.37
Adults	2.16	1.92	1.98	2.33
Children	0.24	0.28	0.36	0.98
Rooms	1.15	1.08	1.09	1.10
Total Clicks	1.09	1.13	1.09	1.11
Two or More Clicks (percent)	0.06	0.07	0.06	0.05
Total Transactions	0.55	0.56	0.64	0.47
Random ranking (percent)	0.41	0.35	0.31	0.45
<i>Observations</i>	<i>2,541</i>	<i>2,184</i>	<i>1,393</i>	<i>943</i>

B.10 Magnitude of Search Costs

In this section, I show evidence that the fact that search impressions contain mostly one click, together with the fact that the data set does not provide enough information to link search impressions made by the same consumer lead to large search cost estimates. To this end, I re-estimate the model on two data sets. First, in Table B.8, I estimate the model on the subset of search impressions that contain at least two clicks. Compared to the baseline search cost coefficient of -1.1136 in Table 9.1 for the Random ranking in the first destination, the coefficient is twice as large and equals -2.2885 . Equivalently, in dollar terms, the baseline level of search costs decreases by as much as 89.55%, making these estimates comparable to those in the literature.

Second, I re-estimate the model on the WCAI companion data set for Manhattan. The advantage of this data set is that it includes information that allows linking different queries made by the same consumer. In the table below, I show that most search impressions (77%) contain no clicks and that the average (median) number of clicks in search impressions with at least one click is 1.80 (1), comparable to the average (median) of 1.12 (1) in the Expedia data set. However, this number is significantly increased if information at the consumer level is available. More precisely, consumers who make at least one click, make on average 3.28 clicks (median is 2 clicks). Thus, estimating search costs at the level of a search or at the consumer level may significantly affect results, as I show in Table B.10.⁵ I estimate the model twice: first, assuming no information on consumers is available (column 1) and then linking search impressions made by the same consumer (column 2). The results in column 1 mirror those in Table 9.1 showing relatively large search costs of \$71.37. In contrast, at

5. The WCAI data set is also used by Chen and Yao (2014). To make the results comparable to theirs, I estimate the model on search impressions ending in a transaction (for a comparison, see Table 6 in Chen and Yao, 2014). Note that this assumption requires normalizing the price coefficient to -1 for identification. However, there are two notable differences in my estimation compared to that of Chen and Yao (2014): (i) the models we estimate are slightly different, as in their model the click probability is separate from the purchase probability; and (ii) I focus the estimation on one destination, Manhattan, while they estimate their model on all four destinations available.

Table B.8: Estimation Results: Search Impressions with at least Two Clicks

Destination	(1)	(2)	(3)	(4)
<i>Panel A: Coefficients</i>				
<i>Preferences (u)</i>				
Price (\$100)	-0.2001 (0.1115)	-0.3039*** (0.0752)	-0.4991** (0.1786)	-0.4662 (0.2688)
Stars	0.2031*** (0.0404)	-0.0162 (0.0651)	0.0648 (0.0562)	0.1669 (0.1281)
Review Score	-0.1615 (0.1322)	-0.0801 (0.0510)	-0.0760 (0.0836)	-0.1473 (0.1442)
Location Score	-0.2441* (0.1109)	0.0674*** (0.0195)	-0.0521 (0.1033)	0.0681 (0.0611)
Chain	0.0613 (0.1264)	-0.0635 (0.1119)	-0.0521 (0.2853)	0.0681 (0.3174)
Promotion	0.0573 (0.1075)	-0.0789 (0.1244)	0.0873 (0.1926)	0.0214 (0.1616)
<i>Search Cost (c)</i>				
Position	0.0048 (0.0112)	0.0104 (0.0105)	-0.0064 (0.0194)	-0.0133 (0.0177)
Constant	-2.2885*** (0.0052)	-2.2012*** (0.0016)	-1.9339*** (0.0175)	-1.6052*** (0.0058)
Observations	1,725	1,100	375	575
Log-likelihood	-255	-164	-56	-84
<i>Panel B: % Change cf. Full Sample (\$)</i>				
Baseline Search Cost	-58.48%	-81.60%	-89.55%	-82.36%

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

the consumer level in column 2, estimated search costs are 63% lower, equaling \$30.20. This estimated search cost is comparable to the one found in Chen and Yao (2014) of \$25.28. I conclude that the magnitude of search costs estimated in Table 9.1 is inherited from the limitations of the data set.

Table B.9: WCAI Data set: Number of Clicks at Search Impression and Consumer Level

Number of clicks	Search impression level	Consumer level
0	11,611 (77%)	2,793 (59%)
1	2,304 (15%)	880 (19%)
2	683 (5%)	392 (8%)
>2	573 (3%)	687 (14%)
Avg. clicks if at least 1	1.80	3.28
Median clicks if at least 1	1	2

Table B.10: WCAI Data Set: Estimation Results

<i>Panel A: Coefficients</i>	Search Impression Level	Consumer Level
<i>Preferences (u)</i>		
Price (normalized)	-1	-1
Stars	56.0870*** (7.2774)	67.9450*** (6.0577)
Review Score	51.7600*** (6.7878)	29.8880*** (7.5547)
Chain	3.0093 (6.6991)	9.0931 (5.4615)
Promotion	-19.4270** (6.8075)	-14.6760** (5.2664)
<i>Search Cost (c)</i>		
Position	0.0196*** (0.0000)	0.0218*** (0.0047)
Booking Window	-0.0118*** (0.0000)	-0.0572*** (0.0009)
Constant k	4.2678*** (0.0000)	3.4078*** (0.0100)
Observations	665	6,740
Log-likelihood	-171	-677
<i>Panel B: \$</i>		
Baseline search costs	71.37	30.20

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$