THÈSE

En vue de l'obtention du

DOCTORAT DE L'UNIVERSITE DE TOULOUSE

Délivré par l'Université Toulouse Capitole

École doctorale : **Sciences Economiques-Toulouse School of Economics**

Présentée et soutenue par

PETROVA Milena

le 14 décembre 2018

Essays in Empirical Industrial Organization

Discipline : **Sciences Economiques**

Unité de recherche : **TSE-R (UMR CNRS 5314 – INRA 1415)**

Directeur de thèse : Christian BONTEMPS, Professeur d' économie, TSE

JURY

Rapporteurs Madame Helena PERRONE, Professeure d'économie à l' Université Pompeu Fabra
Madame Grazia CECERE, Professeure d'économie à l' Institut Mines Telecom

Suffragants Monsieur Wilfried SAND, Professeur d'économie à TSE
Monsieur Christian BONTEMPS, Professeur d'économie à TSE

Essays in Empirical Industrial Organization

Toulouse School of Economics

Toulouse School of Economics
École d'économie de Toulouse

Milena J Petrova

December 10, 2018

Acknowledgement

I regard my as a doctoral student at the Toulouse School of Economics (TSE) as one of the most exciting and intellectually challenging times of my life. I had the chance to be part of a stimulating and creative environment, populated with some of the brightest researchers and the most motivated students in the field of Economics. The thesis presented here has been made possible thanks to the contribution of support of various people at TSE and outside.

First and foremost, I am indebted to Jacques Crémer, the Director of the Digital Chair at the Jean Jacques Laffont Foundation, for giving me the chance to work on what I am passionate about and to show what I am capable of. Without his clear vision and encouragement completing this thesis would not have been possible, and through his example he has inspired me to grow as a person and as a professional.

I am sincerely grateful to Christian Bontemps for his dedicated advice, guidance and support during the very critical period that was the last year of my doctorate. I have also been fortunate to collaborate with and learn from Miguel Urdanoz, whom I count as a dear friend. My work has benefited greatly from the constructive and earnest discussions with Shuti Singha and Daniel Ershov. I am deeply appreciative of my fellow PhD students, Laura Lasio, Tuba Tuncel, Vitalijs Jascisens, Jakob Henig and Emil Palikot, for being outstanding colleagues and friends.

I would like to thank my jury - Helena Perrone, Grazia Cecere and Wilfried Sant-Zantman - for taking on this responsibility and for their invaluable insight on my work. Two of the chapters in this thesis are based on data generously provided by MaistorPlus, for which opportunity I am thankful to Stoyan Irinkov and his team.

Last but not least, my thesis is dedicated to my parents, Yulia and Jan, and to Thomas.

Summary

In the past couple of decades, digitization has affected the strategy of economic players and the structure of markets across the board by lowering the cost of storing, sharing and analyzing data. This has given rise to a new field of economics, the economics of digitization, which touches upon the fields of industrial organization, market design, information economics, and labor economics. For industrial economists, these new questions and challenges coupled with new types of data, have led to vigorous research on the topics of reputation, search, rankings, matching, and online auctions. Following this line of research, the first two of the chapters in my thesis are on the topics information frictions and reputation systems in online service markets, and the third chapter proposes a novel methodology for modeling transaction prices motivated by competition on online distribution channels.

Information Frictions on an Online Services Market

Information frictions can significantly affect the level of economic activity and efficiency, and this is especially true in services markets because of the number of dimensions along which the economic agents have to transmit information. Consider the home repairs market, where traditionally there is no central marketplace where homeowners and repairmen meet. A homeowner would have to spend time and effort to lookup in the yellow pages, ask around for recommendations, or search in other ways for repairmen. Next, he would have to evaluate their quality, either through their reputation or self-reported experience. Another complication arises because the repairmen have limited capacity, and whether they are available is privately observed. If available, their willingness to take on a project depends on its characteristics relative to other potential projects. Most of the time, homeowners do not have the time and resources to search, and they accept the terms of the repairmen. The substantial information frictions cause service markets to be characterized by lowered transaction levels, costly search and inefficient matches. Because of this and the little data availability, economists had little opportunity to study these markets in the past.

Service markets have greatly benefited from digitization enabling the low cost transmission of information in online marketplaces. Platforms such as Uber and AirBnB, valued at billions, would not exist in the 'brick and mortar' economy. Economists became concerned with understanding how economic agents interact in these markets, and how the

different sources of information frictions affect market outcomes. Cullen and Farronato (2015), Fradkin (2017) and Horton (2017) are notable contributions to this small and ambitious literature. Their work considers how service buyers search for service sellers and how a match is eventually formed either at the aggregate level and without agent heterogeneity, or using reduced form approaches.

In this paper, I construct and estimate a structural model at the level of the individual service buyer-seller interaction, which involves two stages: search and matching. The differentiation on both the supply and demand side, and the one-to-one assignment of service buyer to service seller make service markets akin to matching markets. I use the estimated model primitives to study the effect of information transmission features in reducing frictions by constructing counterfactual platform designs. Comparing their outcomes, I quantify the effects on the probability and intensity of search, as well as the probability and efficiency of a match. My model is the first to specifically consider how seller fixed capacity affects the economic outcomes, both through seller availability and seller willingness to transact.

To estimate the model, I use data from a home services platform, MaistorPlus. The data contains information on the characteristics of buyers and sellers, the process of search, and whether there is a match. As prices are formed outside the platform, after a process of comparing offers and negotiation, the data does not contain offer or transaction prices. While price data can be incorporated, my model can be used in settings where the researcher does not observe prices, either because he does not have access to company data or because this data is private information of the economic agents as in my case. The model can also be used when the researcher has access to price data but is not allowed to use it for economic research, as in Fradkin (2017) and Horton (2017).

After estimating the model fundamentals, I construct two information-poor scenarios. I take two steps back and consider the marketplace with no reliable information on seller characteristics, eliminating the platform-verified seller profile and reputation, and no information on whether the seller is currently available. This scenario is akin to choosing a service provider from the yellow pages. The resulting level of economic activity is extremely low, indicating information on who the available sellers are is absolutely crucial. In the second information-poor scenario, the buyer knows which sellers are available but cannot differentiate between them. Compared to the baseline of MaistorPlus, the search probability for buyers is two-thirds lower but the search intensity is two times higher. Information on seller characteristics, at the same time, does not affect outcomes conditional on search taking place: the probability of a match and match efficiency are comparable.

Finally, I consider a full information scenario, where the last remaining friction - the seller reservation price for matching with the given buyer - is eliminated. This scenario represents a limit on the short-run economic activity due to sellers relying on demand from outside the platform. The probability of a match is double and the matches are 22 percent more efficient than the baseline.

Reputational Incentives under Heterogenous Demand Fluctuations

Joint work with Jakob Henig, PhD, TSE

As economic activity moves online, the majority of transactions are anonymous and one-time only. Under these conditions, sellers have an incentive to cheat when the quality of the good or the level of effort is private information. This is why the success of online platforms such as eBay and Amazon, but also Uber and AirBnB more recently, can be credited to a large extent to their successful employment of reputation systems. As in the 'brick and mortar' economy, online reputation systems, or mechanisms, create dependence between current and future demand, this way disciplining sellers. What is new in the digital environment is the platform's ability to design and optimize such systems, in particular with respect to how reviews are aggregated and displayed. Athey and Luca (2018) note that platforms have to consider more carefully the way in which review information is made available to buyers, and how this affects seller incentives.

The objective of this paper is to document how a simple reputation system, not adapted to the economic environment, can fail to provide incentives that are consistent across sellers and uniform across time. Similarly to other online platforms such as GoogleMaps and Yelp, the reputation system of MaistorPlus does not discriminate between reviews when aggregating and displaying them, which may not be optimal. In the home services sector, there are significant demand fluctuations due to seasonality of certain types of demand, to which the sellers are heterogeneously exposed. The returns to reputation also fluctuate with the upcoming demand, making current effort more or less valuable for sellers depending on their idiosyncratic demand conditions. We test this hypothesis using a difference in differences econometric model. The results, supported by robustness checks, document a lower level of effort, measured by the probability of a positive review, at the end of the high demand season for sellers experiencing high demand fluctuations. The same sellers would exert higher effort before the start of the high demand season. This result can inform a better design of the MaistorPlus reputation system, as the platform has access to the seasonality information. More generally, this result brings attention to the issue of optimal review aggregation and display.

Airline Cooperation Effects on Airfare Distribution: An Auction Approach

Joint work with Prof. Marc Ivaldi, TSE and Prof. Miguel Urdanoz, TBS

The third paper in my dissertation proposes a new method to analyze transaction prices and is motivated by many markets moving product distribution to online sales channels. Competition online can be considered as more price-focused because transport or search costs are much lower, and buyers can easily 'comparison shop' using dedicated websites. The airline industry has been among those most affected, with virtually all transactions moving online, and the widespread use of meta-search engines such as Kayak (founded 2004) and Priceline (founded 1997). In any given moment, sellers are competing in prices in an open environment which is equivalent to a reverse English auction (Klemperer (2004)). According to Klemperer, internet markets may be thought of as auction markets under certain conditions, and their analysis and estimation can be aided by auction theory and econometrics.

Under this premise, we propose an econometric model based on the equivalence between competition online and auctions, and estimate it using data from the airline industry. We use transaction prices from the DB1B survey, an industry standard used in most research studying airlines. Our application considers how airline alliances affect the transaction price distribution and we contribute to the literatures on airline alliances and on price dispersion. We are the first to jointly model price levels and price dispersion in alliance markets, while the relevant literature typically focuses on price levels only. This allows us to comment on how alliances affect price variability and to consider a link to price discrimination. We are also the first to study cooperation, in the form of alliances, within the literature on price dispersion that typically focuses on the effect of competition. Our results indicate that alliances are better at price discriminating, as dispersion is higher in alliance markets, but they pass efficiency gains to consumers as average prices are lower in alliance markets.

Contents

List of Figures

List of Tables

Chapter 1

Information Frictions on an Online Services Market

Abstract

Service markets are typically characterized by significant information frictions because there are many small and heterogeneous buyers and sellers, but also because seller capacity limited and private information. Surprisingly, these frictions can persist online and can have substantial effect on market outcomes. This paper contributes to the young and growing empirical literature studying information frictions in online service environments by proposing for the first time a structural framework that makes explicit the process by which individual buyers and sellers search and match, as well as the different sources of information frictions. Simulating coutherfactual scenarios, I evaluate the effect of these frictions on buyer search probability and intensity, and on match probability and surplus.

1.1 Introduction

While e-commerce is synonymous with the incredibly successful goods marketplaces Amazon and eBay, in recent years service platforms have also gained prominence. The two most well-known such platforms are AirBnB and Uber, pioneers of the asset sharing economy. In other service sectors, there are Upwork (for IT and business services, formally Elance and oDesk), Thumbtack (for diverse local services, for example music lessons or party DJ-ing), and TaskRabbit (for low skill domestic services).

Compared to goods markets, service markets are more susceptible to inefficiency stemming from information frictions. One reason is because there are many small and heterogeneous sellers and buyers.[1] Consider the home repairs market, with no "brick and mortar" marketplace where buyers and sellers meet to announce their demands and offers. When a homeowner requires a service, he starts looking for the service providers: he can contact a past provider, asks friends, or look at the yellow pages. The homeowner then may compare the sellers based on information about their experience and expertise. On their end, the service providers are typically small and may not be available. Their willingness to take on a job depends on its characteristics, and how it compares on the potential jobs they would be missing out on.

In many service markets, transmitting and accessing information about characteristics, availability and preferences is not straightforward. Indeed, the information frictions may be so significant that some service markets, such as home (AirBnB) and car sharing (Uber) services, did not exist before the "moving" online. Online service markets are a meeting place, they maintain public seller (and buyer) profiles and reputation systems, and transmit information on availability whenever possible. Nevertheless, substantial information frictions remain even online, especially with relation to the private seller information on availability and willingness to transact. Horton (2018) documents that on oDesk, buyers' search and contacting efforts to lead to a rejection in 45 percent of the times, due to sellers are not available (48 percent) or not willing (29 percent). On AirBnB, seller availability is public information by means of a calendar, yet buyer requests are still rejected in 42 percent of the time, and mainly because the sellers are not willing to provide the service (Fradkin (2017)). Both Fradkin (2017) and Horton (2017) show that information frictions impact the number of matches formed on the platform negatively: buyers are more likely to stop searching and exit the platform after a rejection.

[1] Through out the paper, I refer to the homeowners or platform clients as the buyers, and to the service provider or professionals as sellers.

With this paper, I propose for the first time a structural model of the buyer-seller interaction, which makes explicit the different sources of information frictions and allows for straightforward evaluation of their economic impact through construction of counterfactual scenarios. The model uses data from the marketplace MaistorPlus, an online home services market from Bulgaria, and can potentially be adapted to similar environments. MaistorPlus functions in the following way: the buyers post a job which is sent to all subscribing sellers; in turn, the available sellers make themselves known to the buyer by sending him a massage. The buyer inspects available seller profiles and contacts a number of them, they discuss offers, after which a buyer-seller pair may reach an agreement. Because the seller capacity is small and limited, their willingness to transact depends on the value of other potential jobs they could be doing, or their *reservation price*. On MaistorPlus, the reservation price is based on jobs from traditional channels of demand, such as client referrals. However, the reservation price as a measure of seller willingness is common for many other service markets with capacity constrained sellers. In the context of oDesk, Horton (2016) shows that sellers' activity on the platform is irregular, suggesting outside employment plays an important part. For AirBnB, the seller willingness to rent is based on the possibility to lend the property for a longer stay, or at a better price later on.

I formalize the above described interaction in a two stage search and matching model. The *First stage* is the buyer *search*: he observes the available sellers and their characteristics. The buyer decides which sellers to search in a directed, simultaneous (fixed sample) manner, paying a constant search cost for each searched seller. Seller *reservation prices* are private information as they are based on his private seller-specific demand from outside the platform. In the *Second stage* of the model, the searched sellers and the buyer exchange offers and possibly there is a stable *match*. The equilibrium concept for matching games is *stability* (Roth and Sotomayor (1992)), and the conditions it implies allow me to construct and estimate the *Second stage* parametrically. To identify and estimate the primitives of the *First stage* - the bounds on the seller search cost - I assume seller-optimal surplus allocation of the match.

The structural approach bypasses the need for transaction price or bid (offer) data, which may not always be available to researchers. It is useful to construct counterfactual settings that quantify the effect of information frictions. I step back and evaluate the contribution of the marketplace in resolving two considerable information frictions: eliciting seller availability and maintaining seller profiles and a reputation system. As a full information benchmark, I simulate a scenario where the buyers observe the seller reservation price, which is the last informational asymmetry remaining in the current set-up of MaistorPlus. The economic outcomes I compare are the search probability and intensity (number of sellers searched) and the match probability and surplus.[2]

My paper is positioned in the small but ambitious literature studying information frictions in online service markets. Cullen and Farronato (2015), Fradkin (2017) and Horton (2017) use proprietary data from TaskRabbit, AirBnB and oDesk, respectively, to document the existence of said frictions and their significant impact on the functioning of the marketplace. Horton (2017) studies seller rejection for buyers sending a single (early) invitation, representing about 30 percent of the data as many buyers contact multiple sellers. He does not explicitly model the process of search or matching, nor the source of information frictions. Fradkin (2017) defines rejection due to host preferences as the main information friction in this set-up. He estimates two reduced form models: a discrete choice model to predict the searcher's decision and a model of host screening and rejection. Cullen and Farronato (2015) construct a macro-style structural model of aggregate matching with frictions in the Diamond-Mortensen-Pissarides framework with a focus on demand and supply elasticity. Both Horton (2017) and Fradkin (2017) quantify welfare effects related to the information friction in their respective settings. Horton (2017) does a "back-of-the-envelope" evaluation of the semi-successful capacity signaling feature introduced by oDesk. Fradkin (2017) evaluates a pre-AirBnB scenario and an improved AirBnB scenario, where ranking incorporates the likelihood of acceptance.

In the context of this work, my first contribution is to construct structural model at the individual interaction level, which incorporates buyer and seller heterogeneity and explicitly models the search and match formation processes. The literature up to now has either approached this issue by reduced form analysis of the buyer or seller individual choices (Horton (2017), Fradkin (2017)), or by aggregate matching with no heterogeneity and no explicit information frictions (Cullen and Farronato (2014)). Apart from the aggregate matching model of Cullen and Farronato (2014), the existing studies do not consider, let alone model, how the match surplus is split and the transaction price is formed. Nor do they distinguish between the two sources of information friction related to capacity: seller availability and reservation price. Lastly, the structural model allows me to construct counterfactuals which quantify the economic impact of the frictions. The results can inform online service market design, and the model has the potential to be generalized and applied to other online service markets with similar features.

[2] My counterfactuals allow me to focus on the economic impact of information frictions without burdening the analysis with two-sided market issues (Rochet and Tirole (2006)). A more active role of the platform would make strategic the buyer and seller decision to subscribe, the buyer decision to post a job, the seller decision to indicate availability, and the joint decision to form a match. However, the potential for the platform to further improve efficiency by taking a more active role warrants future research.

This paper is also related to the intersecting literatures on search and matching, and in particular to the structural models which combine them. Chade, Eeckhout and Smith (2017) review the foundations and recent advances in with a focus on empirics. Key differences with this literature stem from the nature of the economic environment generating my data, described in *Section 1.3*. My paper is also related to the consumer search literature. Some recent contributions are De los Santos, Hortacsu and Wildenbeest (2012) on testing search models using online search patterns and Moraga-Gonzalez, Sandor and Wildenbeest (2015) on methods to estimate discrete-choice demand with search on the match value of utility. The theoretical search model that I develop has the flavor of the stochastic portfolio problem of Chade and Smith (2010), with key differences made clear in **Section 1.3**. The consumer search literature does not consider seller capacity constraints, to the best of my knowledge.

The paper has the following structure. In **Section 1.2**, I describe and solve for the equilibrium of the two-stage search and matching game. **Section 1.3** demonstrates the identification of the primitives of interest. **Section 1.4** details the steps I take to estimate the model. **Section 1.5** contains the results of the estimation, and **Section 1.6** contains the counterfactual analysis. I conclude in **Section 1.7**.

1.2 Data

MaistorPlus are based in Bulgaria and started operating in 2012.[3] The marketplace connects home owners, referred to as buyers, to home service professionals, referred to as sellers. Similar business model exists in many other countries: Thumbtack, Angie's List in the US; RatedPeople, MyBuilder, and Home Jane in the UK; MyHammer, Blau Arbeit, and Haus Helden in Germany; and Travaux in France.

The marketplace works in the following way. Buyers post jobs for free, and sellers pay a monthly subscription fee to have access to them.[4] The marketplace is additionally financed by advertising. The platform maintains verified seller profiles and a public reputation system. All sellers are notified when a job is posted, and those available can message the buyer. The sellers do not commit to a price before being contacted or searched; the

[3] The website of MaistorPlus is: http://maistorplus.com/

[4] The fee design affects only the extensive margin (how many agents have subscribed, or participation) and not the intensive margin (the activity of the subscribed agents, or usage). It allows me to consider the seller and buyer activity (sellers sending messages of availability, buyers searching the sellers) as decisions not affected by strategic considerations related to the fee structure.

buyer decides which of the available sellers to search based on their profile and reputation. Searching entails discussions of the job over the phone and visits. Note that I do not observe offers, transaction prices, and the process by which the price is set. I only observe which sellers are available, which sellers are contacted by the buyer, and which seller is hired.

1.2.1 Jobs, activity, seller profiles

In the complete sample, I have 4,167 jobs posted on the platform between January 2013 and June 2015, or about 231 jobs posted each month. The total suggested budget by the buyers is 12 million euro, with the median job budget of 250 euro. On average, 80.6 sellers are notified when a job is posted. The average activity for each job is the following: there are 5.4 available sellers and the buyer searches 1.4 of them. The buyer does not search in 38.7 perfent of the time (1,611 jobs). Conditional on searching, the match probability increases from 0.27 percent to 0.44 percent. This information is available in *Table 1.1*.

Table 1.1: Activity at the level of the job

	Observations	Mean	St. dev.	Min.	Max.
Full sample of jobs					
Notified sellers	4,167	80.2	46.4	1	397
Available sellers	4,167	5.37	5.5	1	61
Pr(Search)	4,167	0.61	0.5	0	1
Searched sellers	4,167	1.29	1.6	0	10
Pr(Match)	4,167	0.27	0.4	0	1
Jobs with search					
Notified sellers	2,556	81.2	46.8	1	397
Available sellers	2,556	5.54	5.5	1	61
Pr(Search)	2,556	1.00	0.0	1	1
Searched sellers	2,556	2.09	1.6	1	10
Pr(Match)	2,556	0.44	0.5	0	1

These descriptive statistics are in line with search and matching on other service provision platforms: buyers search multiple service sellers and the probability to be rejected is considerable. For example, the data of Fradkin (2017) for activity on AirBnB shows that buyers who search view 5.5 percent of all available listings (73 listings), and search 2.4 listings on average. Overall, 42 percent of all searches are rejected. Similalry, buyers on oDesk invite 2 sellers on average to apply to their job, the invitation acceptance rate is 55 percent, and the probability that the job opening is eventually filled is 55 percent (Horton

(2017)). Cullen and Ferronato (2014) report similar results for TaskRabbit: of all posted tasks, 78 percent receive at least one offer, the average being 2.8 offers per job; 49 percent of tasks result in a match.

There are a total of 817 active sellers in the sample. The seller *profile* contains the following information: categories of activity, profile description, references from previous buyers and pictures from past jobs. These characteristics are fixed over time; summary statistics are provided in *Table 1.2*. Characteristics describing the seller *experience* on the marketplace and the seller's *message* indicating availability are measured at the moment the job is posted and for those sellers who are available. The experience variables are the seller tenure on the marketplace (in months), the total times the seller was hired, and the percent positive reviews he has received.[5] The message-related variables are message length (measured in characters) and the time of the message (measured in hours since the job was posted on the online marketplace). Summary statistics for these variables are also presented in *Table 1.2* for the total of 22,379 times any seller was available in the sample.

Table 1.2: Seller characteristics at time of availability for any job in the sample.

	Observations	Mean	St. dev.	Min.	Max.
Profile (fixed)					
Active categories	817	4.80	4.7	1	21
References	817	0.14	0.6	0	3
Profile description length (chars.)	817	548.77	497.6	0	3,645
Pictures	817	10.76	24.5	0	490
Experience (at availability)					
Percent positive reviews	22,379	0.33	0.4	0	1
Marketplace tenure (months)	22,379	7.48	7.1	0	37
Total times hired	22,379	3.78	7.1	0	46
Message (at availability)					
Message length	22,379	240.47	290.8	0	11,932
Message time (hours)	22,379	3.58	14.7	0	577

Lastly, the buyer provides a textual description, indicates the job category (one of 21 categories such as carpentry, roof repairs, construction, etc), expected start date (one of 8 categories), and proposed budget (one of 12 categories). The frequencies of the different characteristics (job category, start and budget) can be found the **Appendix**.

[5] The percent positive reviews is a reputation measure is defined as in Tadelis (2016). I count missing reviews of completed jobs as non-positive reviews, as it has been demonstrated that buyers prefer not to leave any feedback when they are not fully satisfied (Tadelis (2016)).

1.2.2 Reduced form analysis

The reduced form analysis presented in this section provides grounding for the theoretical model. First, I consider outcome variables at the different stages of the process: how many sellers were available (regression A), how many sellers were searched by the buyer (regression S), and the probability of a match (regression M) at the individual job level. As explanatory variables, I am interested in the following exogenous covariates: job characteristics (expected budget and start), year fixed effects and demand factors. The *High demand season* dummy equal to 1 for the period June-November.[6] I also control for the previous stage's activity. The first two processes are modeled by a Poisson distribution, and the second - a probit.[7]

The results are presented in *Table 1.3*. Regression A indicates that: MaistorPlus' ability to bring sellers on board is important for the buyers, as more notified sellers leads to more available sellers. The *High demand season* lowers the number of available sellers, as can be expected when sellers are capacity constrained. In regression B, which describes buyer search, the demand conditions do not appear to be significant for the buyer search process. The last regression, M, demonstrates that the probability of a match depends on the demand conditions even after controlling for the number of available and searched sellers. As the hiring decision is jointly made, this result suggests that the value of the job to the available seller depends on the state of demand. Because capacity is fixed, the sellers must consider the value of giving up outside options, which are undoubtedly more when demand is high, before agreeing on a match. Job characteristics such as the job start and budget also have significant predictive power in all thee regressions.

In *Table 1.4*, I present evidence that buyer search is directional: the sellers are contacted based on how they compare to each other. I model the probability that an available seller is contacted by the buyer using a logit regression. As independent variables, I have seller and available competitor characteristics.[8] In the first model, $S1$, all seller characteristics have a significant impact on the probability that a seller is searched, except for marketplace tenure. The estimated coefficients have the expected signs: fuller profiles, more experience and faster and more detailed messages make sellers more attractive. Includ-

---

6 Significantly more jobs are posted on the platform during that period, due to construction and other categories being outside work, as well as homeowners preparing for the fall and winter months. A priori, there is no reason to believe that demand fluctuations on the marketplace are different from those outside of the marketplace.

