Search, Matching, and the Role of Digital Marketplace Design in Enabling Trade: Evidence from Airbnb

Abstract

Two-sided marketplaces are distinguished by the fact that both sides have preferences regarding each others' non-price characteristics. This paper studies how digital platform design affects transaction costs and volume in these markets by analyzing the decisions of guests and hosts to search and match with each other on Airbnb. I show that the two-sided nature of the market is important. Through 2014, rejections of guests by hosts occur for 42% of inquiries regarding booking and these rejections causally decrease the rate at which guests eventually book on the platform by 43% to 70%. Rejections are primarily caused by stale vacancies and the screening of guests by hosts. I use data on search and communication to estimate a model of guest and host choices. I apply this model to study the effects of search engine design and find that, by tracking listing availability, Airbnb reduces rejections by 59%. I then show that incorporating host preferences into rankings can further increase match rates and discuss how Airbnb's subsequent innovations reflect these findings.

1 Introduction

There are always individually owned assets not being used and people with free time who are not working. Without transaction costs, many of these assets and people could be employed in productive ways, if only for a short time. These transactions could provide value by expanding market capacity, offering better matching products, and providing new sources of income for sellers. However, the costs of transacting between unacquainted individuals have been so large in developed economies that people instead preferred to exclusively transact with traditional firms in most markets. Over the past 20 years, peer-to-peer marketplaces have greatly expanded transaction volume between individuals in markets such as short-term apartment rentals (Airbnb), used goods (Ebay), labor services (Upwork and Taskrabbit), and rides (Uber and Lyft).[1]

In this paper, I use the setting of Airbnb to study the transaction costs which prevent trade in these markets, the role search engine design in reducing those costs, and the potential for further improvements in search and matching. The motivation for this study is two-fold. First, designing better matching systems is a practical problem for many digital marketplaces. I provide new results regarding the causes and business consequences of transaction costs on these platforms. In particular, I show that rejections of guests by hosts occur 42% of the time and that they decrease the transaction probability of searchers by over 40%. I then study interventions designed to reduce the frequency of these rejections and to increase transaction rates. Second, as digital intermediaries are becoming more prominent, it is increasingly important to understand how they generate value. Peer-to-peer marketplaces implement many features, including payments processing, messaging, reputation, and search. I show that the design of the search engine, namely its role in reducing rejections due to inaccurate availability information, is especially important in making Airbnb succeed.

To begin with, unlike traditional markets for accommodations (e.g. hotels), peer-to-peer mar-

[1] See Einav, Farronato and Levin (2016) for a general overview of peer-to-peer markets and Farronato and Fradkin (2018) for an empirical analysis of the effect of Airbnb on the market for short-term accommodations. Horton and Zeckhauser (2016) and Fraiberger and Sundararajan (2015) describe equilibrium models of asset ownership and use as a function of a reduced form transaction cost parameter.

ketplaces involve not just search, but search and *matching*. The reason for this is that the sellers on these platforms have limited capacity and sometimes have preferences towards which buyers they would like to transact with. As an illustrative example consider the contrast between a hotel and a host on Airbnb. The hotel may have hundreds of similar rooms and would like to rent them out every night. The typical Airbnb host has one unique apartment or room and does not always wish to rent it out. That host may also care about the length of stay or the type of guest. This results in two challenges faced by peer-to-peer platforms and not faced by other digital marketplaces.

First, consumers face much more heterogeneous options in peer-to-peer markets. The median Airbnb searcher to Chicago between 2013 and 2014 browses pages that load just 4.2% of the thousand plus distinct listings that match a given search parameters and have not yet been booked for the search dates. Even the more motivated searchers who send contacts regarding booking are shown only 5.5% of these listings. This search is highly directed and heterogeneous: 57% of searchers filter for a location within the city, 70% filter for a room type, and 52% apply the maximum price filter. The process of search consequently takes time. Searchers in my sample spend an average of 58 minutes browsing the site before sending an inquiry whereas a consumer who books on 'booking.com' spends just 34 minutes browsing before booking.[2] The fact that search is time consuming and limited suggests that the search engine may have an important role in determining whether a searcher matches.

Second, unlike in retail markets, consumers in peer-to-peer markets face a chance of being rejected by a seller, either due to the lack of availability or due to seller preferences. Rejection occurs frequently: 42% of inquiries regarding booking are rejected. I classify rejections into three categories, each of which have differing implications for market design. Stale vacancy rejections, which happen when the listing is eventually marked as unavailable by the host and not booked for the dates of an inquiry, occur 15% of the time. Another 8% of rejections occur due to congestion, which happens when listings are booked by inquiries from guests who sent earlier inquiries for the set of dates.[3] Lastly, the remaining 14% of inquiries are rejected due to 'screening', which is

[2] This figure was calculated using the time spent browsing in the two days prior to purchase in the 2013 comScore web panel.

driven by host preferences regarding the characteristics of the searcher or the trip.