7 The results using the subsample of jobs for which the buyer searches at least one seller are consistent.

8 Although not presented in the table, I control for job characteristics, number of notified, available and searched sellers, job category fixed effects, the high demand season and year fixed effects.

Table 1.3: Job-level activity: available sellers, search and probability of a match

	A	S	M
Model	Poisson	Poisson	Probit
N. notified sellers	0.434***	0.072*	0.038
	(0.026)	(0.037)	(0.046)
N. available sellers		0.568***	-0.448***
		(0.036)	(0.053)
N. searched sellers			1.181***
			(0.042)
High demand season	-0.451***	0.061	-0.261***
	(0.026)	(0.046)	(0.059)
Job start	0.177***	(0.028)	-0.330***
	(0.029)	(0.043)	(0.056)
Job budget	0.113***	0.090**	-0.169***
	(0.021)	(0.038)	(0.051)
Constant	0.765***	-1.955***	-0.151
	(0.139)	(0.292)	(0.327)
Year fixed effects	Yes***	Yes**	Yes***
Job category fixed effects	Yes***	Yes***	Yes***
Pseudo R^2	0.33	0.09	0.19
Observations	4,167	4,167	4,167

Significant at: $p < 0.1$: *; $p < 0.05$: **; $p < 0.01$: ***. Continuous variables are transformed by taking the natural logarithm. More detailed results are available upon request.

ing the corresponding average competitor characteristics in $S2$ shows that any seller is less likely to be searched the more attractive his competitors are. Finally, in $S3$ I include seller fixed effects to control for unobserved seller features. The fit of the regression, as measured by the pseudo-R^2, does not appear to improve much, which suggests that the buyers do not have private information on the sellers that is different from what is already included in the seller profile, experience and message.[9]

Lastly, the hiring decision in service markets is joint, which is an important feature that sets these markets apart from goods markets. It is not always the case that the buyer is able to match with, or hire, the most ex-ante attractive seller, even if that seller is available. I denote the predicted probability of search from regression $S3$ seller *attractiveness*. I rank sellers by this probability, with the most attractive seller is ranked as 1, and less attractive sellers have lower ranks $(2, 3$ etc). The data shows that on average the buyers match with sellers of attractiveness 2.3, possibly because the highest ranking seller asked for too much.

1.3 Model

Consider a simple model of one-to-one matching with transfers (prices), where on one side there is a single buyer indexed by i and on the other side there are multiple sellers indexed by j. Let N_i be the number and $\{N_i\}$ be the set of sellers *available* for job i. Similarly, let n_i be the number and $\{n_i\}$ be the set of *searched* sellers such that $\{n_i\} \subset \{N_i\}$ and $n_i \leq N_i$.

When a seller is not available for a given job, I consider that their capacity is fully utilized and it is physically impossible for them to be hired at the given moment. Whenever they are available, their willingness to take on the job depends on whether they will be compensated for their *reservation price*. As seller capacity is fixed, taking on one project means giving up on another. The reservation price represents the value of all work opportunities the seller would give up to take on a given job. For the sellers on MaistorPlus, the reservation price is based on demand coming through traditional channels such as client referrals, repeated clients, and yellow pages. This is why there are only 5.4 sellers are available on average for any job, out of the 80.2 sellers that are notified.

[9] Any private information on the sellers is only revealed after the job is completed, hence it does not affect the decision of the buyer to contact or hire the seller. This information can be then transmitted to future buyers though the reputation system. Repeated hiring, which may be affected by the buyer's private information on seller type, happen outside of the platform.

Table 1.4: Probability of searching any available seller.

	$S1$	$S2$	$S3$
Profile			
Seller active categories	-0.054**	-0.050*	
	(0.024)	(0.029)	
Competitor active categories		0.005	-0.027
		(0.054)	(0.053)
Seller references	0.123***	0.136***	
	(0.039)	(0.044)	
Competitor references		-0.045	0.028
		(0.076)	(0.076)
Seller profile decription length	0.020***	0.025***	
	(0.007)	(0.007)	
Competitor profile description length		-0.038**	-0.014
		(0.019)	(0.019)
Seller pictures	0.031***	0.035***	
	(0.010)	(0.011)	
Competitor pictures		-0.018	-0.017
		(0.021)	(0.021)
Experience			
Seller percent positive reviews	0.298***	0.381***	0.441***
	(0.057)	(0.068)	(0.088)
Competitor percent positive reviews		-0.431***	-0.381***
		(0.115)	(0.119)
Seller total times hired	0.124***	0.162***	0.033
	(0.019)	(0.022)	(0.032)
Competitor total times hired		-0.090***	-0.086**
		(0.033)	(0.034)
Seller marketplace tenure	-0.02	-0.023	0.032
	(0.018)	(0.020)	(0.026)
Competitor marketplace tenure		-0.016	-0.043
		(0.038)	(0.039)
Message			
Seller message length	0.022**	0.054***	0.046***
	(0.011)	(0.014)	(0.015)
Competitor message length		-0.132***	-0.141***
		(0.022)	(0.023)
Seller message time	-0.298***	-0.381***	-0.376***
	(0.017)	(0.020)	(0.021)
Competitor message time		0.181***	0.185***
		(0.020)	(0.021)
Seller fixed effects			Yes***
Pseudo R^2	0.42	0.43	0.45
Observations	22,379	22,379	22,379

Significant at: $p < 0.1$: *; $p < 0.05$: **; $p < 0.01$: ***. Continuous variables are transformed by taking the natural logarithm. More detailed results are available upon request.

Yet, the sellers are not equally in demand: sellers with more experience and better reputation are more attractive to buyers, as demonstrated in *Table 1.4*. Demand conditions also matter: the sellers experience higher demand in the summer and fall, which allows them to be more picky about jobs. Let $\tilde{r}_{ij}$ denote *average reservation price*: given the observed seller and job characteristics (including the period at which the job was posted), what is the expected minimum compensation. As this information is visible to the buyer, so is $\tilde{r}_{ij}$. On the other hand, the exact incidence, profitability and size of demand faced by any seller at any point in time is random and privately observed. This random demand fluctuation determines the seller *private reservation price* r_{ij}, not observed by the marketplace nor the buyer.

Once matched together, the buyer-seller pair generates *match output* $\tilde{f}_{ij}$. The match output is a function of the job characteristics and seller characteristics. It is therefore also observed by both the buyer and the sellers. Note that certain seller characteristics may affect both the match output $\tilde{f}_{ij}$ and the average reservation price $\tilde{r}_{ij}$, such as for example seller experience. A more experienced seller will both do a better job, and require a higher compensation. However, with no data on transaction or offer prices, the econometric model does not identify the effect of seller experience on $\tilde{f}_{ij}$ and $\tilde{r}_{ij}$ separately. I proceed by grouping them in the observed (by the buyer and econometrician) *net match output*: $f_{ij} = \tilde{f}_{ij} - \tilde{r}_{ij}$. Finally, the surplus generated by the match is the net match output less the private seller reservation price, $s_{ij} = f_{ij} - r_{ij}$. As the buyer does not observe r_{ij}, he does not observe s_{ij} but only f_{ij}.

I argue that the realizations of the private seller reservation prices r_{ij} are independent across the different sellers available for the same job, and for the same seller across different jobs.

Assumption 1. $r_{ij} \perp r_{ik}$

Assumption 2. $r_{ij} \perp r_{i'j}$

By **Assumption 1**, the r_{ij} are independently distributed across different sellers j and k within the same job i. The sellers are small relative to each other and relative to the market, which implies that they face idiosyncratic demand shocks. They are not in direct competition on or off the platform except for the given job i. This assumption also rules out a "common value" aspect to r_{ij}, in the language of auction economics.[10]

[10] A common value a set up would entail that sellers do not know the true, common reservation price and they estimate it based on a private signal. I do not find this to be a realistic representation of the economic environment: the sellers have access to the same information, hence there is no private signal;

By **Assumption 2**, the seller j reservation prices r_{ij} and $r_{i'j}$ across jobs i and i' are independent. As the net output f_{ij} already contains the average seller reservation price $\tilde{r}_{ij}$, the remaining private component r_{ij} is random at the moment of posting the job i: the specific realization of idiosyncratic seller outside demand net of the average. The two assumptions on the nature of r_{ij} allow me to treat the realizations of r_{ij} across the different sellers available within a market as independent, and the markets i as independent.

Denote in the upper case R the random reservation price, whose realizations are the r_{ij}. In addition to independence, I assume this variable is identically distributed across sellers and jobs:

Assumption 3. *R is identically distributed across sellers and jobs with cumulative distribution function (CDF) $G_R(.)$.*

The contribution of observed job and seller characteristics to the match output is f_{ij}, therefore the random reservation price R is mean independent from this. **Assumption 3** is stronger, implying all moments of the distribution do not depend on the observed characteristics, and therefore on f_{ij}: $G_R(.|f_{ij}) = G_R(.)$. With no information on the realizations of R, the match surplus is also random: $S_{ij} = f_{ij} - R$. Its distribution is:

$$\begin{aligned} G_{ij}(s) = Pr(S_{ij} \leq s|f_{ij}) = Pr(f_{ij} - R \leq s) = \\ = Pr(f_{ij} - s \leq R) = 1 - Pr(R \leq f_{ij} - s) = 1 - G_R(f_{ij} - s) \end{aligned}$$

The search and matching model is set-up in the following way. In *Zero stage*, which is non-strategic, *Nature* draws the buyer i and the set of available sellers $\{N_i\}$. The buyer and the available sellers observe the job and seller characteristics, therefore the f_{ij}'s for all sellers in $\{N_i\}$. The available sellers have private information about their reservation prices r_{ij}. In the *First stage*, the buyer searches the optimal set of sellers $\{n_i\}$ in a simultaneous and directed manner, paying a positive search cost c_i for each seller. In the *Second stage*, the sellers in $\{n_i\}$ compete by making utility offers to the buyer. A match between a seller and the buyer may be formed if it is stable. *Figure 1.1* displays the sequence. The game is solved by backward induction.

all material costs are borne by the buyer , hence there is no uncertainty about a common cost. Any unexpected developments during the completion of the job are re-negotiated by the parties. Thus, the predominant concern with taking on a job comes from their capacity limitations and private idiosyncratic demand.

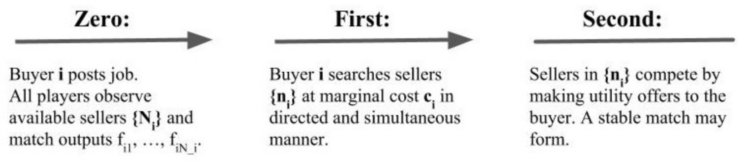

Figure 1.1: Stages of the search and matching game.

1.3.1 Second stage

In the *Second stage*, the searched sellers $\{n_i\}$ compete to form a match with the buyer. As the sellers are vertically differentiated, the buyer is willing to pay more for a seller with more experience or a deeper specialization. I consider the sellers making offers to the buyer in utility space. For a given match output, there is a one-to-one correspondence between the utility offer and the transaction price, which is not observed in the data. The equilibrium concept is *match stability*: the match must satisfy individual rationality for each party and must assign the buyer to the seller who generates the highest match surplus (Roth and Sotomayor (1992)).

The surplus created by a stable match may be split between buyer and seller through the price in infinitely many ways. MaistorPlus, however, does not record prices or provide an explicit framework for how they are set.[11] The surplus split affects the ex-ante expected utility of the buyer in the *First stage* of the model, therefore the optimal search set $\{n_i\}$ and the estimation of the search costs.

I consider the *seller-optimal surplus allocation*: the winning seller has offered the buyer a utility that makes the second-highest surplus seller indifferent. The mechanism by which this can be achieved is the reverse English auction, or open ascending auction in utility (Roth and Sotomayor (1992)).[12] As a mechanisms, the English auction has two important features. Firstly, it guarantees that the assignment, or the match, is efficient, i.e. the seller who generates the highest match surplus will be matched with the buyer. This is true with

[11] The flexibility of interaction environment is necessary for the home service industry, where sellers are unwilling to commit to a price before speaking with the buyer, and where both sides of the interaction typically do not wish to reveal the transaction value to the platform. Other service platforms may employ a different degree of structure on how the partners negotiate and agree to split the surplus. For example, on AirBnB, service sellers typically set a daily rate, which is recommended by the platform algorithm. Interactions on Thumbtack and TaskRabbit (formerly oDesk and Elance) are structured as an open auction. oDesk was previously structured in a similar way as MaistorPlus, but offers are exchanged via the messaging service of the platform.

[12] In the private values framework, the English auction format is strategically equivalent to the Second price sealed bid auction format (Krishna (2009)). The English auction appears to be more suitable in the particular economic environment.

ex-ante asymmetric sellers, as in our game. Secondly, the English auction equilibrium is in dominant strategies and I do not have to make assumptions about what the sellers know or learn about each other.

For the moment, consider a single buyer and drop the subscript i. The buyer utility of transacting with seller j is $u_j = f_j - p$ and the seller j utility is $v_j = p - r_j$.[13] The strategies of the sellers are defined as utility offers they will make to the buyer as in Laffont-Tirole (1993). Let the searched sellers $\{n\}$ be ordered by match surplus: $s_1 = f_1 - r_1 \geq ... \geq s_n = f_n - r_n$. The match stability requires two things. First, the buyer forms a match with the seller who generates the highest match surplus, s_1. Second, that the individual rationality constraints are satisfied. For the buyer it is $IR_j^b : u_j = f_j - p \geq 0$ when transacting with seller j, and for seller j it is $IR_j^s : v_j = p - r_j \geq 0$. For this to be possible, it must be that the generated surplus is at least positive, $u_1 + v_1 = f_1 - r_1 = s_1 >= 0$.

The English auction works in the following way. The auctioneer (the buyer) starts from an utility offer of zero and raises it. The sellers remain in the auction while they agree to the offer, and the game ends when only one seller remains. The transaction utility is that at which the second-last seller drops out of the game. The players' weekly dominant strategies are to remain in the game up to the point they are indifferent (Vickrey (1961)). In other words, player j generating match surplus $s_j = f_j - r_j$ remains up to the utility offer $u_j = s_j$ and drops out afterwards.

The game can be summarized by the following three cases:[14]

1. $0 < s_2 \leq s_1$: 1 wins, gives the buyer utility $u_1 = s_2 > 0$
   - transaction price is determined by $f_1 - p = f_2 - r_2$ so $p = f_1 - f_2 + r_2$
   - $IR_1^s : v_1 = p - r_1 \geq 0$ is satisfied because $s_1 = f_1 - r_1 \geq s_2 = f_2 - r_2$
2. $s_2 \leq 0 \leq s_1$: 1 wins, gives the buyer utility $u_1 = 0$
   - transaction price is determined by $f_1 - p = 0$ so $p = f_1$
   - $IR_1^s : v_1 = p - r_1 \geq 0$ is satisfied because $s_1 = f_1 - r_1 \geq 0$
3. $s_2 \leq s_1 < 0$: no transaction[15]

[13] Here the observed seller reservation price $\tilde{r}_j$ is also zero. This does not change the equilibrium but makes the exposition clearer.

[14] When $n_i = 1$, there is only the last two cases.

[15] Search without hiring is very common in the data, which raises the question why sellers with $s_{ij} \leq 0$

The transaction takes place only in the first two cases, and in only the first case the buyer gets positive utility.

1.3.2 First stage

Given the expected utility of the *Second stage*, in the *First stage* the buyer decides which sellers to search at a positive and constant marginal cost c. I model the search process as directed and simultaneous for the following reasons. On MaistorPlus, the candidates indicate their availability soon after the job is posted and the buyer can direct his search by their observed characteristics. This is different from labour markets, where Chade, Eeckhout and Smith (2017) consider sequential random search because of the much longer horizon over which candidates arrive, and the fact that the identity of the candidates is not known ex-ante. My setting favors simultaneous search because the transaction price is determined in the *Second stage* of the interaction. Thus, the buyer has an incentive to bring the sellers together so that they can compete in making offers. Simultaneous search models are favored in settings of consumer search, as in the online online book industry studied by De los Santos, Hortacsu, Wildenbeest (2012). The search model I construct is close to the optimal portfolio model of Chade and Smith (2010). The difference is that the probability of hiring and expected surplus of any seller added to the search set in my model is a function of all other sellers already added to the search set through the expected highest surplus, while in their model the success rate and ex-post payoff of the any of the sellers are fixed and independent.

The buyer must choose among N differentiated stochastic options, which is a combinatorial optimization problem. I show that the buyer searches sellers in order of decreasing observed match surplus f_j and up to the point where the marginal benefit of an additional seller is less than c.[16] Let the random variable $S^{2:\{L\}}$ be the second-highest realization of

indicate availability. It is possible that a common component of $\tilde{f}_j$ is only revealed in the search stage by the buyer. For example, details about the project which affect all sellers in the same way could be omitted from the buyer description. As long as the sellers have no private signal about it (i.e. the "common values" framework, which I exclude by assuming no correlation of the r_{ij}'s across j) this does not cause selection on the part of the sellers.

[16] The cost of the first search is normalized to zero, similarly to Hortacsu and Syverson (2004) and Dubois and Perrone (2015). When the buyer contacts only one seller, the expected utility of the buyer is zero. Ex-ante no buyer would be willing to contact only a single seller at a positive search cost, but many such observations exist in the data. The assumption of zero search cost for the first search makes buyers indifferent between searching 1 and 0 sellers, which is consistent with the data. One could think of the decision to search at least one professional as incorporated into the decision to post the job on the marketplace.

match surplus from a random set $\{L\}$ of searched sellers. The cumulative distribution of this variable, which is an order statistic, is:

$$G^{S^{2:\{L\}}}(s) = Pr(S^2 \leq s|\{L\}) = \sum_{j=1}^{L}[1 - G_j(s)]\prod_{k\neq j} G_k(s) + \prod_{j=1}^{L} G_j(s)$$

The buyer anticipates the three potential outcomes of the *Second stage*. Only in the first case he receives positive expected utility, which is equal to $E[U] = E[S^{2:\{L\}}|S^{2:\{L\}} > 0]Pr(S^{2:\{L\}} > 0)$. It is derived as:

$$E[S^{2:\{L\}}|S^{2:\{L\}} \geq 0]Pr(S^{2:\{L\}} \geq 0) =$$
$$= \frac{\int_0^\infty s\frac{d}{ds}G^{S^{2:\{L\}}}(s)ds}{1 - G^{S^{2:\{L\}}}(0)}.(1 - G^{S^{2:\{L\}}}(0)) = \int_0^\infty s\frac{d}{ds}G^{S^{2:\{L\}}}(s)ds$$

Now I compare the distributions of the second highest match surplus from the sets $\{L\} + \{l\}$ and $\{L\}+\{l'\}$, where $f_l > f_{l'}$. If the difference between $G^{S^{2:\{L\}+\{l\}}}(s)$ and $G^{S^{2:\{L\}+\{l'\}}}(s)$ is negative, by the property of first order stochastic dominance the random variable distributed by $G^{S^{2:\{L\}+\{l\}}}(s)$ has a higher expected value.

$$G^{S^{2:\{L\}+\{l\}}}(s) - G^{S^{2:\{L\}+\{l'\}}}(s) = [G_R(f_{l'} - s) - G_R(f_l - s)]\Big(\sum_{j=1}^{L}[1 - G_j(s)]\prod_{k\neq j} G_k(s)\Big)$$

Since $f_l > f_{l'}$, I know $G_R(f_l - s) > G_R(f_{l'} - s)$ because $G_R(.)$ is an increasing function. This makes the first part of the expression negative. The distribution of the second highest match surplus from $\{L\} + \{l\}$ first order stochastic dominates that from $\{L\} + \{l'\}$. Adding a seller with higher f_j to any set $\{L\}$ is optimal as it leads to a higher expected utility for the buyer. By induction, this holds for sets of any size and composition.

The second order condition - that the marginal benefit of each seller added to the search set decreases - cannot be proven analytically for a general function $G_R(.)$. Instead, in the **Appendix** I demonstrate that it holds for the specific functional form of $G_R(.)$ given in

Assumption 4 in the next section and for the values of f observed in the data.

1.4 Identification and estimation

In this section, I describe what is observed in the data, the model, and how I am able to identify the primitives.

1.4.1 Observables, primitives and assumptions

In the data, I observe buyers $i \in \{1, ..., M\}$ posting jobs on the marketplace. For each observation i and the corresponding available sellers $\{N_i\}$, both I (the econometrician) and the buyer observe the characteristics of the job and seller pair. The job characteristics, constant across sellers j, are labeled Z_i. These are the start date, category, proposed budget for the job and date fixed effects. The seller characteristics for each job-seller pair ij are labeled X_{ij}. These are the seller profile, experience and message variables presented in *Table 1.2*.

The match surplus is modeled as the sum of the observed net output f_{ij} and unobserved reservation price r_{ij} already introduced in **Section 1.3**: $s_{ij} = f_{ij} - r_{ij}$. The net output is made up of two components: the match output $\tilde{f}_{ij} = Z_i'\gamma_1 + X_{ij}'\beta_1$ and the average seller reservation price $\tilde{r}_{ij} = Z_i'\gamma_2 + X_{ij}'\beta_2$, where some covariates may enter both. Without transaction and offer prices, I am not able to identify how the covariates affect these two components separately, but only their net effect: $f_{ij} = \tilde{f}_{ij} - \tilde{r}_{ij} = Z_i'(\gamma_1 - \gamma_2) + X_{ij}'(\beta_1 - \beta_2) = Z_i'\gamma + X_{ij}'\beta$. The match surplus is therefore parametrized in the following way:

$$s_{ij} = f_{ij} - r_{ij} = Z_i'\gamma + X_{ij}'\beta - r_{ij}$$

The primitives of interest from the *Second stage* estimation are the coefficients γ of Z_i and β of X_{ij}. The primitives of interest from the *First stage* estimation are the search cost bounds $[\underline{c}_i, \overline{c}]_i$.

Lastly, consider the following parametric assumption:

Assumption 4. $-R \sim$ *Type 1 Extreme value*

The ex-ante random negative reservation price $-R$ follows an Type 1 Extreme Value distribution, with probability function $g_{-R}(a) = -e^{-a}e^{-e^{-a}}$ and cumulative density $G_{-R}(a) = -e^{-e^{-a}}$.[17] As in other discrete choice models, the scale of the coefficients, and therefore of the surplus, is not identified and must be set by the econometrician. The distributional assumption here sets the level of the variance of the error term to $\pi^2/6$, which automatically sets the scale.

1.4.2 Second stage

The sample of M jobs is comprised of independent observations, indexed by i. Each job can have one of the following outcomes: either no one is hired, or one of the searched sellers is hired. Let the probability of the first event be $Pr(Y_i = 0)$, and the probability that seller $j \in \{n_i\}$ is hired be $Pr(Y_i = j)$. Let d_{ij} be an indicator equal to 1 when buyer i hires seller j, and zero otherwise. The indicator for the event that no one is hired is $d_{i0} = 1 - \sum_{j=1}^{j=n_i} d_{ij}$. The probability for any outcome for job i is:

$$L(\gamma, \beta | X_i, Z_i, \{n_i\}) = Pr(Y_i = 0)^{1-\sum_{j=1}^{j=n_i} d_{ij}} \prod_{j=1}^{j=n_i} Pr(Y_i = j)^{d_{ij}}$$

The derivation of the Maximum Likelihood can be found in the **Appendix**. The Likelihood function for the whole sample is:

$$L(\gamma, \beta | X, Z, \{n\}) = \prod_{i=1}^{i=M} L(\gamma, \beta | X_i, Z_i, \{n_i\}) =$$
$$= \prod_{i=1}^{i=M} \left(\left(e^{-e^{f_{ij}}} \right)^{1-\sum_{j=1}^{j=n_i} d_{ij}} \prod_{j=1}^{j=n_i} \left(\frac{1 - e^{-\sum_j e^{f_{ij}}}}{1 + e^{-f_{ij}} \sum_{k \neq j} e^{f_{ik}}} \right)^{d_{ij}} \right)$$

Although this form of the Likelihood is reminiscent of the standard multinomial choice model, in fact there is an important difference rooted in the equilibrium of the matching game. In standard multinomial choice, the buyer chooses the best seller among a number of options. The identification strategy is based on utility differences between the options, thus it identifies the effects of variables that are different across the options. In my specification, these are the variables X_{ij} their corresponding coefficients β.

[17] Working with the distribution of the negative reservation price results in a cleaner maximum likelihood derivation, without affecting the rest of the model.

In the matching game, however, match stability requires that whenever a seller is hired, the match surplus is positive. This introduces an additional condition on the overall level of match surplus, which identifies the effects of the variables specific to the buyer, or the job. In our model, these are variables Z_i and their corresponding coefficients γ.

1.4.3 First stage

The sellers are added to the search set $\{n_i\}$ in order of decreasing f_{ij}, and $\{n_i - 1\}$ and $\{n_i + 1\}$ denote the sets with one less and one more seller, respectively. The optimal search set $\{n_i\}$ of the buyer implies the following equilibrium inequalities on the search cost c_i:

$$\underline{c_i} = E[S^{2:\{n_i+1\}}] - E[S^{2:\{n_i\}}] \leq c_i \leq E[S^{2:\{n_i\}}] - E[S^{2:\{n_i-1\}}] = \overline{c_i}$$

For each job i, these bounds will be different because a different set of sellers $\{N_i\}$ is available. Using the results of the *First stage*, I can construct the seller fitted net output $\hat{f}_{ij} = Z_i'\hat{\gamma} + X_{ij}'\hat{\beta}$. The CDF of the match surplus S_{ij} for any job-seller pair ij is:

$$\hat{G}_{ij}(s) = 1 - G_R(\hat{f}_{ij} - s)$$

This allows me to construct $\hat{G}^{S^{2:\{n_i\}}}$. I construct the upper and lower bound on the search cost for each individual job as follows:[18]

$$\hat{\overline{c}}_i = \hat{E}[S^{2:\{n_i\}}] - \hat{E}[S^{2:\{n_i-1\}}]$$

$$\hat{\underline{c}}_i = \hat{E}[S^{2:\{n_i+1\}}] - \hat{E}[S^{2:\{n_i-1\}}]$$

[18] Some jobs are special cases, for example jobs where the set of available sellers $\{N_i\}$ is equal to the set of searched sellers $\{n_i\}$. They are discussed in more detail in the **Appendix**.

1.5 Results

This section contains the results of the estimation: the coefficients $\hat{\beta}$, on variables X_{ij}, and $\hat{\gamma}$, on variables Z_i, and the individual search cost bounds $[\hat{\underline{c}}_i, \hat{\bar{c}}_i]$.

1.5.1 Second stage

Table 1.5 presents the results from the MLE estimation of the *Second stage*. The variables X_{ij} are split into the seller fixed characteristics X_j that describe the seller's profile, and the seller characteristics measured at the time the job was posted X_{ij} that describe the seller experience and the message he sends to the buyer. I present four models $(M1)$ to $(M4)$, each with a different subset of covariates. Likelihood ratio tests indicate that model $(M4)$, with the full set of variables, fits the data significantly better than the models with less explanatory variables. Robustness checks with seller fixed effects are presented in the **Appendix**

The estimated coefficients represent the net effect of the given variable on the net output f, this way combining the effect on the output $\tilde{f}$ and the effect on the average seller reservation price $\tilde{r}$. This makes their interpretation less straightforward. For example, the results indicate a significant negative effect of the number of job categories in which the seller is active. This could be because the seller is less specialized, therefore his expertise does not contribute to the match output $\tilde{f}$ as much. Alternatively, it could be that a seller working in many categories has higher demand on average, and therefore higher $\tilde{r}$. For other variables, net coefficients that are not significant may be due to the two effects having opposing directions.

I find that the majority of seller fixed profile characteristics, such as the number of pictures, references, and the length of the description, are mostly not significant, while the experience-related variables all have significant effects. The first result could arise from seller characteristics having equal but oposite effects on the output and average seller reservation price. The second result suggests that the marketplace reputation system is a meaningful source of information and incentives. It is also interesting to note that the variables describing the seller message to the client, the message length and time, do not appear to have a very strong significance in the last regression. One reason could be that they are "cheap talk", and therefore not taken as a meaningful signal by the buyer. Indeed, inspection of the messages proves that the majority are formulaic and repeat information from the seller's profile.

Table 1.5: Results from the MLE estimations.

	M1	M2	M3	M4
Xj				
Seller pictures	-0.039	-0.047*	0.011	0.003
	(0.024)	(0.024)	(0.026)	(0.027)
Seller references	-0.286**	-0.109	-0.199	-0.116
	(0.101)	(0.104)	(0.105)	(0.109)
Seller profile description length	0.011	0.007	0.021	0.011
	(0.017)	(0.018)	(0.018)	(0.019)
Seller active categories	-0.081	-0.469***	-0.069	-0.264***
	(0.047)	(0.056)	(0.052)	(0.064)
Xij				
Seller percent positive reviews		0.879***		0.546***
		(0.140)		(0.144)
Seller marketplace tenure		-0.223***		-0.163***
		(0.046)		(0.049)
Seller total times hired		0.373***		0.188***
		(0.044)		(0.049)
Seller message length		-0.037		0.056*
		(0.024)		(0.025)
Seller message time		-0.176***		0.048
		(0.050)		(0.124)
Zi				
Job budget			-0.272***	-0.257***
			(0.068)	(0.068)
Job start			-0.350***	-0.342***
			(0.075)	(0.076)
High demand season			-0.383***	-0.397***
			(0.078)	(0.079)
Year fixed effects			Yes***	Yes***
Job category fixed effects			Yes***	Yes***
Constant	-1.031***	-0.155	0.305	0.718
	(0.149)	(0.194)	(0.501)	(0.528)
Likelihood	-2,592	-2,461	-2,200	-2,170
Observations	2,556	2,556	2,556	2,556

Significant at: $p < 0.1$: *; $p < 0.05$: **; $p < 0.01$: ***. Continuous variables are transformed by taking the natural logarithm. More detailed results are available upon request.

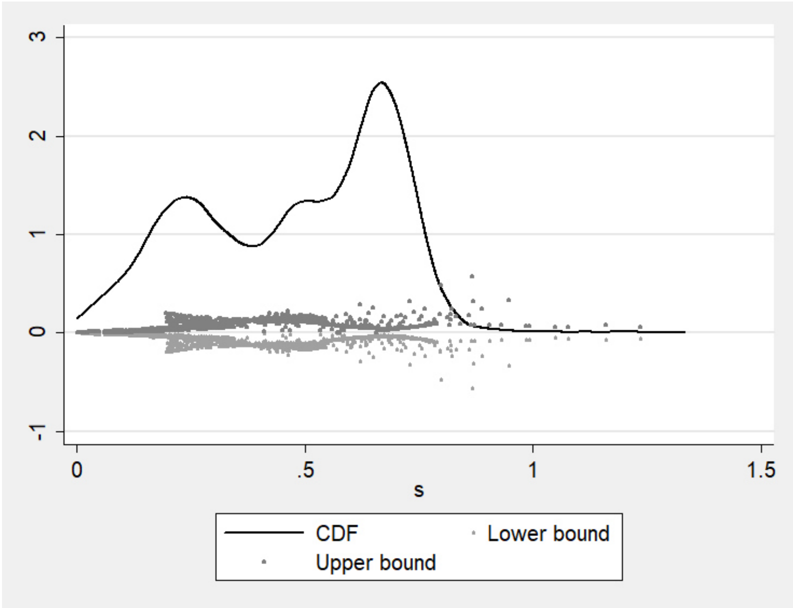

Figure 1.2: Distribution of the mid-point of the estimated search costs, average upper and lower bounds on the mid-point.

Lastly, a larger budget, a more delayed start and the job being posted in the high demand season decrease the match surplus. It is not surprising to see that periods of higher demand contribute to a lower surplus, because seller average reservation prices must also be higher due to the numerous job opportunities. It likely that bigger projects lead to delays and complications, and that projects further away in time are less certain, thus associated with a higher average reservation prices.

1.5.2 First stage

For each buyer, I estimate the individual search cost bounds $[\underline{c}_i, \overline{c}_i]$ using the *Second stage* results and the buyer's optimal search decision. *Figure 1.2* presents a kernel density estimate of the probability distribution of the estimated mid-point, $c_i^m = (\hat{\overline{c}}_i + \hat{\underline{c}}_i)/2$. It also displays the average $\hat{\overline{c}}_i$ and $\hat{\underline{c}}_i$ for observations of the search cost falling in each of the bins over which the density is estimated. The distribution has two peaks, and there is a positive mass of sellers with zero search costs.

1.6 Counterfactuals

MaistorPlus provides the buyers with the following benefits. Firstly, the platform has the contact information of multiple subscribing sellers. Secondly, it notifies these sellers simultaneously about the buyer's job and those who are available indicate that to the buyer, thus saving the buyer from searching unavailable sellers. Lastly, the website contains verified seller profile information and a reputation system.

I construct counterfactual scenarios to evaluate how the last two features affect the economic outcomes on the platform. The *Random* scenario is very information poor, and akin to the buyer picking out sellers at random from the phone book: there is no information on seller availability, no verified information on their expertise and previous buyer experiences. In the *Random available* scenario, the buyer now knows which sellers are available, but again any information they provide about their expertise and previous buyer experiences is cheap talk. I denote as the *Directed available* the current set-up of MaistorPlus.

Despite the marketplace alleviating most information frictions, the buyers are not able to observe costlessly the private reservation prices of the sellers. As a benchmark with no information asymmetry, I consider the *Frictionless* scenario. This can be achieved either by the platform finding a way to lower buyer search costs to zero, or the platform finding a way to make sellers truthfully reveal their reservation prices.

1.6.1 Method

In my counterfactuals I sample repeatedly the observations-markets as presented in *Zero stage* of the model: the jobs posted on the platform with the associated available sellers. This means sampling jointly the following: the job characteristics Z_i, the set of notified $\{\tilde{N}_i\}$ and available $\{N_i\}$ sellers, their respective characteristics $X_i = X_{i1}, ..., X_{in_i}$. For the optimal search decision of the buyer, I use the estimated midpoint of seller search cost $\hat{c}_i^m$ and I construct $\hat{f}_{ij}$ using the estimated coefficients $\hat{\beta}$ and $\hat{\gamma}$. The reservation prices are drawn randomly using the distribution in **Assumtion 4**.[19]

In all scenarios, the *First stage* follows the buyer optimal search strategy derived in **Sec-**

[19] Unfortunately, the data is not of sufficient size to estimate and sample from the joint distribution of the Z_i, $\tilde{N}_i$, $\{N_i\}$, X_i, and $(\overline{c}_i - \underline{c}_i)$. Therefore, counterfactuals represent the effect of information frictions on the a sample of markets I observe.

tion 1.3. In the *Random available* and *Random* scenarios, the sellers are ex-ante identical, and the distribution of $\hat{S}^{2:n_i}$ is derived appropriately. In the *Random* scenario, the buyer is also not sure if the searched seller is available or not. From the data, I find the probability a notified seller is available for the given job, which I assume the buyer knows. Using this, I modify the expected utility from adding the n^{th} seller by multiplying it by the probability that he is available.

In the *Second stage*, the buyer may form a match with a seller from the optimal search set. In the *Directional available* scenario, this is the n_i sellers with the highest $\hat{f}_{ij}$. In the random scenarios, the sellers are identical ex-ante and I construct the search set using draws from a uniform distribution on $U[0, 1]$. For the *Random available* scenario, I take the sellers with the top n_i random draws within the available set $\{N_i\}$. For the *Random* scenario, I take the sellers with the top n_i random draws within the notified set $\{\tilde{N}_i\}$. In all cases, a match is formed with the highest surplus *available* seller whenever $s_{i1} \geq 0$. In scenario *Frictionless*, the buyer observes the match surplus s_{ij} and his search set is *de facto* the set of available sellers $\{N_i\}$. He is matched with the highest surplus seller in $\{N_i\}$ whenever this seller generates a positive surplus: $s_{i1} \geq 0$.

1.6.2 Results

The results from the counterfactuals can be found in *Table 1.6*. The outcomes of *Directed available*, which represents how the platform is organized currently, are very similar to those presented in *Table 1.1*. This indicates that the estimated coefficients and search cost are a good fit. In particular, the uncoditional probability of search in the data is 0.61 percent, while in the simulated counterfactual this is also 0.61 percent; the probability of a match, conditional on search, is 0.44 percent in the data, and 0.46 in the simulation. Slightly more sellers are searched in the data: 1.21 versus 1.30 overall, and 2.09 versus 2.12, conditional on search.

Next, consider the *Random available* scenario. Compared to *Directed available*, note that the probability of search is much lower (0.22 versus 0.61), but the number of searches is higher (4.02 versus 2.12) conditional on search. The reason is the following. In *Random available* search, the sellers appear identical to the buyer in the *First stage*. The marginal benefit of search is lower initially than in the *Directed available* scenario, where the buyer can start his search with the better sellers. As a result, buyers with higher search costs who search in the *Directed available* scenario will not search in the *Random available* scenario, as the initial marginal benefit is not sufficient to offset the cost. The marginal benefit of search in the *Random available* scenario also decreases less quickly than in the *Directed available*, where the buyer exhausts the better sellers and is quickly left with bad ones. Buyers who search in the *Random available* will search for longer both because they have lower search costs on average and because the marginal benefit of search does not decrease as quickly as in the *Directed available* scenario.

The ex-ante probability of a match is substantially lower in the *Random available* scenario - only 0.11. However, conditional on search, the probabilities of a match in the *Directed available* and *Random available* scenarios are comparable. Overall, the results indicate that both buyers and sellers benefit significantly from the seller profile and reputation information, as more buyers are inclined to search, which leads to more matches overall. Buyers also search less intensely, and therefore at a total lower cost, when they know more about the sellers.

The completely *Random* scenario fares very badly compared to the previous two, with almost no search and very few matches. The function of extracting seller availability is essential to the functioning of this specific marketplace. Any fee or design feature that decreases the incentives for sellers to indicate their availability would be ill advised.

Next, consider the *Frictionless* scenario, where seither the marketplace manages to incentivize the sellers to reveal their private reservation prices at the onset of the game or it lowers buyer search costs to zero. The overall probability of a match is 0.61, or more than twice the unconditional match probability in the status quo. The increase in number of matches is limited by the sellers' outside employment opportunities. This is a partial equilibrium result: it is likely that a platform with frictions would attract more buyers, and the sellers would eventually become less reliant on their outside demand.

The counterfactual results also allow me to say something about how match efficiency is affected by the information frictions. In theory, contacting sellers at random could both lead to higher or lower match surplus compared to directional search. This is because the seller ordering by the net output f is not the same as the ordering by the match surplus s, which is not observed at search. For jobs that result in a match in both *Directed available* and *Random available* scenarios, the match surplus is 3 percent lower if the buyers do not observe seller characteristics. In other words, the inefficiency stemming from not fully observing s at the search stage is not substantial, conditional on a match being formed. The same exercise in reverse shows that surplus increases by 22 percent when buyers observe reservation prices for matches formed under both the *Frictionless* and *Directed available* scenarios. This reveals a significant potential for improvement in match efficiency if more information on seller reservation prices is made public or costless to access.

Table 1.6: Counterfactual analysis results.

	Frictionless	Directed available	Random available	Random
N. sellers searched	-	1.30	0.90	0.002
Pr(Search)	-	0.61	0.22	0.001
N. sellers searched | Search	-	2.12	4.02	1.21
Pr(Match)	0.61	0.28	0.11	0.000
Pr(Match|Search)	-	0.46	0.51	0.43
Change in surplus relative to *Directed available* (%)	22	-	-3	

1.7 Conclusion

As we see more service markets moving their activity online, the availability of data and the scope for improving the economic outcomes in such settings open up many interesting research questions. The contribution of this paper is to model structurally for the first time how a buyer-seller match is formed, given the special features of service markets: small and differentiated participants, private seller information, and costly buyer search. I explicitly the model information frictions associated with seller fixed capacity - availability and reservation price - in a general way that can apply to other similar environments. Both oDesk and AirBnB experience a significant impact on the match outcome from sellers rejecting buyers due to preferences. The model can be further enriched by price and offer data, which would allow to estimate search costs without assumptions on how the match surplus is allocated and to identify separately how certain characteristics affect the output and the average seller reservation price.

The structural model allows me to construct counterfactual information scenarios and to evaluate the efficiency contribution of the different features of the online marketplace. Information on seller availability has a substantial impact on all outcomes, and any design or fee structure that affects this should be carefully considered. Lack of information on seller characteristics and experience does not substantially affect match efficiency for those matches that would be formed if this information were present. However, this it greatly affects the number of sellers who search and the duration of search with its associated total search costs. Removing the last information friction, the private seller reservation prince, has the potential to double the number of matches formed on the platform and to significantly improve efficiency for already existing matches.

1.8 References

Chade, Hector, and Lones Smith. "Simultaneous search." Econometrica 74.5 (2006): 1293-1307.

Chade, Hector, Jan Eeckhout, and Lones Smith. "Sorting through search and matching models in economics." Journal of Economic Literature 55.2 (2017): 493-544.

Chiappori, Pierre-Andre, and Bernard Salanie. "The econometrics of matching models." Journal of Economic Literature 54.3 (2016): 832-861.

De los Santos, Babur, Ali Hortaçsu, and Matthijs R. Wildenbeest. "Testing models of consumer search using data on web browsing and purchasing behavior." American Economic Review 102.6 (2012): 2955-80.

Dubois, Pierre, and Helena Perrone. "Price dispersion and informational frictions: Evidence from supermarket purchases." (2015).

Cullen, Zoë, and Chiara Farronato. "Outsourcing tasks online: Matching supply and demand on peer-to-peer internet platforms." Job Market Paper (2014).

Fradkin, Andrey. "Search, matching, and the role of digital marketplace design in enabling trade: Evidence from Airbnb." (2017).

Hortaçsu, Ali, and Chad Syverson. "Product differentiation, search costs, and competition in the mutual fund industry: A case study of S&P 500 index funds." The Quarterly Journal of Economics 119.2 (2004): 403-456.

Horton, John. Buyer Uncertainty about Seller Capacity: Causes, Consequences, and Partial Solution. Working Paper, 2016.

Krishna, Vijay. Auction theory. Academic press, 2009.

Moraga-Gonzalez, Jose, Zsolt Sándor, and Matthijs Wildenbeest. "Consumer search and prices in the automobile market." (2015).

Roth, Alvin E., and Marilda Sotomayor. "Two-sided matching." Handbook of game theory with economic applications 1 (1992): 485-541.

Stigler, George J. "The economics of information." Journal of political economy 69.3 (1961): 213-225.

Tadelis, Steven. ”Reputation and feedback systems in online platform markets.” Annual Review of Economics 8 (2016): 321-340.

Weitzman, Martin L. ”Optimal search for the best alternative.” Econometrica: Journal of the Econometric Society (1979): 641-654.

Ye, Lixin. ”Indicative bidding and a theory of two-stage auctions.” Games and Economic Behavior 58.1 (2007): 181-207.

1.9 Appendix

1.9.1 Job descriptive statistics

Table 1.7: Jobs buyer expected start.

	Freq.	Percent
After reviewing offers	185	4.44
Immediately	564	13.53
In 2 days	806	19.34
In 2 months	396	9.50
In 2 weeks	1,267	30.41
Just checking offers	456	10.94
More than 2 months	91	2.18
Unknown	402	9.64
Total	4,167	100

Table 1.8: Jobs buyer proposed budget.

Job budget	Freq.	Percent
Missing	482	11.57
25	283	6.79
50	325	7.80
100	467	11.21
250	454	10.90
500	481	11.54
1,000	404	9.70
2,000	263	6.31
2,500	150	3.60
4,000	106	2.54
8,000	77	1.85
30,000	675	16.20
Total	4,167	100

Table 1.9: Job categories.

Category	Freq.	Percent
Architecture and design	59	1.42
Bathroom repair	475	11.40
Building restoration and insulation	227	5.45
Carpentry	188	4.51
Cleaning services	56	1.34
Construction	253	6.07
Demolish, clean and transport	61	1.46
Doors and barriers	157	3.77
Dry construction	133	3.19
Electrical repairs	288	6.91
Equipment repair	98	2.35
Floors: parquet and tiles	295	7.08
Furniture	234	5.62
Kitchen repair	103	2.47
Landscaping	37	0.89
Metalworking	36	0.86
Painting and decoration	493	11.83
Roof repairs	282	6.77
Sewage and sanitation	390	9.36
Smithery services	92	2.21
Window pane and glass repairs	210	5.04
Total	4,167	100

1.9.2 Additional derivations

Second order condition check

Using our distributional **Assumption 4** and the estimated $\hat{f}$ for each job, I calculate the marginal benefit of search for each job for up to 10 available sellers. I follow the optimal seller search rule, by which the buyer adds sellers to the search set in order of decreasing $\hat{f}$. Let n be the number of sellers in the search set. According to the results in the table below, the marginal benefit of search is decreasing as more sellers are added to the search set.

Table 1.10: Average change in the marginal benefit of search for the jobs in our sample.

	Obs.	Average change
MB(n=2)-MB(n=1)	3,414	1.29
- % obs. where this is negative		29
MB(n=3)-MB(n=2)	2,737	-0.28
- % obs. where this is negative		100
MB(n=4)-MB(n=3)	2,155	-0.12
- % obs. where this is negative		100
MB(n=5)-MB(n=4)	1,718	-0.06
- % obs. where this is negative		100
MB(n=6)-MB(n=5)	1,363	-0.04
- % obs. where this is negative		100
MB(n=7)-MB(n=6)	1,112	-0.03
- % obs. where this is negative		100
MB(n=8)-MB(n=7)	910	-0.02
- % obs. where this is negative		100
MB(n=9)-MB(n=8)	748	-0.02
- % obs. where this is negative		100
MB(n=10)-MB(n=9)	601	-0.01
- % obs. where this is negative		100

The exception is going from 1 sellers to 2 sellers. Searching 0 and 1 sellers gives the buyer expected utility of zero. In the first case, this is because he takes his outside option, which I set to zero for all buyers. In the second case, because he seller extracts all the surplus. Therefore, $MB(n = 1) = 0$. In the cases where $MB(n = 2) > 0$, the buyer will search $n > 0$ sellers. In the cases where $MB(n = 2) < 0$, the buyer is indifferent between searching 0 and 1 sellers.

MLE derivation

The sample of M jobs is comprised of independent observations, indexed by i. Each job can have one of the following outcomes: either no one is hired, or one of the contacted sellers is hired. Let the probability of the first event be $Pr(Y_i = 0)$, and the probability that seller $j \in \{n_i\}$ is hired be $Pr(Y_i = j)$. Let d_{ij} be an indicator equal to 1 when buyer i hires seller j, and zero otherwise. The indicator for the event that no one is hired is $d_{i0} = 1 - \sum_{j=1}^{j=n_i} d_{ij}$. The probability for any outcome for a job i is:

$$Pr(Y_i = 0)^{1-\sum_{j=1}^{j=n_i} d_{ij}} \prod_{j=1}^{j=n_i} Pr(Y_i = j)^{d_{ij}}$$

The likelihood of any particular observation in the sample is the above probability expressed as a function of the unknown parameters γ and β, conditional on the set of covariates: $X_i = X_{i1}, ..., X_{in_i}$ and Z_i, as well as the set of contacted sellers $\{n_i\}$. By independence of the individual observations i, the likelihood of the full sample is the product of the individual probabilities of the observations:

$$: L(\gamma, \beta | X, Z, \{n\}) = \prod_{i=1}^{i=M} L(\gamma, \beta | X_i, Z_i, \{n_i\}) =$$

$$= \prod_{i=1}^{i=M} \left(Pr(Y_i = 0)^{1-\sum_{j=1}^{j=n_i} d_{ij}} \prod_{j=1}^{j=n_i} P(Y_i = j)^{d_{ij}} \right)$$

The MLE estimator works by maximizing $L(\gamma, \beta | X, Z, \{n\})$ with respect to the parameters γ, β. The first step is to derive the closed form expressions for the two events $Pr(Y_i = 0)$ and $Pr(Y_i = j)$. First, consider the event of hiring seller j with probability $Pr(Y_i = j)$. From match stability, I know that this event will take place under the following two conditions: no better match can be formed with any other available seller and both buyer and seller are better off in the match than unmatched. The probability of hiring seller j is therefore the following:

$$Pr(Y_i = j) = Pr(S_{ij} > 0 \cup S_{ij} > S_{ik}, \forall k \neq j)$$

The first part of the probability is a condition on the level of the net match surplus: it must be positive in order to at least compensate the buyer and seller for giving up their outside options. It will allows to identify the overall level of surplus generated by the match, hence both the coefficients on variables X_{ij} that differ across sellers and the coefficients of the variables Z_i that are constant across buyer-seller pairs. Let $-R = \rho$ and express this as:

$$f_{ij} + \rho_{ij} > 0 \Leftrightarrow \rho_{ij} > -f_{ij}$$

The second part of the probability states that the surplus created by seller j is greater than the surplus created by any of the other sellers $k \neq j$. This is a condition on the difference between the match surplus generated by the different sellers, and therefore it identifies coefficients of variables whose values differ across sellers, the X_{ij}. Consider comparing two alternatives j and k, implying:

$$f_{ij} + \rho_{ij} > f_{ik} + \rho_{ik} \Leftrightarrow \hat{\rho}_{ijk} = \rho_{ij} - \rho_{ik} > -(f_{ij} - f_{ik}) = -\hat{f}_{ijk}$$

where $\hat{\rho}$ is the difference between two randomly drawn realizations of ρ and $\hat{f}$ is the difference between two net match outputs. I can therefore re-write $Pr(Y_i = j)$ as:

$$Pr(Y_i = j) = Pr(\rho_{ij} > -f_{ij} \cup \hat{\rho}_{ijk} > -\hat{f}_{ijk}, \forall k \neq j)$$

The two events in this probability are not independent because the realization of the random variable ρ_{ij} enters both. Using the distributional assumption for ρ, I derive the closed form for the joint distribution of the random variables ρ and $\hat{\rho}$. To simplify the exposition, I present the derivation for the case where $n_i = 2$. The joint distribution of ρ_1 and ρ_2, both distributed independently and identically by a Type 1 Extreme value distribution, is the product of the marginals:

$$Pr(\rho_1 = a \cup \rho_2 = b) = g_{\rho_1,\rho_2}(a,b) = g_{\rho_1}(a)g_{\rho_2}(b) = e^{-a}e^{-e^{-a}}e^{-b}e^{-e^{-b}}$$

Now, consider the following transformation: $\rho = h_1(\rho_1, \rho_2) = \rho_1$ and $\hat{\rho} = h_2(\rho_1, \rho_2) = \rho_1 - \rho_2$. With a slight abuse of notation, the inverse transformation functions are $\rho_1 = h_1^{-1}(\rho, \hat{\rho}) = \rho$ and $\rho_2 = h_2^{-1}(\rho, \hat{\rho}) = \rho - \hat{\rho}$.

The Jacobian of the matrix of derivatives of the inverse functions h^{-1} with respect to their arguments:

$$J = \begin{vmatrix} \frac{\delta h_1^{-1}}{\delta \rho} & \frac{\delta h_1^{-1}}{\delta \hat{\rho}} \\ \frac{\delta h_2^{-1}}{\delta \rho} & \frac{\delta h_2^{-1}}{\delta \hat{\rho}} \end{vmatrix} = \begin{vmatrix} 1 & 0 \\ 1 & -1 \end{vmatrix} = -1$$

The joint probability distribution of the transformed random variables ρ and $\hat{\rho}$ is:

$$g_{\rho,\hat{\rho}}(a,b) = g_{\rho_1,\rho_2}(h_1^{-1}(a,b), h_2^{-1}(a,b))|J| = g_{\rho_1,\rho_2}(a, a-b) = e^{-a}e^{-e^{-a}}e^{-(a-b)}e^{-e^{-(a-b)}}$$

Using this, I derive the cumulative distribution of the following event:

$$: Pr(\rho > -f \cup \hat{\rho} > -\hat{f}) = \int_{-f}^{\infty}\int_{-\hat{f}}^{\infty} e^{-a}e^{-e^{-a}}e^{-(a-b)}e^{-e^{-(a-b)}}dbda =$$
$$= \int_{-f}^{\infty} e^{-a}e^{-e^{-a}} \int_{-\hat{f}}^{\infty} e^{-(a-b)}e^{-e^{-(a-b)}}dadb = \int_{-f}^{\infty} e^{-a}e^{-e^{-a}}e^{-e^{-(a+\hat{f})}}da =$$
$$= \int_{-f}^{\infty} e^{-a}e^{-e^{-a}(1+e^{-\hat{f}})}da = \frac{1-e^{-e^{f}(1+e^{-\hat{f}})}}{1+e^{-\hat{f}}}$$

For a larger set of contacted sellers, the absolute value of the matrix J remains 1. Thus, the probability $Pr(Y_i = j)$ for any number n_i of contacted sellers j is:

$$: Pr(Y_i = j, n_i) = Pr(\rho_{ij} > -f_{ij} \cup r\hat{h}o_{ijk} > -\hat{f}_{ijk}, \forall k \neq j) =$$
$$= \frac{1-e^{-e^{f_{ij}}(1+\sum_{jk} e^{-\hat{f}_{ijk}})}}{1+\sum_{jk} e^{-\hat{f}_{ijk}}} = \frac{1-e^{-\sum_j e^{f_{ij}}}}{1+e^{-f_{ij}}\sum_{k\neq j} e^{f_{ik}}}$$

Now consider the event that no one is hired. When a single seller j is contacted, the probability that he is not hired, $Pr(Y_i = 0)$, can be expressed as:

$$Pr(Y_i = 0) = Pr(\rho_{ij} < -f_{ij}) = e^{-e^{f_{ij}}}$$

By the independence of the ρ_{ij}'s across j, with n_i sellers j this is simply the following:

$$Pr(Y_i = 0, n_i) = \prod_j Pr(\rho_{ij} < -f_{ij}) = e^{-\sum_j e^{f_{ij}}}$$

MLE implementation

I construct a set of indicator variables I_{ij} for $j = 1, ..., \overline{n}$, where each I_j is equal to 1 whenever seller j is contacted for the job, and zero otherwise. $\overline{n}$ is the maximum number of sellers that were ever contacted for any job. In the theoretical model, the sellers j are ordered by descending f_{ij} for any job i, but the position/order of the sellers is not important for the construction of the probability. This way $\sum_{j=1}^{j=\overline{n}} I_{ij} = n_i$ whenever n sellers are contacted for job i.

Using these indicators, the MLE formula can be generalized in the following way:

$$Pr(Y_i = 0, n_i) = e^{-\sum_{j=1}^{j=\overline{n}} I_{ij} e^{f_{ij}}}$$

$$Pr(Y_i = j, n_i) = \frac{1 - Pr(Y_i = 0, n)}{1 + I_{ij} e^{-f_j} \sum_{k \neq j} I_{ik} e^{f_{ik}}}$$

Consider the full form of the likelihood derived in **Section 1.4**. There are j cases for each of the j sellers that can potentially be hired: $d_{ij} Pr(Y_i = j, n_i)$, where d_{ij} is 1 whenever j is hired for job i and zero otherwise. In the data, any seller can take the position j for any job i: the likelihood derivation is not affected by the seller order. With a random ordering of the sellers, each d_{ij} can contain 1's across the sample of observations. Thus, I would have to code all probabilities $Pr(Y_i = j, n_i)$. However, as the order/position of the sellers does not matter for the likelihood derivation, I can simplify the problem in the following way: for each *observation* i whenever someone was hired, always order that seller who was hired first: $j = 1$ whenever $d_{ij} = 1$. Thus, only the indicator d_{ij} is ever non-zero and all indicators $d_{ij'}$ for $j' \neq j$ are always zero in the sample. The cases $d_{ij'} Pr(Y_i = j', n_i)$ for any sellers $j' \neq j$ are always zero and drop out of the likelihood.

MLE probabilities sum to 1

Consider the case of $n = 2$ and the 3 probabilities:

$$Pr(Y = 0, n = 2) = e^{-e^{f_1} - e^{f_2}} = A$$

$$Pr(Y = 1, n = 2) = \frac{1 - e^{-e^{f_1} - e^{f_2}}}{1 - e^{-f_1} e^{f_2}} = \frac{1 - A}{B}$$

$$Pr(Y = 2, n = 2) = \frac{1 - e^{-e^{f_1} - e^{f_2}}}{1 - e^{-f_2} e^{f_1}} = \frac{1 - A}{C}$$

I sum them:

$$Pr(Y = 0, n = 2) + Pr(Y = 1, n = 2) + Pr(Y = 2, n = 2) = \frac{C(1 - A) + B(1 - A) + BCA}{BC}$$

However, I can show that:

$$BC = 1 - e^{-f_1} e^{f_2} - e^{-f_2} e^{f_1} + e^{-f_1} e^{f_2} e^{-f_2} e^{f_1} = 1 - e^{-f_1} e^{f_2} - e^{-f_2} e^{f_1} + e^{0} = B + C$$

Therefore, the sum of the probabilities is exhaustive:

$$\begin{aligned} & Pr(Y = 0, n = 2) + Pr(Y = 1, n = 2) + Pr(Y = 2, n = 2) = \\ & \qquad = \frac{C(1 - A) + B(1 - A) + (B + C)A}{B + C} = 1 \end{aligned}$$

Special cases for estimating seller search costs

We have the following 5 types of jobs:

1. Jobs i for which $n_i < N_i$ and $n_i > 1$. Both upper $\hat{\bar{c}}_i$ and lower$\hat{\underline{c}}_i$ bounds on the search cost c_i can be estimated. There are 1,036 jobs of this type.

2. Jobs i for wihch $n_i = N_i$ and $n_i > 1$. Only the upper bound on the search cost, $\bar{c}_i$ can be estimated. The lower bound cannot be estimated because we do not have another available seller to construct the set $\{n_i + 1\}$. Therefore, $\underline{c}_i$ is approximated using a conditional average of $\hat{\underline{c}}$ across all other jobs i', such that $\hat{\underline{c}}_{i'} <= \hat{\bar{c}}_i$. There are 381 jobs of this type.

3. Jobs i for which either $n_i = 0$ or $n_i = 1$ and $N > 1$. For these jobs, the buyer is indifferent between searching 1 or 0 sellers because in both cases his expected utility is zero. Only the lower bound on the search cost, $\underline{c}_i$ can be estimated for $\{n_i + 1\} = 2$. The upper bound cannot be estimated because there cannot be a set $\{n_i - 1\}$ that gives positive expected utility to the buyer. Therefore, $\bar{c}_i$ is approximated using a conditional average of $\hat{\bar{c}}$ across all other jobs i', such that $\hat{\bar{c}}_{i'} >= \hat{\underline{c}}_i$. There are 2,038 jobs of this type.

4. Jobs i for which either $n_i = 0$ or $n_i = 1$ and $N = 1$. The buyer gets zero expected utility in either $n_i = 0$ or $n_i = 1$, and cannot search more than that because there are no more available sellers. Therefore, neither search cost bound can be estimated. $\bar{c}_i$ is approximated using a conditional average of $\hat{\bar{c}}$ across all other jobs i', such that $\hat{\bar{c}}_{i'} >= \hat{\underline{c}}_i$. $\underline{c}_i$ is approximated using a conditional average of $\hat{\underline{c}}$ across all other jobs i', such that $\hat{\underline{c}}_{i'} <= \hat{\bar{c}}_i$. There are 737 jobs of this type.

Robustness checks

Table 1.11: Results from the MLE estimations with seller fixed effects.

	M1	M2	M3	M4
Xij				
Seller percent positive reviews		0.860***		0.437*
		(0.173)		(0.176)
Seller marketplace tenure		-0.083		-0.179**
		(0.050)		(0.059)
Seller total times hired		0.000		0.000
		(0.000)		(0.000)
Message length		-0.070**		0.032
		(0.027)		(0.029)
Message time		-0.143**		0.038
		(0.051)		(0.126)
Zi				
Job budget			-0.252***	-0.261***
			(0.069)	(0.070)
Job start			-0.350***	-0.363***
			(0.078)	(0.079)
High demand season			-0.328***	-0.295***
			(0.084)	(0.086)
Year fixed effects			Yes***	Yes***
Job category fixed effects			Yes***	Yes***
Seller fixed effects	Yes***	Yes***	Yes***	Yes***
Constant	-2.228*	-1.751	-0.165	-0.248
	(1.031)	(1.036)	(1.297)	(1.220)
Likelihood	-2,347	-2,327	-2,073	-2,066
Observations	2,556	2,556	2,556	2,556

Significant at: $p < 0.1$: *; $p < 0.05$: **; $p < 0.01$: ***. Continuous variables are transformed by taking the natural logarithm. More detailed results are available upon request.

Chapter 2

Reputational Incentives under Heterogeneous Demand Fluctuations

Abstract

Using data from an online home services marketplace, we study reputational incentives in an environment with heterogeneous seller demand fluctuations. We empirically corroborate the hypothesis that sellers with an upcoming demand downturn value their reputation less, which in turn affects their effort negatively. Our robustness checks rule out that this is due to selection of sellers into jobs, higher costs of seller effort, or adverse selection. Neither the economic literature nor the platforms designing reputation systems have sufficiently considered seller heterogeneity and the optimal way to account for it in terms of the information available to the buyer and the seller incentives.

2.1 Introduction

Online platforms have successfully expanded the economic activity in many diverse markets such as the sale and re-sale of various new and used goods (eBay, Amazon) and short-term rentals of cars and property (Uber, AirBnB). These markets were previously impeded by the issues of anonymity and lack of trust: the sellers and buyers did not know anything about each other, and their interaction was one-time only. The incentives to cheat, by providing a low quality good or service, were significant enough to make economic activity in these markets non-existent or significantly depressed.

To a large extent, the success of online platforms can be credited to their employment of *reputation systems*. *Reputation systems* (or *reputation mechanisms*) are a set of rules by which buyer-seller interactions are reviewed publicly.[1] Even if individual buyers interact only once with the seller, the mechanism establishes a dependence between current and future demand. Consider a setting with *moral hazard*, where effort is costly and unobserved. By subscribing to a marketplace with a public reputation system, the seller agrees that the outcome of all interactions is made public, which in turn informs the decision of potential buyers. Knowing that the sellers value the ability of reputation to generate demand in future time periods, the buyers are less hesitant to hire them in the current period.

Despite how ubiquitous reputation systems have become, the approach to their design is still ”one size fits all”, with limited consideration of seller heterogeneity and the economic environment more generally. The predominant way in which individual reviews are aggregated and displayed is by taking a simple average, thus making each review equally improtant. For example, the biggest review systems for restaurants - Tripadvisor, Yelp and GoogleMaps - use this simple approach. Amazon is one of the few online retailers that recognize the need for a more fine-tuned reputation system: it gives higher visibility of reviews received in the last 12 months to prevent sellers lowering quality when a large amount of positive reviews is already accumulated. In their working paper on the economics and the tech industry, Athey and Luca (2018) note that platforms have to consider more carefully the way in which review information is made available to buyers, and how this affects seller incentives. Economists have only recently started to addressed

[1] The role of reputation is not new to economic theory and it is studied in the following two cases of asymmetric information with regards to the seller. In settings with unobserved seller actions (*moral hazard*), reputation is a sanctioning device that motivates the seller to take appropriate actions. In settings with unobserved seller types (*adverse selection*), reputation is a signaling device that reveals the true type of the seller.

optimality of reputation systems in specific economic environments, and how this can be achieved by review aggregation and display (see Dai, Jin, Lee and Luca (2018)).

In this paper, we provide evidence that when the economic environment and actors are not homogeneous, a simple reputation system fails to incentivize a consistent effort level. We consider an environment with demand fluctuations and sellers who are differentially exposed to them. Because reputation works by linking current and future demand, when future transaction opportunities are few, the value placed on a good reputation by the seller is low. This lowers the incentive to exert effort, which in turn leads to a lower probability that the buyer is satisfied by a project's outcome. Consistent with the hypothesis, the empirical results indicate that the sellers who experience a more significant drop in demand in future months are less likely to receive a positive review.

We work with proprietary data from MaistorPlus, a Bulgarian online marketplace for home services. MaistorPlus is similar to other well-known platforms connecting clients, or *buyers*, and professionals, or *sellers*, such as oDesk (IT and business professionals), TaskRabbit (for home owners and low skilled labor), and Thumbtack (for a large variety of local services). The seller's account has information on his scope of expertise, but most importantly it contains a public record of his activity on the platform: the number of times he was hired and the corresponding average of the received reviews.[2] This information is important in determining if he is contacted and eventually hired by buyers. Our preliminary empirical analysis demonstrates the positive returns to reputation when demand is high.

The source of variation in our data comes from the seasonality of demand for certain jobs, which is due to weather conditions permitting outdoor work and to weather-proofing homes for the fall, winter and spring seasons. The majority of jobs are posted between July and November, which we call the high demand season. At the same time, the sellers themselves specialize in different services and this means that they are heterogeneously exposed to the demand seasonality. For example, someone specializing in plumbing services would see a more steady flow of demand compared to someone specializing in roof repairs, which mostly take place in the summer and autumn. As a result, the seasonal fluctuation of demand, which translates to a reputational incentive to exert effort, varies at the level of the individual seller. The analysis focuses on the seller' behavior at the end of the high demand season, when the discounted sum of work opportunities in the

[2] We refer to the individual interaction feedback as a *review*, while the cumulative of the reviews is the seller's *reputation*. In this sense, the review is an incremental change in reputation.

immediate future is the lowest.[3]

Our measure of demand heterogeneity is seller demand *Smoothness*: the average monthly demand in the low period relative to the average monthly demand in the high period. A smoothness measure closer to one implies lower exposure to seasonality, but although most sellers are affected at least to some extent. We construct this variable using demand in the categories in which the seller is active at any point in time in our sample.[4] This way, *Smoothness* measures the exogenous demand conditions faced by the seller rather than that seller's endogenous activity. This avoids the following reverse causality problem: after a positive review, the sellers may increase (or decrease) his activity, which in turn would affect the calculation of demand smoothness in that period.

On its own, the end of the high demand season is associated with a lower likelihood of a positive review across all sellers. We do not attribute this finding to the imminent reduction in overall demand, as the end of the high demand season is potentially associated with other non-effort determinants of project success, such as the disruption in material supplies, bad weather or last-minute projects that are harder to perform well. By interacting seller-specific demand *Smoothness* and the *High season end*, we are sure to find an effect which is driven by heterogeneous demand conditions at the level of the individual seller. We find that the sellers whose demand drops substantially in the upcoming low demand season are less likely to receive a positive review. The average seller in our sample has a 4 percent lower probability of a positive review in the last months of the high season. For sellers with low demand smoothness at the 10th percentile of the distribution, we see a drop of 16 percent, while sellers at the 90th percentile see an increase of 18 percent in the probability of a positive review. The results in the reverse direction are weekly significant in some specifications: the end of the low demand season is associated with a higher likelihood of a positive review, but less so for sellers with more seasonal demand.

Given the current reputation mechanism, our results can be used to inform buyers better about seller incentives. Consider roof repairs, a highly seasonal job due to weather conditions. Towards the end of the high demand season, a homeowner would be better off to hire a seller specializing in plumbing but also fixing roofs (high *Smoothness*), rather than the roof specialist (low *Smoothness*), all else equal. Alternatively, the homeowner looking at the profile of the roof repairman should not be as concerned by a negative re-

[3] The reverse effect - a higher probability of a positive review towards the end of the low demand season for high seasonality sellers - is also present in the results but significant only in certain specifications.

[4] In the context of buyer-seller interaction structure on the marketplace, the seller activity in a category is indicated by the seller responding to a job in that particular category.

view that was received in November as much as by one received in June, especially if he is looking to hire the roof repairman at the start of the high demand season. The results are suggest that design the system itself may be optimized in a way that provides more consistent incentives for seller effort. For example, the reputation system can aggregate seller reviews through a weighted average that puts higher weights on reviews received under circumstances similar to the ones facing the particular seller.

Our hypothesis assumes implicitly that the channel through which demand affects the outcome is seller effort. We perform a number of robustness checks to rule out alternative channels of causality. In a first robustness check, we consider selection of professionals into jobs, where some jobs require more effort, especially towards the end of the high demand season. We also rule out the possibility that sellers who drop out of the platform at the end of the high season, and may be of a ”bad” type, drive our results. Additionally, we propose two specifications that test indirectly whether the cost of effort, as a function of demand or of the size of the seller, confounds the results.

Our work is related to the new literature which uses data from online marketplaces to test formal reputation models, to study reputation dynamics, and to evaluate the effectiveness of reputation systems. The surveys of Bajari and Hortacsu (2004), Bar-Isaac and Tadelis (2008), Cabral (2012) and most recently Tadelis (2016) document this body of work. A substantial number of papers in the literature attempt to test whether the source of informational asymmetry in the market is due to moral hazard or adverse selection. For example, Cabral and Hortacsu (2010) examine the dynamics of seller reputation on eBay and find strong evidence for moral hazard: the first negative feedback received lowers the returns to reputation once and for all, and sellers react by lowering their effort and eventually exiting. However, the authors cannot rule out an underlying moral hazard plus adverse selection model. While the objective of our paper is not to rule out adverse selection, it is seller effort, rather than seller type, which is affected by the demand fluctuation. In the robustness checks, we rule out two alternative explanations for the effect we find that are based on seller type (seller size and exiting sellers).

The effectiveness of reputation systems has raised considerable research interest, with the majority of work focusing on the ability of reputation mechanisms to elicit truthful reviews. For example, Fradkin, Grewal and Holtz (2018) investigate the informativeness of reviews on a hospitality marketplace, AirBnB, where review reciprocity causes significant bias. Horton and Golden (2015) also document reputation inflation, on oDesk, a service marketplace. Only recently have economists considered more explicitly how seller reviews should be aggregated, with the objective of providing future buyers with the most up to date information and sellers with the right incentives. The work of Dai, Jin, Lee and Luca (2018) considers the optimal design of a reputation system in an environment with changing service quality. Because quality is exogenous, the effect of the reputation mechanism design on seller incentives is not a part of their analysis. We are not aware of other empirical work studying the heterogeneity of reputational incentives across sellers.

Although their work concerns the credit rating industry, Bar-Isaac and Shapiro (2013) are closer to the topic of our work as they investigate how the exogenous variation of the business cycle, a common demand trend, drives reputational incentives for Credit Rating Agencies (CRAs). The value of reputation is shown to depend on the economic fundamentals varying over the business cycle. Ratings accuracy is counter-cyclical as effort is lower during periods of high economic activity, which is consistent with our findings that the probability of a positive review is influenced by demand fluctuations.

The paper proceeds as follows. **Section 2.2** provides background on our data source, the MaistorPlus marketplace, including information on the categories of jobs available and their seasonal variation, the sellers and the reputation mechanism. **Section 2.3** presents the empirical implementation, as well as the results and robustness checks. **Section 2.4** concludes.

2.2 MaistorPlus and data generation

We work with company data form the MaistorPlus marketplace, founded in Sofia, Bulgaria in 2012, and our sample covers the period between January 2013 and July 2015.[5] The marketplace connects buyers to subscribing home service sellers and is financed by sellers' subscription fees and by advertisements.

2.2.1 Job activity and reviews

Demand on the marketplace is generated by buyers who sign up freely and post what we call $jobs$: home repairs projects for which they want to hire a service provider. We observe 4,167 jobs, as well as data at the job-seller level of interaction. The total value of all jobs in our sample is 12.6 million Euros.

[5] The website of the market place is: http://maistorplus.com/

The buyer provides a description of the job, specifies the job category (such as carpentry, roof repairs, construction, etc), an estimated budget, and a start date. All sellers active in that job category are potentially available and notified of the job; those who are actually available can message the buyer. The buyer is free to inspect the seller profiles and their reputation when he decides who to contact or hire. In the data, the buyers contact 1.3 sellers on average, and hire someone in 27 percent of the time. This information can be found in *Table 2.1*.

Table 2.1: Job activity on MaistorPlus

Data	Observations	Per job
Distinct jobs	4,167	-
- potentially available sellers	333,639	80.0
- available sellers	22,379	5.37
- contacted sellers	5,354	1.28
- hired sellers	1,126	0.27
- positive reviews (total and per hire)	749	0.66
- total value (Euro)	12,650,600	2,785

Only buyers who have hired a seller through the marketplace are allowed to post a review, which can be positive, negative, or neutral, or nothing (1/0/-1/.). As it is customary in the empirical reputations literature, we separate the reviews into positive and non-positive reviews.[6] Out of 1,126 jobs where someone was hired in our sample, we observe that the buyer left a positive review in 749 cases.[7]

For buyer reviews to carry information, it is necessary that the good or service in question can be competently evaluated by the buyer. Services performed by expert professionals, such as home repairs, can have similar aspects to credence goods as the buyer only observes whether the problem is fixed, so to say, but not the quality of the work or whether it was overperscribed. Seller professionalism, as well as explanations of how the problem arose and how to present future problems, are therefore highly valued by the buyers and the predominant themes of the textual part of the reviews. Providing this observable aspect of the service is still costly and affected by fluctuations in the value of reputation.

[6] Consistent with the previous literature (e.g. Nosko and Tadelis (2015) and Dellarocas and Wood (2008)), we view no feedback and neutral feedback as a non-positive review.

[7] The buyers can also post a textual review of the interaction. However, we do not work with this metric because it is highly collinear with the point review.

2.2.2 Service sellers

There is a total of 864 active sellers in our sample, of whom 251 were hired at least once. They subscribe to one of three fixed-term plans, ordered by their annual cost: Start (100 Euro), Pro (150 Euro)and Pro+ (250 Euro). The more costly the plan, the more options the sellers have for the number of categories in which they are active (3/4/6 categories). The sellers can change these at any time and at no cost, potentially adjusting their activity to seasonal demand fluctuations. While the buyers can infer, to some extent, the seller size and scope of activity by observing the current active categories, they do not observe history of category changes, and therefore seller's demand smoothness.

Table 2.2 presents summary statistics on the experience of the sellers present in the main analysis: the sellers who were hired at least once in the duration of our sample. The average seller is hired 4.5 times and has 3 positive reviews, both with high standard deviation. The average time between the last two messages of interest is about 9 days, but some sellers appear to have not been available much longer. The average activity span (time between the first and last message observed in the data) is 19 months out of 38 months in the data. Periods of high inactivity are likely for sellers who work in seasonal job categories, and among those who use the marketplace relatively rarely because they have sufficient job referrals from outside of the marketplace.

Table 2.2: The experience of the MaistorPlus sellers on the marketplace.

At latest observed time period:	Mean	St. dev	Min	Max
- activity span on marketplace (months)	19.19	8.56	0	30
- time between last two messages (days)	8.65	39.5	0	362
- total times hired	4.48	7.02	1	49
- total positive reviews	2.98	5.17	0	38

2.2.3 Seasonality and reputation

Demand in the home services sector is highly seasonal because some types of services, for example outdoor work, can only be done during the months when the weather permits. There are also services that are demanded right before the start of fall and winter, such as services related to heating and insulation. Examples of demand seasonality in the different job categories can be found in *Figure 2.1*.

To establish the duration of the high demand season, we look at the number of jobs posted

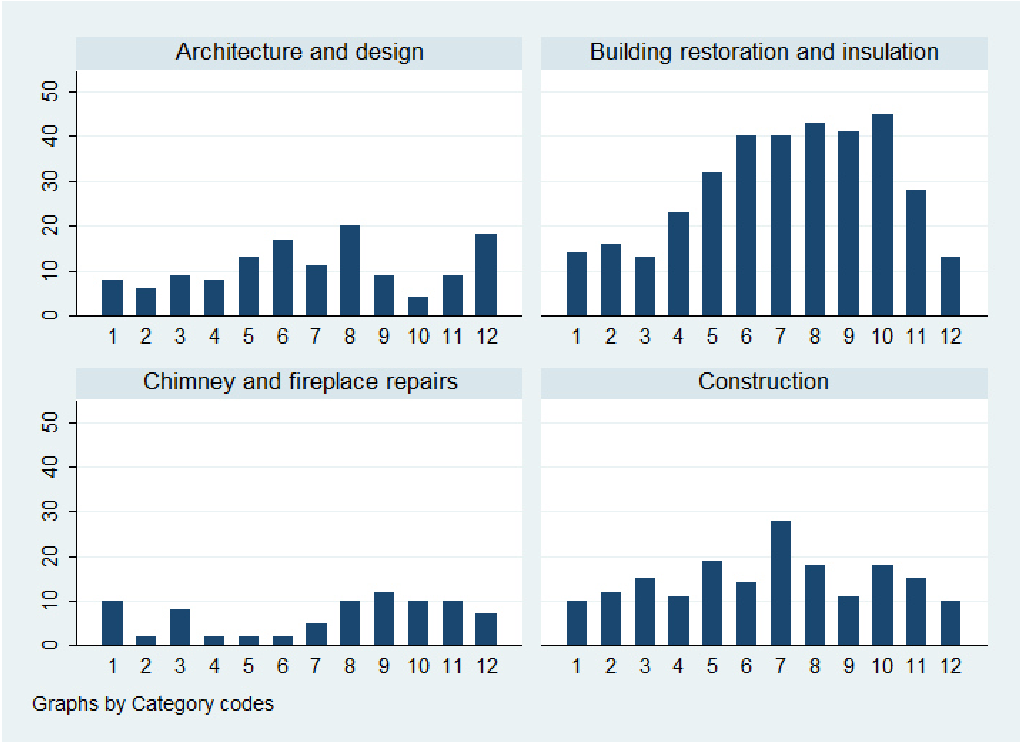

Figure 2.1: Number of jobs posted by both in different categories, year 2014

on the platform each month during the period 2013-2015, $N.\ jobs$. *Table 2.3* has the results from two regressions that help us do that. In the first regression $S1$, we use indicators for the individual months. Months July through November experience a significantly higher level of demand compared to other months. We define a seasonal variable *High season* equal to 1 for months months July-November and use that in the second regression $S2$. The explanatory power of $S2$ is still very high, and in fact the adjusted R^2 is higher compared to $S1$.

Implicit in our hypothesis that demand conditions affect effort is the conjecture that reputation is more valuable when demand is high. Ex-ante, it is not obvious whether reputation is more valuable during high or low demand periods. On the one hand, the value of reputation could be higher in the high demand period because there are more buyers with a high willingness to pay for a seller with a high reputation. On the other hand, reputation could be more valuable in the low demand season when there are fewer jobs and more intense competition between the sellers.

We resolve this question by going to the data. We construct a panel at the professional-date level and a variable measuring the unconditional returns to reputation as the number

Table 2.3: Monthly demand on the marketplace and the *High season* dummy.

Dependent variable	N. jobs	N. jobs
Specification	(S1)	(S2)
Independent variables		
February	-0.061	
	(0.311)	
March	0.236	
	(0.316)	
April	0.277	
	(0.297)	
May	0.319	
	(0.403)	
June	0.410	
	(0.426)	
July	0.988***	
	(0.314)	
August	1.016***	
	(0.338)	
September	0.945***	
	(0.317)	
October	0.828**	
	(0.305)	
November	0.801**	
	(0.323)	
December	0.250	
	(0.322)	
High season		0.711***
		(0.090)
Constant	4.410***	4.615***
	(0.299)	(0.094)
Year fixed effects	Yes	Yes
R^2	0.88	0.85
R^2 Adjusted	0.81	0.83
N. observations	36	36

Significant at: $p < 0.1$: *; $p < 0.05$: **; $p < 0.01$: ***.
Robust standard errors.

of times a seller was contacted by buyers in a given month, *Monthly contacts*.[8] The variables of interest in specification $C1$ are our measure of reputation, *Percent positive reviews*, lagged one month, and the *High season* indicator. Their interaction is included in specification $C2$. Other control variables include the seller's activity on the platform during that time (how many jobs he was invited to, how many times he was available), and date and seller fixed effects. The estimates are presented in *Table 2.4*. The benefits of a good reputation are indeed positive: the number of monthly contacts increases with *Percent positive reviews*. The results of the second regression support the claim that reputation is more valuable when demand is high as the coefficient on the interaction between *Percent positive reviews* and *High season* is positive and significant. If *Percent positive reviews* increases by 20 percentage points, *Monthly contacts* in the *High demand* period increase by 1.3 percent.

Table 2.4: Returns to reputation and the *High season*.

Dependent variable	Monthly contacts	Monthly contacts
Specification	(C1)	(C2)
Independent variables		
Percent positive reviews	0.038*	0.016
	(0.020)	(0.021)
High season	0.115***	0.107***
	(0.018)	(0.018)
High season*Percent positive reviews		0.065***
		(0.014)
Messages of availability	0.422***	0.422***
	(0.007)	(0.007)
Constant	-0.071	-0.069
	(0.051)	(0.051)
Fixed effects		
Date	Yes	Yes
Seller	Yes	Yes
R^2	0.77	0.77
N. observations	16,912	16,912

Significant at: $p < 0.1$: *; $p < 0.05$: **; $p < 0.01$: ***. Robust standard errors. The continuous dependent variables and regressors (except *Percent positive reviews*) are transformed by taking their natural logarithm, and their coefficients can be interpreted as elasticities. The coefficient on *Percent positive reviews* is a semi-elasticity.

[8] The hiring decision is joint and depends on other potential projects that the seller may be considering. Theregore, we prefer to use as dependent variable the number of times the buyers have expressed preference for hiring the seller, i.e. *Monthly contacts*.

2.3 Econometric analysis

Our econometric analysis is based on the difference in differences methodology. The results indicate that indeed sellers who experience a drop in demand are less likely to receive a positive review towards the end of the high season. We also see a weekly significant result in the reverse direction: sellers who experience an increase in demand are more likely to receive a positive review towards the end of the low demand season. We perform a number of robustness checks that corroborate this conclusion.

2.3.1 Demand smoothness

Central to empirical analysis is the variable measuring how the seller's heterogeneous demand fluctuates, $Smoothness_i$. Let k denote the individual job categories and i denote the sellers. The data allows us to observe the set $\{K_i\}$ of categories that seller i is available in. For example, ***Construction*** is one such category for seller i if the seller has expressed availability to a job in that category. Let $Demand_{k,Season}$ denote the number of unique jobs in category k during the respective season.[9] As the platform is still growing during the period of observation, we group two ***Low*** and two ***High*** seasons to improve the representativeness of the demand pattern. Assuming that demand grows multiplicatively, our definition of $Smoothness_i$ is not distorted by this. For the demand in the ***High season***, we use the jobs posted on the marketplace during July 2014-Nov 2014 and July 2015-Nov 2015. For demand during the ***Low season***, we use jobs during Dec 2013-June 2014 and Dec 2014 - June 2015.

For the seller i, the monthly demand during any given season, $Demand_{i,Season}$, is the sum of demand over his active categories $\{K_i\}$ in that season, $Demand_{k,Season}$, divided by the duration of the season.[10] The seller's demand $Smoothness_i$ is the ratio of of demand in the low season to demand in the high season. Constructed this way, the demand measure is exogenous to the seller's actual activity and reputation. In other words, $Smoothness_i$ is defined as:

[9] Our data rangers from January 2013 to December 2015 but there was a significant change in the marketplace rules in July 2015. The change affects the activity of the sellers, hence we cannot use data after it for the main analysis. However, it does not affect demand incidence.

[10] The low and high demand seasons have different duration, 5 and 7 months respectively. Dividing the total demand by the season length is necessary to avoid a mechanical difference due to the different length of the seasons.

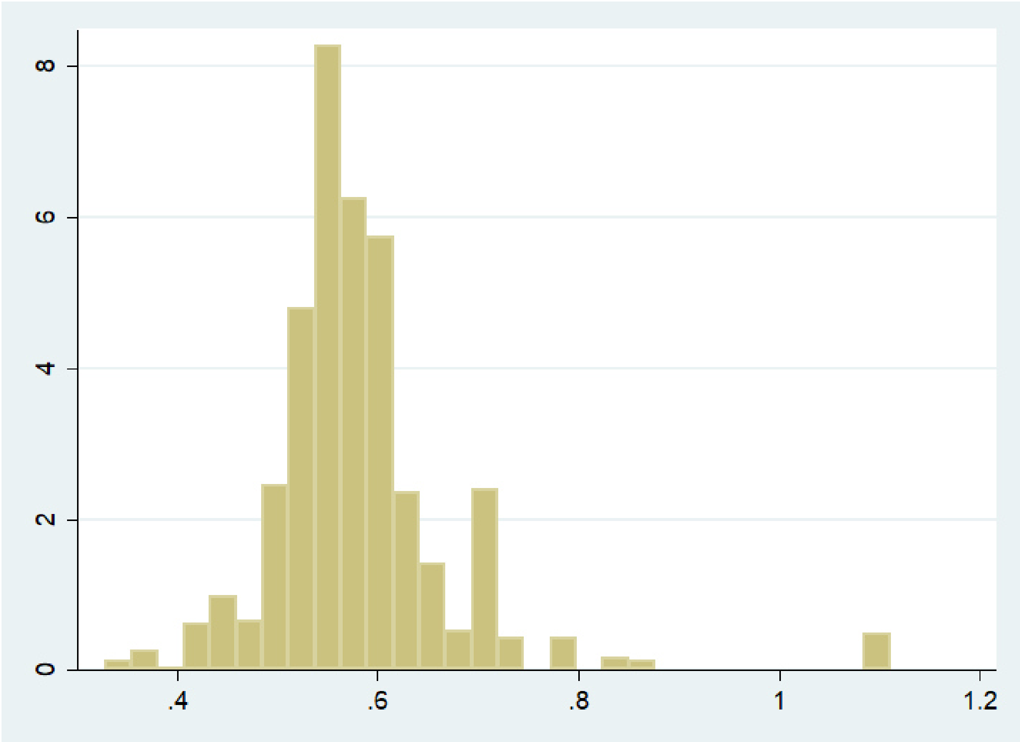

Figure 2.2: Histogram of $Smoothness_i$ for all sellers in the sample.

$$Smoothness_i = \frac{Demand_{i,Low}}{Demand_{i,High}} = \frac{\frac{1}{5}\sum_{k\in K_i} Demand_{k,Low}}{\frac{1}{7}\sum_{k\in K_i} Demand_{k,High}}$$

A demand $Smoothness_i$ value of 1 would indicate no fluctuation: seller i faces the same demand conditions during the $Lowseason_t$ and the $Highseason_t$. The lower is the value of $Smoothness_i$, the more the seller is exposed to the seasonality of demand. Most sellers experience a drop in demand during the low season, and the average value of smoothness is 0.59. A few sellers experience reverse seasonality with $Smoothness_i$ maximum of 1.59. *Figure 2.2* contains the histogram of $Smoothness_i$. The distribution is can be described as relatively bell-shaped with a longer tail on the right.

2.3.2 Main specification and results

We test the hypothesis that seller demand heterogeneity affects seller incentives to exert effort, which in turn affects the probability that a job is done well and the seller receiving a positive review. This effect should be especially strong at the end of the high demand season, when the discounted sum of future demand is the lowest.

To this end, we work with a sample of completed jobs, indexed by the hired seller i, the date defined as month-year t, and the job characteristics k. The outcome variable of interest is the probability that the seller i receives a positive review for job k in period t, $Pr(Positive\ review)_{ikt}$. The variable $High\ season\ end_t$ which indicates if the job was posted in the two months of the high demand season (October and November) and it is the common trend among all sellers. We interact it with seller $Smoothness_i$ to test whether sellers heterogeneous demand fluctuation affect the probability of $Pr(Positive\ review)_{ikt}$.

Specification (1) is:

$$
\begin{aligned}
Pr(Positive\ review)_{ikt} = \quad & \beta_1 High\ season\ end_t + \beta_2 Smoothness_i * High\ season\ end_t \\
& + \gamma X_{it} + FE_i + FE_k + FE_t + \epsilon_{ikt}
\end{aligned}
$$

The coefficient of interest is β_2. If it is positive, sellers with lower demand fluctuations (a higher $Smoothness_i$) are more likely to receive a positive review at the end of the high season. Or, the more the seller is exposed to seasonal demand fluctuations, the less likely he is to receive a positive review at the end of the high season.

As controls, we include covariates measuring seller experience and a number of fixed effects. The variable X_{it} is the *Percent positive reviews*. We also include fixed effects for each seller i, for the job category, budget and start date of the job k, and for the time period t. Since $Smoothness_i$ is constant over the sellers, it is absorbed by the seller fixed effects.

In specification (2), we allow for the complementary effect by including $Low\ season\ end_t$ dummy equal to 1 for the months May and June and its interaction with $Smoothness_i$. We expect coefficient β_4 to be negative, suggesting sellers who do not experience significant demand fluctuations are is less likely to make exceptional effort to improve their reputation before the beginning of the high season.

Specification (2) is:

$$
\begin{aligned}
Pr(Positive\ review)_{ikt} = \quad & \beta_1 High\ season\ end_t + \beta_2 Smoothness_i * High\ season\ end_t \\
+ \quad & \beta_3 Low\ season\ end_t + \beta_4 Smoothness_i * Low\ season\ end_t \\
& + \gamma X_{it} + FE_i + FE_k + FE_t + \epsilon_{ikt}
\end{aligned}
$$

We employ a linear probability model to make the interpretation of the estimated effects easier. Our estimates of (1) and (2) can be found in *Table 2.5*.[11] We find that the $High\ season\ end_t$ is associated with a general drop in the probability of a positive review. This effect is countered by demand $Smoothness_i$, as the estimated coefficient on the interaction is positive. For the median seller with $Smoothness_i = 0.57$, the probability of receiving a positive review at the end of the high season is 4 percent lower: $-0.99 + 0.57 \times 1.66 = -0.04$. sellers at the 10th percentile see a drop of 16 percent, while those at the 90th percentile actually see an increase of 18 percent in the probability of positive review at the end of the high season.

The results of specification (2) are in line with specification (1), but the additional variables are not significant. It is possible that this is because of the small sample size. The $Low\ season\ end_t$ is associated with a higher likelihood of a positive review, and sellers with higher demand $Smoothness_i$ are less likely to receive a positive review at the end of the low demand season.

Conditional on the seller being hired, we find that the variable *Percent positive reviews* is associated negatively with the probability of receiving another positive review. A potential explanation is that sellers are strategic in building and using up their reputation. Strategic reputation building is indirect evidence of moral hazard, and moral hazard as predominant source of asymmetric information has been recorded in online environments such as eBay by Klein, Lambetz and Stahl (2016).

2.3.3 Robustness checks

We consider the robustness of our results to alternative channels of causality: selection between sellers and jobs, adverse selection, and cost of effort.

Selection between sellers and jobs

As already mentioned in **Section 2.2**, the buyers do not observe the heterogeneous sellerl demand fluctuations, hence the opportunities for selection on the part of the buyers is limited. However, it is possible that sellers of different $Smoothness_i$ self-select into different jobs k, at different times t, and this interferes with our results.

[11] The results of probit and logit specifications are consistent and available upon request.

Table 2.5: Linear probability model for receiving a positive review.

Dependent variable	Positive review	Positive review
Specification	(1)	(2)
Independent variables		
High season end	-0.986**	-0.919**
	(0.439)	(0.446)
Smoothness*High season end	1.659**	1.532**
	(0.653)	(0.666)
Low season end		1.057
		(0.860)
Smoothness*Low season end		-2.319
		(1.448)
Percent positive reviews	-0.272***	-0.274***
	(0.065)	(0.064)
Constant	1.381***	1.396***
	(0.265)	(0.264)
Fixed effects		
Job category, budget, start	Yes, Yes, Yes	Yes, Yes, Yes
Date	Yes	Yes
Seller	Yes	Yes
R^2	0.41	0.41
N. observations	1,126	1,126

Significant at: $p < 0.1$: *; $p < 0.05$: **; $p < 0.01$: ***. Robust standard errors. The continuous regressors (except *Percent positive reviews*) are transformed by taking their natural logarithm, and their coefficients can be interpreted as elasticities. The coefficient on *Percent positive reviews* is a semi-elasticity.

Consider the indicator $Seasonal\ category_k$ equal to 1 for jobs with high demand incidence in June to November.[12] It may be that jobs in seasonal categories are more difficult to complete successfully. It is also possible that their successful completion is affected by the time period, for example: seasonal jobs may suffer from worsening weather conditions at the end of the high demand season. In addition, sellers in more seasonal lines of work (and therefore lower demand $Smoothness_i$) may be more likely to take on seasonal jobs, especially at the end of the high demand season. These effects present an alternative relationship between the probability of a positive review and the explanatory variables of interest: job seasonality affects the effort level needed for a positive review, and low demand $Smoothness_i$ sellers select into seasonal jobs towards the end of the high demand season. To test this, we augment (1) by interacting the $High\ season\ end_t$ and $Smoothness_i$ variables individually and jointly with the $Seasonal\ category_i$ dummy.[13] Our first robustness specification $(R1)$ is:

$$
\begin{aligned}
Pr(Positive\ review)_{ikt} = \quad & \beta_1 High\ season\ end_t + \beta_2 Seasonal\ category_k \\
& + \beta_3 Smoothness_i * High\ season\ end_t \\
& + \beta_4 Seasonal\ category_k * Smoothness_i \\
& + \beta_5 Seasonal\ category_k * Smoothness_i * High\ season\ end_t \\
& + \gamma X_{it} + FE_i + FE_k + FE_t + \epsilon_{ikt}
\end{aligned}
$$

The results of $(R1)$ can be found in *Table 2.6*. Jobs in seasonal categories are not more or less likely to get a positive review, neither in general nor only at the end of the season. The probability of a positive review for a seasonal job is also not affected by the demand $Smoothness_i$ of the hired seller, neither overall nor at the end of the high demand season. The main effects of interest - the $Highseasonend_t$ and its interaction with $Smoothness_i$ - remain significant and of similar higher magnitude.

[12] To classify the jobs, we regress monthly job incidence on the $Highseason_t$ dummy. These regression are available upon request.

[13] Individual category fixed effects are no longer included because of collinearity with the *Seasonal category* dummy.

Table 2.6: Robustness check with seasonal category fixed effect and interactions.

Dependent variable	Positive review	Positive review
Specification	(1)	(R1)
Independent variables		
High season end	-0.992**	-1.507**
	(0.413)	(0.480)
Seasonal category		-0.037
		(1.111)
Smoothness*High season end	1.626**	2.333***
	(0.683)	(0.421)
Seasonal category*High season end		0.859
		(1.033)
Seasonal category*Smoothness		-0.058
		(1.885)
Seasonal category*Smoothness*High season end		-1.310
		(1.646)
Percent positive reviews	-0.220***	-0.211***
	(0.063)	(0.061)
Total times hired	-0.030	-0.030
	(0.037)	(0.036)
Constant	1.360***	0.942
	(0.249)	(4.857)
Fixed effects		
Job category, budget, start	Yes, Yes, Yes	No, Yes, Yes
Date	Yes	Yes
Seller	Yes	Yes
R^2	0.41	0.39
N. observations	1,126	1,126

Significant at: $p < 0.1$: *; $p < 0.05$: **; $p < 0.01$: ***. Robust standard errors. The continuous regressors (except *Percent positive reviews*) are transformed by taking their natural logarithm, and their coefficients can be interpreted as elasticities. The coefficient on *Percent positive reviews* is a semi-elasticity.

Adverse selection

In the period for which we have data, the platform is relatively young and some sellers in our sample become inactive after a while. For these sellers, the incentive to exert effort in the last months of their subscription is even lower, as they may be ”cashing in” their reputation. If the end of their use of the platform coincides with the end of the high demand period, it may be that adverse selection is partially driving our results.

We restrict our sample to only those sellers who are still active on the platform in 2015 and 2016, which brings down our observations to 867 from 1,126. The results from estimating our model on the restricted sample, $(R2)$, are presented in *Table 2.7*; all coefficients on the variables of interest are very similar to those in our main specification. This is not surprising because the mean of demand $Smoothness_i$ of the two groups, the sellers who continue to be active and those who drop out, is not statistically different.

Table 2.7: Robustness check with sellers still active in 2015 and 2016.

Dependent variable	Positive review	Positive review
Specification	(1)	(R2)
Independent variables		
High season end	-0.992**	-1.035***
	(0.413)	(0.396)
Smoothness*High season end	1.626**	1.685***
	(0.681)	(0.649)
Percent positive reviews	-0.220***	-0.230***
	(0.062)	(0.067)
Total times hired	-0.029	-0.024
	(0.037)	(0.039)
Constant	1.360***	1.265***
	(0.249)	(0.245)
Fixed effects		
Job category, budget, start	Yes, Yes, Yes	Yes, Yes, Yes
Time (year)	Yes	Yes
seller	Yes	Yes
R^2	0.37	0.31
N. observations	1,126	867

Significant at: p < 0.1: *; p < 0.05: **; p < 0.01: ***. Robust standard errors. The continuous regressors (except *Percent positive reviews*) are transformed by taking their natural logarithm, and their coefficients can be interpreted as elasticities. The coefficient on *Percent positive reviews* is a semi-elasticity.

Cost of effort

A potential confounding factor in our results is the cost of effort, which may also change when demand is high, and differentially so for for sellers with different demand smoothness. With no information on seller costs, we are unable investigate this directly. Instead, we propose two regressions providing indirect evidence that the cost of effort is the driving factor of our results.

Firstly, we conjecture that the cost of effort fluctuates with current demand, and not with the discounted sum of demand in future periods. In $(R3)$, we use the full set of bi-monthly indicators on their own and interacted with the *Smoothness* variable.

$$\begin{aligned} Pr(Positive\ review)_{ikt} = & \sum_{t=2}^{T=6} \beta_j Bimonthly\ dummy_t \\ & + \sum_{t=2}^{T=6} \alpha_j Smoothness_i * Bimonthly\ dummy_t \\ & + \gamma X_{it} + FE_i + FE_k + FE_t + \epsilon_{ikt} \end{aligned}$$

The estimation results of $(R3)$ indicate no evidence to support a cost of effort based on the level of current demand. Indeed, the probability of positive review is significantly lower only in October and November. Furthermore, we see no differential effect for sellers of different demand $Smoothness_i$.

Another way to look at the issue of fluctuating costs is to consider the relationship between the size of the seller's operations and $Smoothness_i$. If sellers with lower $Smoothness_i$ are also smaller, they may be relatively more overstretched during the high demand season. In this sense, our estimations would be suffering from omitted variable bias, where $Smoothness_i$ picks up the effect of the omitted variable measuring seller cost of effort. To remedy this, we use the mean budget size for all jobs for which the seller was available, $Average\ budget_i$, as a metric for seller size, and interact it with the $High\ season\ end_t$ dummy.

Table 2.8: Robustness check with bimonthly fixed effects.

Dependent variable	Positive review	Positive review
Specification	(1)	(R3)
Independent variables		
High season end	-0.992**	
	(0.413)	
Smoothness*High season end	1.162**	
	(0.486)	
Feb& March		1.000
		(1.230)
Apr&May		0.888
		(1.222)
Jul& Jun		-1.369
		(1.062)
Aug& Sep		-1.750
		(1.192)
Oct& Nov		-2.035*
		(1.050)
Feb& March*Smoothness		-1.701
		(2.103)
Apr& May*Smoothness		-1.531
		(2.096)
Jun& Jul*Smoothness		2.403
		(1.805)
Aug& Sep*Smoothness		2.782
		(2.026)
Oct& Nov*Smoothness		3.316*
		(1.763)
Fixed effects		
Job category, budget, start	Yes, Yes, Yes	Yes, Yes, Yes
Date	Yes	Yes
Seller	Yes	Yes
R^2	0.37	0.39
N. observations	1,126	1,126

Significant at: $p < 0.1$: *; $p < 0.05$: **; $p < 0.01$: ***. Robust standard errors. Constant, *Percent positive reviews* and *Total times hired* omitted from table for brevity.

$$
\begin{aligned}
Pr(Positive\ review)_{ikt} = \quad & \beta_1 High\ season\ end_t + \beta_2 Seasonal\ category_k \\
& + \beta_3 Smoothness_i * High\ season\ end_t \\
& + \beta_4 Average\ budget_i * High\ season\ end_t \\
& + \gamma X_{it} + FE_i + FE_k + FE_t + \epsilon_{ikt}
\end{aligned}
$$

Table 2.9: Robustness check with measures of seller size.

Dependent variable	Positive review	Positive review
Specification	(1)	(R4)
Independent variables		
High season end	-0.992**	-1.363
	(0.415)	(1.078)
Smoothness*High season end	1.626**	1.814**
	(0.683)	(0.705)
Percent positive reviews	-0.220***	-0.222***
	(0.063)	(0.063)
Total times hired	-0.029	-0.030
	(0.037)	(0.037)
Average budget*High season end		0.033
		(0.090)
Constant	1.360***	0.349
	(0.249)	(0.820)
Fixed effects		
Job category, budget, start	Yes, Yes, Yes	Yes, Yes, Yes
Date	Yes	Yes
Seller	Yes	Yes
R^2	0.37	0.37
N. observations	1,126	1,126

Significant at: $p < 0.1$: *; $p < 0.05$: **; $p < 0.01$: ***. Robust standard errors. The continuous regressors (except *Percent positive reviews*, total offers and average budget) are transformed by taking their natural logarithm, and their coefficients can be interpreted as elasticities. The coefficient on *Percent positive reviews* is a semi-elasticity.

The results of $(R4)$ are presented in *Table 2.9*. $Average\ budget_i$ interacted with $High\ season\ end_t$ is not statistically significant, which suggests that our results are not suffering from bias due to omitting seller cost of effort that is also correlated with seller demand heterogeneity.

2.4 Conclusion and outlook

Reputation systems are essential for online economic activity because they provide buyers with information and sellers with incentives. However, the approach to designing such system is uniform and without much consideration for the economic environment and seller incentive heterogeneity. The economic literature has only recently started to consider the question of optimal aggregation and display of review information, and the effect that would have on seller behavior. Our work aims to inform this question by documenting how heterogeneous seller demand fluctuations lead to inconsistent effort provision on an online services marketplace.

Apart from the general contribution of our results, there are specific implications for reputation mechanism design in environments with heterogeneous demand fluctuations. For example, firms in the tourism industry (restaurants, hotels, etc) are similarly susceptible to heterogeneous and fluctuating demand and therefore incentives. The majority of buyers in these markets, who are tourists, rely exclusively on online reputation when making choices. Buyers choosing between hotels in tourist and residential areas may not be fully aware of the different demand conditions, and therefore incentives, that these establishments face. Making certain reviews more prominent, either by increasing their weight or visibility, can provide buyers with more useful information and sellers with more consistent incentives.

2.5 References

Athey, Susan, and Michael Luca. Economists (and Economics) in Tech Companies. No. w25064. National Bureau of Economic Research, 2018.

Bajari, Patrick, and Ali Hortacsu. ”Economic insights from internet auctions.” Journal of Economic Literature 42.2 (2004): 457-486.

Bar-Isaac, Heski, and Steven Tadelis. ”Seller reputation.” Foundations and Trends® in Microeconomics 4.4 (2008): 273-351.

Bar-Isaac, Heski, and Joel Shapiro. ”Ratings quality over the business cycle.” Journal of Financial Economics 108.1 (2013): 62-78.

Cabral, Luis. ”Reputation on the Internet.” The Oxford handbook of the digital economy (2012): 343-354.

Cabral, Luis, and Ali Hortacsu. ”The dynamics of seller reputation: Evidence from eBay.” The Journal of Industrial Economics 58.1 (2010): 54-78.

Dai, Weijia, Ginger Z. Jin, Jungmin Lee and Michael Luca. Optimal aggregation of consumer ratings: an application to yelp. com. No. w18567. National Bureau of Economic Research, 2012.

Dellarocas, Chrysanthos, and Charles A. Wood. ”The sound of silence in online feedback: Estimating trading risks in the presence of reporting bias.” Management science 54.3 (2008): 460-476.

Fradkin, Andrey, Elena Grewal, and David Holtz. ”The determinants of online review informativeness: Evidence from field experiments on Airbnb.” (2018).

Horton, John, and Joseph Golden. ”Reputation inflation: Evidence from an online labor market.” Work. Pap., NYU 1 (2015).

Klein, Tobias J., Christian Lambertz, and Konrad O. Stahl. ”Market transparency, adverse selection, and moral hazard.” Journal of political economy 124.6 (2016): 1677-1713.

Nosko, Chris, and Steven Tadelis. The limits of reputation in platform markets: An empirical analysis and field experiment. No. w20830. National Bureau of Economic Research, 2015.

Tadelis, Steven. ”Reputation and feedback systems in online platform markets.” Annual Review of Economics 8 (2016): 321-340.

Chapter 3

Airline Cooperation Effects on Airfare Distribution: An Auction Approach

Abstract

Assessing airline cooperation through its impact on aggregated prices could be challenging due to potential issues of heteroscedasticity and heterogeneity. To properly address this question, we represent air ticket sales by a reverse auction using individual transaction prices. Applicable to other industries where sellers compete in prices and when only price data are available, this approach allows us to reconsider the effect of airline alliances on the distribution of airfares in the US domestic market: We find a statistically significant lower price mean and higher price dispersion in markets where airlines belong to an alliance. Hence, we hope our methodology would help competition authorities to refine their competitive analysis of cooperative agreements among airlines but also in other industries.

3.1 Introduction

Airline cooperation plays a unique and crucial role in the industrial organization of international and domestic air travel. Most airlines cooperate in some manner, varying from a codeshare agreement on a particular market where airlines not operating a flight are allowed to sell tickets on that flight, to the more integrated arrangement of an alliance. There is no clear agreement among economists on the magnitude or/and the intensity of the effect of cooperation among airlines on prices or welfare. Because airlines face consumers with large heterogeneity in their tastes in terms of service features or preferred purchasing date, they implement complex pricing practices to better extract consumer's surplus, which blurs studying the impact of cooperation on prices. To address this question, the analysis usually proposed in the literature is somehow incomplete as it mainly focuses on how aggregated prices are affected by cooperation. However, as aggregated figures ignore the consumer heterogeneity, their use can create heteroskedasticy which can affect the precision of the measure of impact of cooperation, an issue which is well recognized in most of the literature on airlines.[1]

Here, to contribute to the empirical assessment of airline cooperation's impact on prices, we propose an original approach that allows the use of individual transaction prices to estimate the airfare distribution rather than studying aggregated figures. We argue that the ticket sales process is adequately approximated by price competition, and this allows us to ground our theoretical model on the equivalence between Bertrand competition and the reverse English auction, a result well known in the literature. In this setup, there is one buyer (the passenger or auctioneer) and multiple sellers (the airlines or bidders) who compete to sell their service by proposing prices; fare are observable by all the sellers who can modify their prices according to their competitors' offers.[2] Then, the transaction price at which the passenger buys the ticket is equal to the highest reservation cost among the competitors, where the reservation cost is the minimum acceptable compensation for the airline.[3] Hence, this result allows us to interpret the observed airfares as winning bids and to analyze their distribution by methods pertaining to the econometrics of auctions.

---

[1] For instance, Blundell and Stoker (2005)'s survey presents techniques that allows to attenuate the heterogeneity problems created when using aggregated figures in specific scenarios.

[2] Note that Klemperer (2004) states that theoretically such a "process corresponds exactly to the standard ascending auction among bidders competing to buy an object." He therefore refers to "ascending auctions" even for reverse auctions. We prefer to use the term "reverse auction" which is more coherent with our context.

[3] We use the term of reservation cost (instead of simply cost) to emphasize that we consider both operating and opportunity costs as explained in **Section 3.3**.

The empirical auction literature has developed several methods for the estimation of auction outcomes.[4] Our contribution relies on applying for the first time such methods to describe an internet sale process. In the recent years, we have seen a spectacular increase in the number and popularity of search engines that compare prices among different firms and in the deployment of yield management techniques.[5] More and more sectors see their goods and services traded online, with near-zero cost of price comparison and under heightened price competition. Some examples are car rentals, hotels, trains or more generally any market where firms offer similar products or services and compete exclusively on prices based on different reservation costs, a set-up corresponding to a reverse English auction. To present this new approach in a tractable manner, we focus thereafter on symmetric duopoly markets, i.e., markets with two companies that share a similar cost structure and similar product characteristics such as frequencies.

We apply our methodology to revisit the literature studying airline alliance effects on prices, where we make three important contributions. First, we directly work with individual prices while traditionally, the impact of alliances -or cooperative agreements more generally- is estimated in terms of average prices, aggregated over passengers, per airline, per market and per period. Second, our approach allows for a more comprehensive treatment of the price distribution by jointly modeling airfares' mean and the variance. Finally, we estimate the impact of alliances on the variability of ticket prices, which has not been considered before, neither by the literature on airline cooperation, nor by the literature on the effect of competition on airfare dispersion.

Alliances are partnership agreements between two or more competing firms. There exist a wide range of such agreements in the different sectors of the economy, see for instance the review by Kang and Sakai (2000) on international alliances; our work is focused on airlines alliances. The latter allow the carriers to cooperate, while maintaining certain boundaries and not constituting a merger. Most of the practices that alliance partners can engage in are considered beneficial for consumers: Alliance partners can market their partners' tickets and collaborate in supplying a product ($codeshare$), offering a larger network reach (foreign carriers usually cannot operate within the domestic market known as cabotage); they can coordinate their schedules, improving the service; and they can share frequent flyer programs and promotional campaigns. The alliance may lead to lower costs due to economies of density, because partners share airport equipment and staff.

---

4 For a recent survey, see Gentry et al (2018).

5 $Yield$ or *revenue management* is a variable pricing strategy that allows firms to increase revenues in an environment with fixed capacity that has an expiration date (for instance, the takeoff of a plane) and uncertain demand.

Despite the listed benefits, the impact of alliances over consumers in terms of prices is still open to discussion. There is general agreement that airline alliances can reduce prices for international services, as suggested by Park (1997), Brueckner and Whalen (2000) or Brueckner (2001). Most of the proposed products on international air markets, namely, the connecting flights, combine the services of at least two carriers. For instance, to travel from city A in one country to city C in another country, a stop is required in city B, with the routes AB and BC operated by two different carriers. To the benefit of passengers, their alliance can eliminate the double marginalization problem that appears when each of the carriers prices its service independently from the other. Now, on markets where partners in an alliance offer the same service, that is, on the so-called overlapping markets, the double marginalization problem does not exist and airfares may be higher because of the alliance if there are not enough competitors. As overlapping international markets represent a small percentage of the total number of markets, the social costs of higher prices are in this case largely compensated by the social benefits due to the removal of double marginalization on connecting flights. That is why international alliances are generally approved.

The situation used to be different for U.S. domestic alliances, where carriers are free to provide service between any two cities and their networks can overlap significantly.[6] The competitive effects of alliances in such markets caused concerns for the relevant authorities, one example being the Continental/Northwest/Delta alliance in 2002. The U.S. Department of Transportation (the DOT) argued that the process of communicating the necessary information to organize the codesharing service would facilitate carriers to collude explicitly or tacitly on prices and/or service in the overlapping markets. Despite these allegations, the Department of Justice allowed the formation of domestic alliances that eventually transformed into mergers, while their impact on airfares in the overlapping domestic markets was still uncertain.[7]

To reassess such decisions, we implement our methodology on the US domestic $direct$ and $connecting$ duopoly markets operated by legacy carriers during the third quarter of 2008. The service offered and the alliance status of the airlines define four market types. We label markets as direct when both airlines operate direct flights. Similarly, for connecting markets, both airlines must operate connecting flights, that is, flights connecting airport A to airport C with a stop by airport B. Moreover, markets can be classified as alliance or non-alliance markets. If the two airlines operating in a market belong to the

[6] The number of available slots and their allocation is regulated by the Department of Transportation at only a few airports due to traffic congestion.

[7] We discuss the relevant empirical literature in the next section.

same alliance, we denote it as an alliance market. These are the overlapping markets of the alliance partners. The market is non-alliance if the two airlines do not belong to the same alliance.

We show that, in the considered duopoly markets, airline alliances are associated with lower fares and higher price dispersion, the latter being measured by the standard deviation of the price distribution and the associated coefficient of variation. Prices are lower in alliance markets compared to non-alliance markets; more precisely, prices are 21 percent lower in direct markets and 4 percent lower in connecting markets. This finding suggests that alliance agreements lead to efficiency gains that are passed on to consumers. Now, price dispersion is higher in alliance markets compared to non-alliance markets: It is 13 (9) percent larger in direct (respectively, connecting) markets. The higher price dispersion in alliance markets is potentially a consequence of more efficient price discrimination techniques as allied airlines may share consumer's information or coordinate in other ways.

It is not clear if the combination of the two obtained effects is welfare improving. A reduction of the price mean is considered as welfare enhancing for a given quality level. However, the effects of price dispersion depend on the dispersion source. Since the work by Stigler (1961), there is a general agreement that price dispersion created by consumer search costs is welfare decreasing. However, the origin of price dispersion can also be price discrimination.[8] A policy paper by Geradin and Petit (2005) argues that price discrimination is an important and complex competition concern that should be considered on a case-by-case basis. There is a large variety of price discriminatory practices, which makes it arduous to classify them. A priori, their effect on total or consumer welfare is not clear since it depends on factors such as the market size, the demand shape, the firm's cost structure or the offered quality.[9] Here, thanks to our methodology, we point out that the competitive analysis of alliances, in particular by competition authorities, should not only focus on the effect of cooperation on the mean of the airfare distribution but should also consider the impact on its variance.

Given that this methodology is general, it can be applied to analyze the competition issues of alliances and cooperation in other industries when only price data are available.

[8] There exist other potential sources of price dispersion such as demand uncertainty, costly capacity or peak load pricing. See, for instance, Gale and Holmes (1993), Deneckere, Marvel, and Peck (1996) or Dana (1999), who show that price dispersion can arise as a result of other factors and not be linked to price discrimination.

[9] We direct the reader to Geradin and Petit (2005) or to Armstrong (2008) for a thorough discussion on the price discrimination theory and its effect on total and consumer welfare.

As shown by Kang and Sakai (2000), alliances have been widely implemented in the past. According to KPMG (2017), they are a valuable strategic opportunity for firms: "As critical drivers of growth, strategic alliances should be up there with M&A as a top priority for CEOs." Some examples from the car industry over the last 20 years are the General Motors-Fiat partnership, and the more recent Fiat-Renault and Daimler-Uber partnerships. (See KPMG, 2017.) In competition analysis, it is often the case that data on quantities are not available and a structural estimation using prices and quantities cannot be implemented. For instance, it is relatively easy to observe gasoline prices, but not quantities, at gas stations. Thus, our methodology can be a relevant tool to analyze the impact of a cooperation agreements between firms when only price data is available.

In the next section, we present the background to our work: our novel approach using the standard industry data set, and the literature on airline alliances and price dispersion. In **Section 3.3**, we introduce our theoretical model and the econometric specification. **Section 3.4** presents the data set and variables, and **Section 3.5** provides the empirical results. Lastly, **Section 3.6** concludes.

3.2 Background

In **Subsection 3.2.1**, we present the standard industry data set, the DB1B survey. We discuss how our proposed model of ticket sales allows a more comprehensive analysis compared to how the data has been used traditionally. In **Subsections 3.2.2** and **3.2.3**, we respectively discuss literatures on the effect of alliances on average prices and the effect of competition on price dispersion, as our methodology allows us to investigate both features of the price distribution simultaneously.

3.2.1 The DB1B dataset

The U.S. Department of Transportation (DOT) publishes a comprehensive price data source, the Airline Origin and Destination Survey (DB1B). This survey is a 10 percent sample of all airline tickets sold in the U.S. domestic market. It provides information on the price paid for each ticket sold (called below the *transaction price*) for a given market (or city pair) and for given product characteristics. The product characteristics are the attributes that distinguish different types of flights within the same market, namely, the operating airline, whether the flight is direct or connecting, and so on. Note that information about the purchasing date and the flight characteristics, such as the scheduled flight date and time, are not available. Due to this limitation, other databases aside from the DB1B, are sometimes considered in the airline literature.

Web data-scraping is one way to collect data on posted prices that includes the flight characteristics as well as the date and time at which prices were posted. The structural approach applied to data collected from online sources has great research potential for airline dynamic pricing. See, for instance, Escobari (2012), Lazarev (2013) and Williams (2013). The main limitation of this approach is that only posted prices are observed, but not the transaction prices and the number of transactions. Moreover, structural models using this kind of data are so far limited to the monopoly case because of the high complexity of modeling competition in a dynamic framework. The collection of online data may be blocked, as experienced by McAfee and Te Velde (2006) with American Airlines.

Computer reservation systems (CRS), such as Amadeus or Sabre, can provide information on actual transactions, not only on posted prices, including information on the purchasing date. However, only transactions that occur within the system are registered in this dataset. Information from some airlines may be missing in certain markets, with no clear way to model or reconstruct the missing data. The CRS data is usually sold at high prices to airlines and not generally accessible to researchers. As far as we know, the only exception is the work done by Sengupta and Wiggins (2012, 2014), Hernandez and Wiggins (2014) and Escobari and Hernandez (2015), who had access to one CRS for most of the carriers and domestic routes within US.

For these reasons, the DB1B remains the main source for analyzing different market and product features of the U.S. domestic industry, such as competition, mergers, collusion, entry of low-cost carriers (LCC), hub premium, or loyalty programs, as in Borenstein and Rose (1991), Brueckner and Spiller (1991), Miller (2010), Brueckner, Lee and Singer (2013), Berry, Carnall, and Spiller (1996), and Ciliberto and Williams (2010), respectively. These studies use the average market price or average product price as the dependent variable.

Traditionally, the airline literature studies average prices and price dispersion separately, while our study is the first to propose a joint analysis of the mean price and the price variability through a methodological contribution that allows us to work with individual transaction prices from the DB1B.

3.2.2 U.S. domestic alliances

The literature on domestic airline alliances exclusively uses the DB1B data set, and the outcome variable is the average (at the market or product level) transaction price. The alliance impact is typically measured by comparing the average prices before and after the alliance formation. Bamberger, Carlton and Neumann (2004) focus their analysis on the Continental/America West and Northwest/Alaska alliances; Armantier and Richard (2006) estimate the effect of the Continental/Northwest alliance; Gayle (2007) studies the formation of the Continental/Northwest/Delta alliance. While Bamberger, Carlton and Neumann's (2004) results suggest lower prices for alliance markets, the last two studies find the opposite result. All three studies find an increase in traffic volumes. The authors interpret their results as suggesting that alliance partners are successful at expanding their customer base and employing price discrimination strategies. They conclude that, while the airline alliance can lead to higher overall prices, the outcome is not necessarily collusive or universally welfare reducing for consumers.

To evaluate the overall effect of alliances on consumer surplus, Armantier and Richard (2008) propose a structural discrete choice model, which uses individual transaction prices as well as an auxiliary data set to circumvent the limitations of the DB1B. Their analysis demonstrates that, while consumers using direct flights do face higher prices, this is compensated by the overall improvement of service quality as a result of the alliance. This methodology is not as easily accessible as what we propose below, because it requires detailed data to supplement the DB1B.

Another strand of the literature focuses on the type of cooperation between alliance partners as a product feature. Ito and Lee (2007) distinguish between virtual codeshared products (where one partner operates the flight and the other sells the tickets on that flight) and $traditional$ codeshared products (where both partners are involved in the operation of the flight and both can sell tickets). They report that 85 percent of their sample are virtual codeshare products and they are in direct competition with the airline's own product in 70 percent of the markets. They conclude that alliance products are seen as inferior by consumers in comparison to pure $online$ flights (that is, flights operated and marketed by the same airline) and used by airlines to price discriminate between consumers with different willingness to pay. Gayle (2007) performs a similar exercise, but he focuses on the effect of the presence of traditional and virtual codeshare flights on the average market price; he finds that markets with traditional codesharing products have lower average prices, while markets with virtual codesharing have higher average prices.

While the literature attests that alliances (and more generally cooperation) are a relevant factor influencing prices, the estimated effects on average prices vary according to the employed methodology and the selected data subset. The model that we present in the next section updates this evidence regarding a more recent period in the history of alliances, while complementing the analysis of price means with that of price dispersion.

3.2.3 Price dispersion in the airline industry

A large branch of the literature on airline markets has analyzed price dispersion and how it is affected by different market features. As an outcome variable, these studies use aggregated measures of price dispersion and the majority focus on the effect of competition. Up to our knowledge, no study in this literature analyzes the impact of alliances or other cooperation forms. Alderighi (2010) compiles the main results in this literature, where the Gini index and the coefficient of variation are the most common measures of dispersion. In a seminal paper, Borenstein and Rose (1994) regress the Gini coefficient on factors related to costs. They exploit the difference in the number of carriers across markets to measure the effect of competition on price dispersion, and they find a positive effect. Gerardi and Shapiro (2009) pursue the same objective by implementing a before-after approach that uses fixed market effects to control for unobservable time invariant market characteristics. They find the opposite result – a negative effect of competition on price dispersion. Dai, Liu and Serfes (2013) find that the relationship between competition and price dispersion may be non-monotonic. Despite the methodological differences, the three studies used the DB1B database. Gaggero and Piga (2011) and Siegert and Ulbricht (2014) use web-scraping to collect posted price data for the European market. They find a negative correlation between competition and posted price dispersion, although the latter shows that this correlation is positive when price dispersion is measured at the market level rather than the flight level.

There is no clear consensus on the effect of competition on the variability of prices, or what measure of dispersion is most suitable. Bachis and Piga (2011), Mantin and Koo (2009) and Hernandez and Wiggins (2014) find that price dispersion decreases with the level of competition using other types of price dispersion measures. Gillen and Mantin (2009) and Sengupta and Wiggins (2014) find that competition does not generally affect price dispersion. Recently Chandra and Lederman (2018) find that the relationship depends on consumer heterogeneity and can be U-shaped.

We find it to be an important omission that none of the aforementioned studies analyze the impact of cooperation on price dispersion, despite the alliance and codesharing literature demonstrating that cooperation certainly has a significant effect on price means. Our study includes both competition and cooperation measures, and establishes a link between the literature on price dispersion and that on alliances.

3.3 A model of airline competition

In this section, we detail the assumptions that allow us to establish the observational equivalence of competition in the airline market with an auction model. We discuss the underlying determinants of costs; we outline the derivation of the maximum likelihood estimation (MLE); we describe how to estimate the distribution of prices and how cooperation, alliances in particular, affect it.

3.3.1 Overview

We propose a competitive framework aimed to depict appropriately the current economic environment faced by airlines. The recent trends in the industry, specifically service homogenization and high consumer price sensitivity, motivate our assumption that, in the short run, firms compete in prices given existing capacities.[10]

Assumption 1. The observed transaction price (airfare) is the result of price competition between airlines.

We conjecture that legacy airlines propose similar quality levels and compete by setting prices. Consider two airlines with different minimum prices at which they are willing to provide the service, what we call their *reservation cost*. The latter comprises the *operating cost* as well as the *opportunity cost* of the service. First, the operating cost covers the explicit costs to provide the service on a market.[11] This cost can vary within or across markets when we observe flights with different numbers of connections, flights with different distances or airlines with different economies of scale and scope. Second, the opportunity

[10] We leave aside entry and exit issues, which are studied by Berry (1992). However, we control for the potential bias that this could represent by sensitivity checks that include origin and destination level fixed effects.

[11] Operating costs are defined by the International Civil Aviation Organization to include aircraft or direct operating costs such as fuel, aircraft servicing costs such as handling, traffic service costs such as meals or flight attendance, booking and sales costs and other costs such as advertising or general administrative expenses.

cost is the value an airline places on selling a ticket now, relative to an uncertain sale of a ticket with a potentially higher price closer to the departure date. The opportunity cost varies over time within a market because it is affected by the purchasing patterns on that market, which in turn depends on the remaining capacity on a flight, the number of substitute flights and the expectation of future demand. Implicitly, the opportunity cost incorporates the ability of an airline to price discriminate over the course of ticket sales up to the flight takeoff.

The airline with the lower reservation cost has a competitive advantage – it can provide the service at a lower price than its competitor. The profit-maximizing strategy of this airline is to offer a price that is not unnecessarily low; it ”wins” the sale at the highest price that guarantees a sale. In other words, the most competitive airline makes a sale by offering a price that slightly undercuts the reservation cost of its competitor.[12] Under **Assumption 1** each ticket sale is viewed as a reverse English auction following the approach proposed by Klemperer (2004) for internet sales.[13]

The reservation cost of any airline, at any point in time, can be split into two parts with respect to its statistical nature and its relevance for the airlines. In the language of statistics, there is a deterministic component that is common and observed by all competitors. For example, the fuel cost to cover the distance between the ends of a market. There is also a random component that is private knowledge and has private relevance to the cost of an airline, for example, the idiosyncratic variation in the cost of fuel of each airline. This is why we consider that:

Assumption 2. The random component of reservation costs is an independent private value.

Furthermore, we treat the individual ticket sales as realizations of repeated, independent auction games between symmetric players.

[12] Our model is in line with the widely-spread yield management method of bid price control. In practice, there are several techniques that the airline can use to increase their revenue, some of which involve the estimation of a marginal cost of each seat on a plane, at each moment in time. One such method is *bid price control*, where this marginal cost is used as a bid – an optimum cut-off required to accept a booking. These bids correspond to reservation cost in our model, below which airlines are unwilling to sell tickets. Bid prices are dynamically adjusted over time to reflect the changing reservation cost under dynamic demand. This practice has been analyzed in the operations management literature by Talluri and Van Ryzin (1998) and Adelman (2007), among others.

[13] In a *reverse auction*, the auctioneer is a buyer and the participants are sellers who compete by offering prices (their bids) at which they are willing to provide the service. During an open auction of this kind, known as an *English auction*, competitors can observe each other’s bids (just as they do in our price competition set-up) and react to them.

Assumption 3. Ticket sales are realizations of independent, repeated auction games between symmetric airlines.

Under **Assumption 2** and **3**, the private random component of each airline in each sale is drawn anew from a probability distribution that is independent and identical across airlines and across sales.[14] This allows us to specify a methodology based on the Paarsch (1997)'s approach for estimating auction outcomes as we will explain in the next subsection. Our last assumption allows us to model price variability within a market using market characteristics. The DB1B prices exhibit significant variability driven by the unavailable flight characteristics and purchasing date, and the literature has treated this issue by averaging prices at the market level. We conjecture that flight characteristics and purchasing dynamics are endogenous to the market fundamentals. Unlike previous work, we propose to model this variability by making the following assumption:

Assumption 4. Market characteristics are determining factors of flight characteristics and advance purchasing dynamics, and thus, of price variability within a market.

For example, in a large metropolitan market we would expect multiple flights with different characteristics due to the large and diverse population; as a result, price variability would be higher because different flights have different operating costs. Higher or more diverse population income would affect the advance purchasing patterns. For instance, more last-minute business travels would cause the opportunity costs to vary more significantly. Thus, the market features, in terms of income and size, determine the unobserved flight features and the advance purchasing patterns over time. This rationale allows us to model the variability of airfares as a function of the market features.

3.3.2 The model

In this section, we pattern the equilibrium bidding strategy in a reverse English auction under the independent private values paradigm. We identify the players and describe each player's own information, available strategies and rewards; finally, we characterize the equilibrium. We consider exclusively duopolies for the sake of simplicity.[15]

[14] As auctions are repeated among players with capacity constraints, an ideal model should account for this dynamic; however, the DB1B dataset does not provide any information on the acquisition date which impedes analyzing such dynamics.

[15] Scenarios with more than two players imply asymmetry of the players, as usually only two of them will be in an alliance. Asymmetry would add significant complexity to the analysis without broadening the contribution of our methodology.

Suppose that a single buyer (namely, the consumer or the traveler) wishes to purchase one ticket in a market with two players (that is, the sellers or the airlines). Each player has a reservation cost to provide the ticket, which we denote c. The strategies available to the sellers are their bids (announced, posted or offered price) as a function of the reservation cost. The game proceeds as follows. The consumer only cares about prices and compares the airlines' offers.[16] The players fully observe, and can react to, each other's prices. Whenever it is profitable for them to do so, they can undercut the price of the competitor to win the sale. Each player is willing to lower one's price up to $p = c$, but not lower. The winner is the player with the lowest reservation cost who undercuts slightly the opponent with the highest reservation cost. *The resulting transaction price corresponds to the highest reservation cost among the two players.*

To estimate this model, we implement the MLE methodology for auction data proposed by Paarsch (1997). Given the equilibrium of the game, the observed transaction price is a function of the reservation cost: It is the highest reservation cost from two independent and identical reservation cost draws. Let us assume that the reservation cost has a cumulative distribution F and a respective density f in R. Moreover, we assume that the reservation cost follows a logarithmic normal distribution with mean and standard deviation , its natural logarithm being a normally distributed variable.[17]

Denoting by o (d) an airport at origin (destination, respectively), a market is defined as a directional city pair od.[18] The likelihood of an observation i in the market od is the occurrence probability that an airline sells ticket i at a price p_{iod} for a travel from o to d. For clarity of the exposition, the exact derivation of the likelihood, which is standard, is left for the **Appendix**. The likelihood of a single price observation is written as:

$$L_{iod} = 2F(p_{iod}|\mu_{iod}, \sigma_{iod})f(p_{iod}|\mu_{iod}, \sigma_{iod}) \tag{3.1}$$

The MLE approach yields estimates for the distribution parameters (μ_{od}, and σ_{od}) such that the resulting price distribution approximates the observed sample of individual ob-

[16] For example, consumers may use one of many and very popular websites offering search and comparison services, such as Kayak, Expedia, Orbitz and Travelocity, that allow consumers to enter their trip parameters and obtain a price ranking.

[17] The advantage of the logarithmic-normal distribution is that it allows us to interpret the coefficients of all the continuous explanatory variables as elasticities, as the explanatory variables themselves are transformed by taking their natural logarithm.

[18] The directional definition provides the basis for the delineation of relevant market in many studies. See Gayle (2007) and Berry and Jia (2010) or Luttmann (2018).

servations as close as possible.

The task is now to estimate the parametric effect of variables of interest (notably the presence of alliances) on the distribution of reservation costs. Through the distributional relationship between the reservation cost and the transaction price, we derive how these variables ultimately affect the transaction price distribution. To do so, we discuss in the next subsection how the distribution parameters are identified and estimated.

Model Specification

Let X_{od} be a vector of N variables relevant to the market od. These market factors affect the mean and standard deviation of the reservation cost distribution according to:

$$\mu_{od} = X_{od}^{'}\alpha \tag{3.2}$$

$$\sigma_{od} = X_{od}^{'}\beta \tag{3.3}$$

where α and β are the N-parameter vectors of the underlying reservation cost distribution to be estimated.[19]

Our main interest is how these market factors affect the actual transaction prices. As already discussed, the observed transaction price has a distribution that is a function of the underlying reservation cost, that is, the highest out of two reservation cost draws. Let us denote the corresponding mean and standard deviation of the price distribution as m_{od} and s_{od}, respectively. Then, following the derivation in Nadarajah and Kotz (2008), we can express the parameters of the price distribution in terms of those of the reservation cost distribution.[20] The mean of the transaction price, m_{od}, is a combination of μ_{od}, and σ_{od}, the mean and standard deviation of the reservation cost distribution. The price standard deviation, s_{od}, is simply the scaled standard deviation of the reservation cost. The marginal effects of the set of variables X_{od} on m_{od} and s_{od}, denoted below as a and

[19] The reservation cost comprises the operating cost as well as the opportunity cost of the service. We are not able to disentangle the impact of operating and opportunity cost over α and β due to data limitations. Modeling it would require information on the purchasing date and ticket characteristics of each transaction, which is not available in the DB1B.

[20] The exact forms of the different moments of the distribution of order statistics have been known for a while and are available in many good reference books such as David and Nagaraja (2003).

b, respectively, can then be simply calculated using the marginal effects α and β of the reservation cost distribution as:

$$m_{od} = \mu_{od} + \frac{\sigma_{od}}{\sqrt{\pi}} = X_{od}^{'}\Big(\alpha + \frac{\beta}{\sqrt{\pi}}\Big) = X_{od}^{'}a \tag{3.4}$$

$$s_{od} = \sigma_{od}\frac{\sqrt{\pi - 1}}{\sqrt{\pi}} = X_{od}^{'}\beta\frac{\sqrt{\pi - 1}}{\sqrt{\pi}} = X_{od}^{'}b \tag{3.5}$$

Looking at **Equation 3.4**, we can infer that the average transaction price is larger than the average reservation cost. However, we cannot conclude if the impact of the nth variable in X_{od} will be larger over transaction prices than over reservation costs. The ranking depends on the signs and relative magnitude of α_n and β_n coefficients. If they have the same sign, then the impact on the mean price, a_n, is larger than the impact on the reservation cost. If they have different signs, the overall effect depends on their relative magnitude and significance. It may be the case that both α_n and β_n are significant but of opposite sign, and a_n is insignificant.[21]

According to **Equation 3.5**, the variables X_{od} have a smaller impact on the price standard deviation, s_{od}, than on the reservation cost standard deviation, σ_{od}, since $\sqrt{\frac{\pi - 1}{\pi}} < 1$. The distribution of prices presents a lower standard deviation because observed prices are reservation costs that are selected in a "directional" way – we take the highest from two.

3.4 Data and explanatory variables

We use the DB1B data for the third quarter of 2008, and we select all markets satisfying three conditions. First, we exclusively consider duopolies where only two airlines operate. With more than two airlines, **Assumption 3** of symmetry is violated, as one of the airlines will not be part of the alliance, and this would complexify the exposition and methodology without a significant benefit. Second, all the proposed flights must be direct flights or all the proposed flights must be connecting, and we denote respectively the markets as direct and connecting. Again, this is because our methodology requires that the airlines propose

[21] Coefficients in vectors b and β share the same statistical significance. The significance of coefficients in vector a is calculated by representing them as a combination of two randomly distributed normal variables (α and β).

equivalent products and compete exclusively on prices as indicated in **Assumption 1**.[22] Furthermore, our analysis concerns only markets where major legacy carriers operate. We do not consider LCCs and markets where they operate. The cost structure of LCCs is different from that of legacy carriers, and moreover, they do not enter alliances, which makes them an unsuitable group to use for comparison.[23] American (AA), Alaska (AS), Continental (CO), Delta (DL), Midwest (YX), Northwest (NW), United (UA), and US Air (US) are the airlines that operate in our sample.[24] The presence of each airline in each market type according to the selected sample is found in *Table 3.1*.[25]

Table 3.1: Legacy carriers by market type: number of markets and observations

	Direct		Connecting		Total	
Carrier	Markets	Obs.	Markets	Obs.	Markets	Obs.
American	32	1,283	531	5,648	563	6,931
Alaska	2	116	51	829	53	945
Continental	1	38	257	2,637	258	2,675
Delta	12	708	623	5,746	635	6,454
Northwest	8	642	664	7,230	672	7,872
United Airlines	31	1,362	577	5,092	608	6,454
US Airways	8	303	539	6,870	547	7,173
Midwest	2	117	12	89	14	206
Total	48	4,569	1,627	34,141	1,675	38,710

Note:Number of markets and observations for the airlines in our sample. Each observation represents a ticket sale. Since markets are duopolies, the total number of markets is the sum of all markets where each airline operates, divided by two.

Explanatory variables

Our explanatory variables Xod include market, origin and destination characteristics and are obtained from the DB1B and the U.S. Census Bureau. These variables affect the operating and/or opportunity costs in a non-random manner. The complete list of variables and their definition is displayed in *Table 3.2*. *Distance* is a market-level variable that is

[22] This implies that we exclude duopoly markets where one airline offers connecting flights to attain d while the competitor proposes direct flights. Direct flights are perceived by customers as higher quality products than connecting flights.

[23] Regional-legacy carrier agreements are not considered to be alliances but rather an integrated service. As is standard in the literature, we recode tickets sold by regional carriers as the legacy partner.

[24] The legacy carrier definition here is the same as in Brueckner, Lee and Singer (2013).

[25] Additional information on the cleaning process of the data set is presented in the **Appendix**.

measured in number of miles between the origin and destination cities. The distance between two cities impacts the level of operating costs, as longer distances require more fuel to reach the destination. Distance can also affect the opportunity cost of a ticket, as it affects substitution with land travel.

Our demographic measures – *Population* and *Income* – are measured at the origin and destination cities. Higher income cities have both richer leisure travelers who do not need to plan in advance and more business travelers who book tickets in the last days before departure. We therefore expect high income to lead to a higher average and a lower variability for the reservation cost. Population, on the other hand, is a measure of market size and could be associated with lower operating costs, as larger scale operations are more efficient. However, the effect of population over the reservation cost is not clear since a larger population can also imply higher opportunity costs, as more buyers are expected to arrive closer to the departure date.

Table 3.2: List of variables and their definitions.

Variable	Description
Distance	The distance between the two city end-points of the market measured in miles.
Origin population	The origin city's population.
Destination population	The destination city's population.
Origin income	The origin city's income (GDP) in US dollars.
Destination income	The destination city's income (GDP) in US dollars.
Market volume	Number of tickets sold on the directional market.
Origin volume	Number of tickets sold for any market originating at the origin.
Destination volume	Number of tickets sold for any market originating at the destination.
Origin markets	Number of markets accessible from the origin.
Destination markets	Number of markets accessible from the destination.
Origin hub	Dummy equal to 1 if the origin city is a Hub.
Destination hub	Dummy equal to 1 if the destination city is a Hub.
HHI	Herfindahl-Hirschman Index.
Alliance	Dummy equal to 1 if the two carriers operating on the market are in an alliance.

The *Market volume* variable is a measure of capacity or scale on the market; it is the total number of ticket sales observed in that market. A higher volume is a combination of higher number of flights and/or larger plane size, which would allow the airlines to have lower average operating costs. At the same time, each additional flight is a substitute for the other flights the airline offers, and a larger number of flights may lead to lower opportunity cost variability.

We construct four other variables that describe the market in relation to the network. *Origin volume* and *Destination volume* measure the total number of ticket sales at the origin and destination of the market, respectively, to passengers traveling to any point in the network. This is an alternative "scale" measure to Market volume. To quantify the market's centrality in the network, we build two variables: *Origin markets* and *Destination markets*. Origin markets counts the number of cities directly accessible from the origin, while Destination markets counts the number of cities from which one can fly to the destination. The centrality in a network affects operating costs through scope economies and the alternative use of resources (planes, personnel) in adjacent markets.

As a further measure of the importance of the origin and destination in the network of the operating carrier, we include the dummy variables *Origin hub* and *Destination hub*. The large scale of operations at hub airports may reduce operating costs, but costs may also fluctuate more as the airline allocates aircraft capacity among connecting passengers from different origins and destinations.

We also include the *Herfindahl-Hirschman Index* (HHI) as a measure of the relative size of the market share of the two carriers that relates to the level of competition between them. Finally, we use a dummy variable to indicate the alliance presence, which we discuss in more detail in the next subsection.

Alliance presence

In our definition of an alliance, we follow Ito and Lee (2007). Carriers are alliance partners if passengers on one of the alliance carriers can earn elite-qualifying frequent flyer miles on flights marketed or operated by the other alliance partner and vice versa. The alliance presence is defined at the market level. The market can either be an alliance market ($Alliance_{od} = 1$) if both carriers are in the same alliance, or a non-alliance market ($Alliance_{od} = 0$) if the carriers are not in an alliance together. In this sense, we do not model explicitly the exact type of cooperation products or level of coordination (flight scheduling, sharing equipment and personnel, revenue sharing or else) that occurs within alliance markets. The complete list of alliances by type of market, the number of markets where they operate and the number of observations in the third quarter 2008 is provided in *Table 3.3*.

The literature on the impact of airline alliances has approached the estimation of their effect in two ways. Ito and Lee (2007) look at the prices of different types of alliance products within the same market. Gayle (2007) and Bamberger, Carlton and Neumann (2004), on the other hand, look at the effect of introducing an alliance product in a given

Table 3.3: Presence of alliances in each market type.

	Direct		Connecting	
Alliance	Markets	Obs.	Markets	Obs.
Alaska/Delta			36	890
Alaska/Northwest			1	17
Continental/Delta			22	334
Continental/Northwest			3	36
Delta/Northwest	4	529	74	1,249
Northwest/Midwest	2	346	12	247
US Airways/United Airlines	2	70	67	1,338
Total alliance	8	945	215	4,111
Non-alliance	40	3,624	1,412	30,030

Note: Type and number of markets and observations for each alliance pair.

market. In our methodology, we compare prices across markets (cross-sectionally), rather than before and after the agreement, to estimate how the presence of the agreement affects them. The cross-sectional comparison allows us to focus on more recent and relevant time periods compared to a before/after comparison, which must analyze data around the time of the alliance formation. The model is estimated separately for direct and connecting markets. As these markets have different demand and cost factors, we expect their reservation costs to come from different distributions.

Ex-ante, it is not obvious how the alliance presence could affect the level and variability of the reservation cost, as there are several potential effects working in opposite directions. On the one hand, alliances are allowed to share certain operating costs such as personnel and airport facilities, which could reduce operating costs. On the other hand, the ability to coordinate schedules and to sell tickets on a competitor's flights can make price discrimination more profitable, affecting the opportunity cost and how it evolves over time. The $Alliance$ variable thus measures the overall effect of the alliance on the reservation cost's mean and standard deviation.

An important assumption in our approach is that, after controlling for all observed variables, alliance markets must be comparable to non-alliance markets. In other words, there are no unobservable factors that make the alliance more profitable in the specific markets where the alliance operates. If this were not true, the estimated alliance effect would also contain the effect of the unobserved factors, hence it would be biased. To avoid this problem, Brueckner (2003) uses a model with entry. Another more direct approach that is used by Brueckner and Whalen (2000) and Ito and Lee (2007) is to introduce fixed effects, which we consider in **Subsection 3.5.2**.

Summary statistics

Table 3.4 presents the summary statistics for direct and connecting markets distinguishing between alliance and non-alliance markets. As expected, the distance in direct markets is much shorter compared to that in connecting markets. Still, note that transaction prices in these two types of markets are on average very similar. The summary statistics indicate that, the direct and connecting markets appear to be quite different in their characteristics, and therefore, we expect that the effect of these characteristics on prices to be different. Overall, there do not seem to be significant differences in market characteristics between alliance and non-alliance direct or connecting markets. The only exception is that non-alliance markets have larger populations and higher traffic volumes in both the direct and connecting markets. Although the variables are presented here in levels, for the estimation they are transformed by taking their natural logarithm.[26] This transformation allows us to interpret the coefficients of all continuous variables as elasticities. In other words, each estimated coefficient represents the percentage change in the mean μ_{od} or the standard deviation σ_{od} of the reservation cost given one percent change in the relevant variable. The $Alliance$ coefficient, however, is a dummy, and its interpretation is slightly different; we multiply the estimated coefficient by 100 to obtain the percentage change of the mean or standard deviation of the reservation cost when $Alliance = 1$. This interpretation is also relevant for the price mean and standard deviation, m_{od} and s_{od}, respectively. The effects of the explanatory variables on m_{od} and s_{od} are derived using equations (3.4) and (3.5).

3.5 Empirical results

In **Subsection 3.5.1**, we present the estimation results from equations $(3.2) - (3.5)$ using the full set of covariates measuring the market characteristics. Our main results are presented in **Subsection 3.5.2**, where we control for unobserved factors not included among our explanatory variables by re-estimating our model with fixed effects. These results indicate that omitted variable bias is a valid concern and that fixed effect estimation is preferred. Finally, in **Subsection 3.5.3** we present the results for the coefficient of varia-

[26] This is common in the airline literature. See, for example, Ito and Lee (2007).

Table 3.4: Average values by market type and alliance presence.

	Direct		Connecting	
	Alliance	Non-Alliance	Alliance	Non-Alliance
Price (USD)	607	514	558	524
Distance (Miles)	863	828	2,585	2,366
Origin population	2,064,981	4,787,705	1,431,042	2,297,882
Destination population	2,291,229	3,952,544	1,419,460	2,183,883
Origin income (USD)	59,721	56,373	53,693	53,307
Destination income (USD)	60,444	55,851	55,011	53,978
Market volume	189	141	28	34
Origin volume	7,385,647	17,100,000	3,542,906	6,036,518
Destination volume	7,998,992	14,100,000	4,700,607	5,987,512
Origin markets	95	86	26	31
Destination markets	98	73	27	32
Origin hub	1	0.75	0.21	0.25
Destination hub	1	0.62	0.25	0.26
HHI	0.57	0.58	0.63	0.67

tion.

Estimation with market covariates

Table 3.5 contains the results from the estimation of equations (3.2) and (3.3) on two separate sub-samples, direct and connecting markets.[27] First, the $Alliance$ variable has no significant effect on the mean reservation cost in direct markets, and it has a positive effect of 1.2 percent in connecting markets. Second, there is a significant and positive effect of 5.4 percent on the standard deviation in direct markets and 2.1 percent in connecting markets.

The effects of other covariates are less consistent between market types, and we believe this is due to the different underlying market characteristics and demand. For instance, distance has a positive impact on the mean of the reservation costs for both direct and connecting markets, which is due to the cost of fuel and other operating expenses. At the same time, higher distance is associated with more varied reservation costs in direct markets but less varied reservation costs in connecting markets. This finding can be explained considering the option to substitute air travel with land travel in direct markets: Reserva-

[27] A likelihood ratio test indicates that the jointly estimated model is strongly rejected in favor of the separate estimations.

tion costs are lower for early tourist demand that is sensitive to surface transportation competition and higher for last minute business demand for which surface travel is not an attractive substitute. Substitution with land travel is less feasible in connecting markets. It is likely that longer connecting trips are done predominantly by tourists and planned much in advance, hence on average there is less last-minute demand, and reservation costs are more stable.

Table 3.6 presents the effects of the market factors on the transaction price, calculated using equations (3.4) and (3.5). As previously noted, the estimated parameters of the standard deviation of the distribution of transaction prices are always lower compared to those of the distribution of the reservation costs. Intuitively, the transaction price has less variability because the competition between the carriers leads to the highest reservation cost being the transaction price. The mean coefficients, however, are a combination of both the mean and standard deviation coefficients of the distribution of the reservation costs, and thus, the overall effect depends on their respective significance, signs and relative magnitude. The resulting effect of $Alliance$ on the mean of the distribution of transaction prices is still not significant in direct markets and it has a positive 2.4 percent effect in connecting markets. The effect on the standard deviation is 4.4 percent in direct markets and 1.8 percent in connecting markets. We observe a negative and significant effect of competition, as measured by HHI, on transaction price dispersion in connecting markets. These results are in line with the results obtained by Gerardi and Shapiro (2009), who find that price dispersion decreases with competition.

Estimation with fixed effects

In this subsection, we propose to control for unobserved factors that may correlate with $Alliance$ and could bias its effect by re-estimating our model with origin and destination fixed effects. However, in direct markets the $Alliance$ variable is collinear with some fixed effects and cannot be separately identified. This is due to the small number of direct markets. We solve this issue by including origin fixed effects only and keeping the covariates measured at the destination level. We do not have the same issue in connecting markets, and we include both origin and destination fixed effects in that estimation. The estimated effects of the $Alliance$ variable on the reservation cost and the transaction price are reported in *Table 3.7*.

The results indicate that omitted variable bias is a valid concern. $Alliance$ now has a significant and negative effect on the mean of the reservation cost distribution in both

Table 3.5: Reservation cost distribution: mean and standard deviation

	Direct		Connecting	
	μ_{od}	σ_{od}	μ_{od}	σ_{od}
Alliance	-0.011	0.054***	0.012**	0.021***
	(0.022)	(0.013)	(0.005)	(0.003)
HHI	-0.130**	0.028	-0.071***	0.040***
	(0.054)	(0.033)	(0.009)	(0.005)
Distance	0.172***	0.038*	0.278***	-0.029***
	(0.021)	(0.013)	(0.004)	(0.002)
Origin population	0.073***	-0.01	-0.007*	-0.006*
	(0.013)	(0.008)	(0.003)	(0.002)
Dest. Population	0.092***	0.0	-0.067***	-0.009***
	(0.013)	(0.008)	(0.003)	(0.002)
Origin income	0.499***	-0.099	-0.043***	-0.016*
	(0.114)	(0.065)	(0.012)	(0.007)
Destination income	0.240*	-0.160*	0.069***	-0.002
	(0.104)	(0.056)	(0.012)	(0.007)
Market volume	0.107***	-0.023*	-0.056***	-0.008***
	(0.014)	(0.008)	(0.003)	(0.002)
Origin volume	-0.166***	0.013	0.023***	-0.007***
	(0.017)	(0.010)	(0.002)	(0.001)
Destination volume	-0.175***	-0.007	0.032***	0.007***
	(0.016)	(0.010)	(0.002)	(0.001)
Origin markets	0.335***	0.036*	0.011*	0.024***
	(0.030)	(0.018)	(0.004)	(0.003)
Destination markets	0.386***	0.056*	0.023***	-0.010***
	(0.034)	(0.020)	(0.004)	(0.003)
Origin hub	-0.100*	-0.095*	0.001	0.013***
	(0.051)	(0.030)	(0.006)	(0.004)
Destination hub	-0.240***	-0.073*	0.064***	0.011*
	(0.049)	(0.027)	(0.006)	(0.004)
Constant	-3.933*	2.976*	3.872***	0.961***
	(1.724)	(0.946)	(0.194)	(0.114)
Likelihood	-9,166	-9,166		-47,503
Number of observations	4,569	4,569		34,141

Note: Standard error are given in parentheses. The symbols */**/*** indicate statistical significance at the 5/1/0.1 % level, respectively.

Table 3.6: Transaction price distribution: mean and standard deviation

	Direct		Connecting	
	m_{od}	s_{od}	m_{od}	s_{od}
Alliance	0.02	0.044***	0.024***	0.018***
	(0.021)	(0.011)	(0.005)	(0.003)
HHI	-0.114*	0.023	-0.049*	0.033***
	(0.052)	(0.027)	(0.009)	(0.004)
Distance	0.193***	0.031*	0.261***	-0.024***
	(0.021)	(0.010)	(0.004)	(0.002)
Origin population	0.068***	-0.009	-0.060***	-0.006***
	(0.013)	(0.007)	(0.003)	(0.002)
Dest. population	0.092***	0.0	-0.010***	-0.005*
	(0.013)	(0.006)	(0.003)	(0.001)
Origin income	0.443***	-0.082	-0.072***	-0.007***
	(0.099)	(0.054)	(0.003)	(0.001)
Destination income	0.15	-0.132*	-0.052***	-0.013*
	(0.101)	(0.046)	(0.011)	(0.006)
Market volume	0.094***	-0.019*	0.068***	-0.002
	(0.014)	(0.007)	(0.011)	(0.006)
Origin volume	-0.159***	0.01	0.019***	-0.006***
	(0.015)	(0.008)	(0.002)	(0.001)
Destination volume	-0.179***	-0.006	0.036***	0.006***
	(0.015)	(0.008)	(0.002)	(0.001)
Origin markets	0.355***	0.030*	0.024***	0.020***
	(0.029)	(0.015)	(0.004)	(0.002)
Destination markets	0.417***	0.046*	0.018***	-0.008***
	(0.031)	(0.016)	(0.004)	(0.002)
Origin hub	-0.154***	-0.078*	0.008	0.011***
	(0.044)	(0.024)	(0.006)	(0.003)
Destination hub	-0.281***	-0.060*	0.070***	0.009*
	(0.045)	(0.022)	(0.006)	(0.003)
Constant	-2.254	2.457***	4.414***	0.793***
	(1.561)	(0.781)	(0.185)	(0.094)

Note: Standard error are given in parentheses. The symbols */**/*** indicate statistical significance at the 5/1/0.1 % level, respectively.

direct (-24.3 percent) and connecting markets (-5.9 percent). The effect of $Alliance$ on the standard deviation of the reservation cost is augmented to a positive 5.2 percent in direct markets and a positive 2.8 percent in connecting markets. The resulting effect on the mean of the transaction price distribution is negative and significant as follows: -21.4 percent in direct markets and -4.3 percent in connecting markets. The effect on the standard deviation of the transaction price distribution is positive and significant as follows: 4.3 percent in direct markets and 2.7 percent in connecting markets. Overall, our results indicate that alliances are associated with lower transaction prices and higher price dispersion in both direct and connecting markets, and the effect in direct markets is much stronger.

To ensure the airlines' symmetry assumption, we perform a sensitivity analysis on a restricted sample of symmetric markets, where the market share and flight frequencies are almost equal for both airlines. The negative and significant effects over price means in both direct and connecting markets remains (-16 percent and -4 percent, respectively, reported in *Table 3.7*). There is no significant effect over the standard deviation in connecting markets, but there is a positive effect of 6.9 percent in direct markets.

Table 3.7: The effect of $Alliance$ in different specifications.

			Fixed effects	Reservation cost		Transaction price	
Market	Sample	Origin	Dest.	μ_{od}	σ_{od}	m_{od}	s_{od}
Direct	Full	Yes	No	-0.243***	0.052*	-0.214***	0.043*
				(0.033)	(0.020)	(0.032)	(0.017)
Connecting	Full	Yes	Yes	-0.059***	0.028***	-0.043***	0.027***
				(0.007)	(0.005)	(0.007)	(0.004)
Direct	Sym.	Yes	No	-0.207***	0.084*	-0.160***	0.069***
				(0.049)	(0.027)	(0.046)	(0.022)
Connecting	Sym.	Yes	Yes	-0.041*	0.002	-0.040*	0.002
				(0.016)	(0.010)	(0.015)	(0.008)

Note: Standard error are given in in parentheses; Origin and Destination fixed effects are controlled for, but estimates have been suppressed in the table. Direct markets include covariates measured at the destination level. The symbols */**/*** indicate statistical significance at the 5/1/0.1 $\%$ level, respectively.

Coefficient of variation

As argued in the introduction, the coefficient of variation (CV) is a standard metric for price dispersion in the literature. To be consistent with previous work, we present here the effect of $Alliance$ on the CV of prices. The reservation cost and transaction price are transformed by the natural logarithm in our estimation; then, we construct the CV of transaction prices using properties of moment generating functions. The complete derivation of the CV is presented in the **Appendix**. Note that, in the end, the CV of prices is only a function of the standard deviation of the log-reservation cost, σ, specifically:[28]

$$CV = \left(\frac{exp(\sigma^2)\Phi(\sqrt{2}\sigma)}{2\left(\Phi\left(\frac{\sigma}{\sqrt{2}}\right)\right)^2} - 1 \right)^{1/2}$$

Since $Alliance$ is an indicator rather than a continuous variable, its marginal effect is obtained as the difference between the CV when $Alliance$ is one and zero, that is to say:

$$\Delta CV(Alliance) = CV(Alliance = 1) - CV(Alliance = 0)$$

For testing the significance of the alliance effect, we use the estimated standard deviation of the log-reservation cost evaluated at either $Alliance = 1$ and $Alliance = 0$ and at the sample mean values for all other covariates.

Table 3.8 contains the estimated CVs for direct and connecting markets, both for the model with covariates and the model with fixed effects. In direct markets, $Alliance$ is associated with an increase in CV of 0.05, or approximately 13 percent, regardless of whether the model is estimated with or without the fixed effects. In connecting markets, $Alliance$ is associated with an increase of the CV of 0.02, or approximately 7 percent, in the model with covariates, and 0.025, or approximately 9 percent, in the model with fixed effects.

As in the previous section, we perform a sensitivity check on our restricted sample of symmetric markets. *Table 3.8* presents the impact over the coefficient of variation, which is positive and significant with an increase in CV of 0.083 (approximately 21 percent) for direct markets and no effect on the connecting markets. This result is particularly relevant. The literature has not found a clear relationship between competition (usually measured with HHI) and price dispersion. In the restricted sample, we should not observe price dispersion differences due to competition since all the markets have two competitors with similar HHI levels. However, there is higher price dispersion for direct markets with

[28] In the formula, $exp(.)$ indicates the exponential function and $\Phi(.)$ is the cumulative distribution of the standard normal function.

Table 3.8: The effect of $Alliance$ on the coefficient of variation of transaction prices.

Model with covariates

	Sample	$CV(Alliance = 1)$	$CV(Alliance = 0)$	$\Delta CV(Alliance)$
Direct markets	Full	0.457	0.404	0.053***
Connecting markets	Full	0.315	0.295	0.020*
Direct markets	Symm.	0.447	0.394	0.083***
Connecting markets	Symm.	0.292	0.287	0.005

Model with fixed effects

	Sample	$CV(Alliance = 1)$	$CV(Alliance = 0)$	$\Delta CV(Alliance)$
Direct markets	Full	0.445	0.393	0.052***
Connecting markets	Full	0.295	0.27	0.025*
Direct markets	Symmetric	0.472	0.389	0.083***
Connecting markets	Symmetric	0.262	0.26	0.002

Note: Standard error are given in in parentheses; Origin and Destination fixed effects are controlled for, but estimates have been suppressed in the table. Direct markets include covariates measured at the destination level. The symbols */**/*** indicate statistical significance at the 5/1/0.1 % level, respectively.

an alliance, even after we control for all possible market effects. This result shows that cooperation among airlines affects price dispersion, which should not be omitted when analyzing the relationship between competition and price dispersion. In this sense our work is a new step contributing to the work by Liu and Serfes (2006) and Chandra and Lederman (2018) that tries to reconcile the conflicting results in the earlier literature.

3.6 Conclusion

Airline markets have gone through many transformative changes in the last couple of decades. Low-cost carriers expanded their services and made standard the "no frills" type of service, decreasing the importance of service quality and increasing the homogenization of the product. The spread of the Internet as a sales channel has been another challenge for the industry; consumers' search and comparison costs became negligible using online travel agents and price comparison sites. Furthermore, the recent economic crisis made all travelers, and particularly business travelers, very sensitive toward prices. We are motivated by these recent changes in the industry to propose to represent the ticket sales process by means of an auction model. This approach, which is the main novelty of this paper, is applicable to any online sales process based on prices.

We apply this model to revisit the analysis of airline alliances, a form of cooperation in airline markets that has caused much controversy. Our novel approach allows us to work with the individual data observations of the DB1B and to simultaneously explore the effect of alliances on price means and price variability, the latter being completely novel to the alliance literature. Our results indicate that alliances are associated with lower prices and higher price variability in both direct and connecting markets. These results are in contrast to previous results in the alliance literature by Gayle (2008) and Armantier and Richard (2006), which were relevant to the post-alliance formation period of the late 1990s and early 2000s. This difference in results could be well explained by the emerging competition from LCCs and by the fact that alliance partners were responding by passing their efficiency gains down to the consumers. Additionally, the higher standard deviation of prices suggests alliance partners have improved their price discrimination strategies. Although we are not able to distinguish whether the alliance effect is due to a synchronization of operating costs or a better management of opportunity costs, our analysis is a valuable stepping-stone in the direction of understanding the effect of cooperation agreements.

Our results shed a new light on the debate of the impact of competition on price dispersion. Indeed, the literature studying this issue presented in **Subsection 3.2.2** has not found a clear conclusion. For the first time, our analysis shows that markets with similar levels of competition present higher dispersion levels due to the presence of cooperation, in the form of alliances. The proposed methodology could be applied to analyze the impact over the price distribution of other cooperation forms. In the case of airlines, codesharing has been a frequent issue of concern for competition authorities. Diverse types of alliance or cooperation agreements are present in a wide range of industries, such as Financial services, Pharmaceuticals, Automobile or Software.[29]

Beyond the question of the impact of alliances in airline markets, we hope that our approach based on the econometrics of auction models will be applied to facilitate the analysis of any issue of interest in markets where competition is based on prices, when only price data are available and when the analyst is interested in both price levels and dispersion.

[29] Kang and Sakai (2000) present a review of international strategic alliances between 1990 and 2000 by sector and by country or region.

3.7 References

Adelman, Daniel. 2007. ‘Dynamic Bid Prices in Revenue Management’. Operations Research 55 (4): 647–61. https://doi.org/10.1287/opre.1060.0368.

Alderighi, Marco. 2010. ‘Fare Dispersions in Airline Markets: A Quantitative Assessment of Theoretical Explanations’. Journal of Air Transport Management, Selected papers from the 48th Congress of the European Regional Science Association, 16 (3): 144–50. https://doi.org/10.1016/j.jairtraman.2009.07.004.

Armantier, Olivier, and Oliver Richard. 2006. ‘Evidence on Pricing from the Continental Airlines and Northwest Airlines Code-Share Agreement’. Advances in Airline Economics 1: 91–108.

Armantier, Olivier, and Oliver Richard. 2008. ‘Domestic Airline Alliances and Consumer Welfare’. The RAND Journal of Economics 39 (3): 875–904.

Armstrong, Mark. 2008. ‘Price Discrimination’, Mimeo, Department of Economics, University College London, . http://discovery.ucl.ac.uk/14500/1/14500.pdf.

Bachis, Enrico, and Claudio A. Piga. 2011. ‘Low-Cost Airlines and Online Price Dispersion’. International Journal of Industrial Organization 29 (6): 655–67.

Bamberger, Gustavo E., Dennis W. Carlton, and Lynette R. Neumann. 2004. ‘An Empirical Investigation of the Competitive Effects of Domestic Airline Alliances*’. The Journal of Law and Economics 47 (1): 195–222.

Berry, Steven. 1992. ‘Estimation of a Model of Entry in the Airline Industry’. Econometrica 60 (4): 889–917.

Berry, Steven, Michael Carnall, and Pablo T. Spiller. 1996. ‘Airline Hubs: Costs, Markups and the Implications of Customer Heterogeneity’. National Bureau of Economic Research.

Berry, Steven, and Panle Jia. 2010. ‘Tracing the Woes: An Empirical Analysis of the Airline Industry’. American Economic Journal: Microeconomics 2 (3): 1–43.

Blundell, Richard, and Thomas M. Stoker. 2005. ‘Heterogeneity and Aggregation’. Journal of Economic Literature 43 (2): 347–91.

Borenstein, Severin, and Nancy L. Rose. 1991. ‘Competition and Price Dispersion in the

US Airline Industry'. National Bureau of Economic Research.

Borenstein, Severin, and Nancy L. Rose. 1994. 'Competition and Price Dispersion in the U.S. Airline Industry'. Journal of Political Economy 102 (4): 653–83.

Brueckner, Jan K. 2001. 'The Economics of International Codesharing: An Analysis of Airline Alliances'. International Journal of Industrial Organization 19 (10): 1475–98.

Brueckner, Jan K. 2003. 'International Airfares in the Age of Alliances: The Effects of Codesharing and Antitrust Immunity'. Review of Economics and Statistics 85 (1): 105–18.

Brueckner, Jan K., Darin Lee, and Ethan S. Singer. 2013. 'Airline Competition and Domestic US Airfares: A Comprehensive Reappraisal'. Economics of Transportation 2 (1): 1–17.

Brueckner, Jan K., and W. Tom Whalen. 2000. 'The Price Effects of International Airline Alliances'. The Journal of Law & Economics 43 (2): 503–46.

Brueckner, Jan, and Pablo Spiller. 1991. 'Competition and Mergers in Airline Networks'. International Journal of Industrial Organization 9 (3): 323–42.

Chandra, Ambarish, and Mara Lederman. 2018. 'Revisiting the Relationship between Competition and Price Discrimination'. American Economic Journal: Microeconomics 10 (2): 190–224.

Ciliberto, Federico, and Jonathan W. Williams. 2010. 'Limited Access to Airport Facilities and Market Power in the Airline Industry'. Journal of Law and Economics 53 (3): 467–495.

Dai, Mian, Qihong Liu, and Konstantinos Serfes. 2013. 'Is the Effect of Competition on Price Dispersion Nonmonotonic? Evidence from the U.S. Airline Industry'. The Review of Economics and Statistics 96 (1): 161–70.

Dana, James D. 1999. 'Using Yield Management to Shift Demand When the Peak Time Is Unknown'. The RAND Journal of Economics 30 (3): 456.

David, Herbert A., and H. N. Nagaraja. 2003. Order Statistics. 3 edition. Hoboken, N.J: Wiley-Interscience.

Deneckere, Raymond, Howard P. Marvel, and James Peck. 1996. 'Demand Uncertainty, Inventories, and Resale Price Maintenance'. The Quarterly Journal of Economics 111 (3): 885–913.

Escobari, Diego. 2012. 'Dynamic Pricing, Advance Sales and Aggregate Demand Learning in Airlines'. The Journal of Industrial Economics 60 (4): 697–724.

Escobari, Diego, and Manuel A. Hernandez. 2015. 'Screening and Price Discrimination with Unobserved Consumer Types: Evidence from Airlines'. SSRN Scholarly Paper ID 2552313. Rochester, NY: Social Science Research Network.

Evans, William N., and Ioannis N. Kessides. 1993. 'Localized Market Power in the U.S. Airline Industry'. The Review of Economics and Statistics 75 (1): 66–75.

Evans, William N., and Ioannis N. Kessides 1994. 'Living by the "Golden Rule": Multimarket Contact in the U. S. Airline Industry'. The Quarterly Journal of Economics 109 (2): 341–66.

Gaggero, Alberto A., and Claudio A. Piga. 2011. 'Airline Market Power and Intertemporal Price Dispersion'. The Journal of Industrial Economics 59 (4): 552–577.

Gale, Ian L., and Thomas J. Holmes. 1993. 'Advance-Purchase Discounts and Monopoly Allocation of Capacity'. The American Economic Review, 135–146.

Gayle, Philip G. 2007. 'Is Virtual Codesharing a Market Segmenting Mechanism Employed by Airlines?' Economics Letters 95 (1): 17–24.

Gayle, Philip G. 2008. 'An Empirical Analysis of the Competitive Effects of the Delta-Continental - Northwestern Code-Sharing Alliance'. The Journal of Law and Economics 51 (4): 743-66.

Gentry, Matthew L., Timothy P. Hubbard, Denis Nekipelov, and Harry J. Paarsch. 2018. 'Structural Econometrics of Auctions: A Review'. Foundations and Trends® in Econometrics 9 (2–4): 79–302.

Geradin, Damien, and Nicolas Petit. 2005. 'Price Discrimination under EC Competition Law: The Need for a Case-by-Case Approach'. College of Europe DP 7 (05).

Gerardi, Kristopher S., and Adam Hale Shapiro. 2009. 'Does Competition Reduce Price Dispersion? New Evidence from the Airline Industry'. Journal of Political Economy 117 (1): 1–37.

Gillen, David, and Benny Mantin. 2009. 'Price Volatility in the Airline Markets'. Transportation Research Part E: Logistics and Transportation Review 45 (5): 693–709.

Hernandez, Manuel A., and Steven N. Wiggins. 2014. 'Nonlinear Pricing Strategies and Competitive Conditions in the Airline Industry'. Economic Inquiry 52 (2): 539–61.

Ito, Harumi, and Darin Lee. 2007. 'Domestic Code Sharing, Alliances, and Airfares in the U.S. Airline Industry'. The Journal of Law and Economics 50 (2): 355–80.

Kang, Nam-Hoon, and Kentaro Sakai. 2000. 'International Strategic Alliances: Their Role in Industrial Globalisation'. OECD Science, Technology and Industry Working Papers 2000/05.

Klemperer, Paul. 2004. Auctions: Theory and Practice. Princeton, NJ: Princeton University Press.

KPMG International. 2017. 'Strategic Alliances: A Real Alternative to M&A?'

Lazarev, John. 2013. 'The Welfare Effects of Intertemporal Price Discrimination: An Empirical Analysis of Airline Pricing in US Monopoly Markets'. working paper.

Liu, Qihong, and Konstantinos Serfes. 2006. 'Second-Degree Price Discrimination and Price Dispersion: The Case of the US Airline Industry'. Social Science Research Network, Rochester, NY.

Luttmann, Alexander. 2018. 'Evidence of Directional Price Discrimination in the U.S. Airline Industry'. International Journal of Industrial Organization, March.

Mantin, Benny, and Bonwoo Koo. 2009. 'Dynamic Price Dispersion in Airline Markets'. Transportation Research Part E: Logistics and Transportation Review 45 (6): 1020–29.

McAfee, R. Preston, and Vera Te Velde. 2006. 'Dynamic Pricing in the Airline Industry'. Forthcoming in Handbook on Economics and Information Systems, Ed: TJ Hendershott, Elsevier.

Miller, Amalia R. 2010. 'Did the Airline Tariff Publishing Case Reduce Collusion?' The Journal of Law & Economics 53 (3): 569–86.

Nadarajah, S., and S. Kotz. 2008. 'Exact Distribution of the Max/Min of Two Gaussian Random Variables'. IEEE Transactions on Very Large Scale Integration (VLSI) Systems 16 (2): 210–12.

Paarsch, Harry J. 1997. ‘Deriving an Estimate of the Optimal Reserve Price: An Application to British Columbian Timber Sales’. Journal of Econometrics 78 (2): 333–57.

Paarsch, Harry J., and Han Hong. 2006. An Introduction to the Structural Econometrics of Auction Data. Cambridge, MA: MIT Press.

Park, Jong-Hun. 1997. ‘The Effects of Airline Alliances on Markets and Economic Welfare’. Transportation Research Part E: Logistics and Transportation Review 33 (3): 181–95.

Sengupta, Anirban, and Steven N. Wiggins. 2012. ‘Comparing Price Dispersion on and off the Internet Using Airline Transaction Data’. Review of Network Economics 11 (1).

Sengupta, Anirban, and Steven N. Wiggins. 2014. ‘Airline Pricing, Price Dispersion, and Ticket Characteristics on and off the Internet’. American Economic Journal: Economic Policy 6 (1): 272–307.

Siegert, Caspar, and Robert Ulbricht. 2014. ‘Dynamic Oligopoly Pricing: Evidence from the Airline Industry’. Available at SSRN 2413180.

Stigler, George. 1961. ‘The Economics of Information’. Journal of Political Economy 69.

Talluri, Kalyan, and Garrett Van Ryzin. 1998. ‘An Analysis of Bid-Price Controls for Network Revenue Management’. Manage. Sci. 44 (11): 1577–1593.

Williams, Kevin R. 2013. ‘Dynamic Airline Pricing and Seat Availability’. mimeo.

3.8 Appendix

3.8.1 Order statistics and Maximum Likelihood derivation

To estimate a model using the maximum likelihood approach, we need to specify the distribution by which our data is generated. In the framework that we propose, the price is an order statistic of the reservation cost; it is the highest reservation cost from two randomly drawn reservation costs. Order statistics and their distribution are important elements of auction models, where the winner is chosen based on a ranking of the bids, and the bids are usually a monotonic function of the underlying random costs or valuations. For a detailed exposition on the derivation of order statistics, see Paarsch and Hong (2006), particularly Appendices 1 and 2 in their book. Here, we will explain intuitively how the distributions are derived in our model.

Both carriers l and m offer the same service and draw their reservation costs, c_l and c_m, from the same distribution $F(c)$. A price observation takes the value p in two distinct cases: in the case where player l makes the sale, and in the case where player m makes the sale.

Let us take the last case, where player m makes the sale, which is an outcome of two independent events happening simultaneously. Given the strategies of the players to lower their offers/bids until it is no longer profitable, we know that the "loser", carrier l, which has the higher reservation cost, will have reservation cost $c_l = p$ exactly. At the same time, the "winner", carrier m, or the carrier with the lower reservation cost, must have $c_m < c_l = p$. These are two independent events, and therefore, the probability of observing price equal to p is the product of the probabilities of the two events:

$$P(p|m\ wins) = P(c_m < c_l)P(c_l = p)$$

The probability of the first event is the sum of all probabilities for which $c_m < c_l = p$. With a continuous distribution, this is the cumulative density $F(p)$. The probability of the second event is exactly the density of the distribution at p: $f(p)$. Hence, we have:

$$P(p|m\ wins) = F(p)f(p)$$

Similarly, due to the symmetry of the players, the case of observing p when player l wins has the following probability:

$$P(p|l \; wins) = F(p)f(p)$$

Then, the unconditional event of observing the price p is the sum of the cases where l wins and m wins:

$$P(p) = 2F(p)f(p)$$

3.8.2 Data cleaning

Regional "feeder" or "commuter" carriers are recoded as their major carrier partner. The full table can be provided upon request. Carriers with less than 15 passengers are deleted, since these probably reflect coding errors. We also remove tickets with prices lower than 50 USD and higher than 3000 USD. Most of these happen to be tickets at zero USD, representing frequent flyer purchases. We also focus on markets with more than nine passengers per quarter, as that is equivalent to one passenger per day given that the sample represents 10 percent of the quarterly ticket sales.

Another modification of the data set comprises grouping airports in the same metropolitan area.[30] The six groups of airports are: Dallas-Fort Worth International and Love Field in Dallas, TX; Baltimore/ Washington International, Dulles, and National in Washington, DC; Midway and O'Hare in Chicago, IL; Kennedy, LaGuardia, and Newark in New York, NY; Los Angeles, Burbank, and Long Beach in Los Angeles, CA; San Francisco, Oakland, and San Jose in San Francisco, CA. For example, Chicago Midway and Chicago O'Hare International will represent the same market. Again, this is a standard treatment in the literature (Berry and Jia (2010). Note this modification affects only approximately 15 percent of our observations, as we are working with duopoly markets that are usually markets less central to the network.

Following Evan and Kessides (1993, 1994), we count carriers as operating in a given market if their sales represent at least 1 percent of observations in the data, equivalently 1 percent of total sales.

[30] Note than not all the metropolitan areas that are usually considered in the literature are included in our database such as Miami or Houston.

We study exclusively direct and connecting duopoly markets. In any market, we might observe unusual choices by travelers using long paths (two or three connecting airports) due to capacity constraints in the supply or to random events such as bad weather conditions, technical issues on a plane or strikes. Therefore, we include in our database markets where a third airline exist with a market share smaller than 5 percent or markets where less than 5 percent of the passengers use alternative routes (for instance, flying with 3 coupons), although these passengers are excluded from our analysis.

3.8.3 Deriving the Coefficient of Variation

The coefficient of variation (CV) is defined as the variable's standard deviation divided by the variable's mean. In our case, we are interested on the transaction price p with mean, m, and standard deviation s. The CV can be expressed as:

$$CV = \frac{s}{m} = \left(\frac{E[p^2] - E[p]^2}{E[p]^2}\right)^{1/2} = \left(\frac{E[p^2]}{E[p]^2} - 1\right)^{1/2}$$

Our model analyzes the logarithm of the reservation costs in duopoly markets. We call $log(c_l)$ and $log(c_m)$ the logarithms of the reservation costs of our two competitors, airlines l and m, respectively. The logarithm of the transaction price is the highest of the two reservation costs, $log(p) = max(log(c_l), log(c_m))$.

With the moment generating functions obtained from Nadarajah and Kotz (2008) for the max/min of two random variables, we can compute:

$$E[p] = E[e^{log(p)}] = 2e\Big(\mu + \frac{\sigma^2}{2}\Big)\Phi\Big(\frac{\sigma}{\sqrt{2}}\Big)$$

$$E[p^2] = E[e^{log(p)}e^{log(p)}] = E[e^{2log(p)}] = 2e\Big(2\mu + 2\sigma^2\Big)\Phi\Big(\frac{2\sigma}{\sqrt{2}}\Big)$$

which implies that:

$$CV = \left(\frac{E[p^2]}{E[p]^2} - 1\right)^{1/2} = \left(\frac{e(\sigma^2)\Phi\Big(\frac{2\sigma}{\sqrt{2}}\Big)}{2\Big(\Phi\Big(\frac{\sigma}{\sqrt{2}}\Big)\Big)^2} - 1\right)^{1/2}$$