These rejections represent transaction costs on the platform. First, communication is costly and leads to delay and uncertainty. Second, rejection leads searchers to leave the Airbnb platform. Conditional on being rejected from their first inquiry, searchers are 43% to 70% less likely to eventually book a listing for a given market. I demonstrate that this effect is causal by showing that it persists even when controlling for market-level availability of rooms, guest and listing characteristics, and guest motivation. I also use the presence of stale vacancies, which should be exogenous to host preferences regarding a particular trip, as an instrument for rejection and find that the effect of a rejection on eventual booking by a searcher persists. This effect is also not driven by guest experience with the platform. Those guests who've previously booked experience a higher decrease in booking rates after an initial rejection.

Next, I study how the search engine affects the rates of rejections and successful matches. I find that the digital calendar which tracks previously booked and other unavailable listings plays an especially important role. If Airbnb did not automatically remove already booked listings from search, more guests would try to book these hosts and the rate of rejection would increase by 144%. I also show that the share of accepted searchers can increase by 10% when search ranking algorithms use the expected probability of an acceptance by a host relative to a similar algorithm that does not use this information. Consequently, there is sufficient supply in the market so that most rejected guests would be willing to stay at other available listings.

To simulate these counterfactual outcomes, I estimate predictive models of searcher and host behavior. I use a discrete choice model to predict a searcher's choice of whether and whom to contact from a fixed consideration set. Importantly, I use each searcher's choice of filters to account for preference heterogeneity. For example, searchers who filter for a particular neighborhood in a city are far more likely to send an inquiry to a listing in that neighborhood. On the host side, I use the set of inquiries which are not rejected due to a stale vacancy or due to congestion to estimate a logistic regression predicting rejection as a function of guest and listing characteristics. I find

[3] Note that my definition of congestion is more narrow than that of Roth (2008), whose definition would treat all rejections on the platform as congestion.

that the heterogeneity in selectivity across listings as measured by listing random or fixed effects is large relative to the coefficients on inquiry characteristics. This points to a large role for the platform to direct search to less selective hosts.

My simulation works by studying alternative outcomes under the assumption that each searcher browses pages that load the same number of distinct listings as they did in the data. I first calculate the expected outcomes from a search engine without availability tracking (the removal of previously booked or marked unavailable listings), filtering, and search engine ranking. I find that under this scenario, the share of searchers who choose the outside option increases by 1.6 percentage points and the share of rejected inquiries increases from 32% to 78%.[4] This results in a 68% drop in the rate of accepted inquiries. I then decompose the effects of filtering and availability tracking. I find that in a counterfactual with availability tracking but without filtering and ranking, the share of individuals who send an accepted inquiry drops 25% relative to the status quo. Therefore, availability tracking is relatively more important than ranking and filtering.

The above counterfactuals can be thought of as reverting the search engine to be like Craigslist, which also offers short-term rentals (Figures 1 and 2). Unlike Airbnb, Craigslist operates as a mostly passive listing service that does not track transactions. Consequently, listings on Craigslist may have already been booked when they are shown to searchers. Furthermore, early versions of Craigslist did limit results to relevant geographies, dates, or prices, and displayed listings in chronological order rather than by relevance. In contrast, Airbnb directly mediates the transaction between guest and host (e.g. Varian (2010)). This allows it to remove previously booked listings, display only relevant geographies and prices, and to use a ranking algorithm based on historical data regarding searcher outcomes.

Finally, I consider the effects of ranking algorithms that display alternative consideration sets to searchers. I show that without availability tracking, better ranking algorithms make little difference. However, with availability tracking, matching outcomes would improve by a meaningful amount. For example, I find that a ranking algorithm which personalizes search results with regards

[4]Throughout the paper, the "outside option" refers to a decision by a searcher to choose a hotel option, a non-market option such as staying with friends and family, or to stay at home.

to both expected utility and the probability of screening rejections by hosts would increase the rate at which searchers send a contact and are accepted to 34% from 24%. Because this counterfactual is static, it does not account for the market equilibrium effects of the policy. Fradkin (2015) simulates these effects accounting for the fact that a booking today reduces supply for subsequent searchers and that policies affect the level of congestion in the market.

The models I estimate in this paper are simple relative to those in the literature on endogenous search decisions with rational expectations (e.g. De los Santos and Koulayev (2017)) and equilibrium price adjustment in search and matching markets (e.g. Gavazza (2016)). There are two reasons for this simplicity. First, my approach is computationally tractable, which is important when there are thousands of options, numerous potential consideration sets,[5] and sparse choices. Second, the goal of this model is to demonstrate the role of the search engine rather than to completely characterize equilibrium responses to ranking algorithms. The main results of the above counterfactuals, namely that the search engine plays in important role in reducing rejections and that guests who are rejected typically have accepted an alternative listing that is available, would not go away even if searchers altered their search intensity or sellers adjusted prices.[6]

The analysis conducted in this paper has been influential within Airbnb and much of the work done for this paper was conducted while I was working as a data scientist there. For example, in 2015, Airbnb announced a new policy of using host rejection behavior to rank listings.[7] In particular, as suggested by this work, the implemented algorithm calculates the likelihood of rejection given a set of query parameters and surfaces listings that are less likely to reject.

Another Airbnb intervention focused on the problem of stale vacancies. Specifically, if stale vacancies are caused by the inattention of hosts, then reminders should be helpful in reducing rejections. In 2018, Airbnb announced a successful market design intervention that used data on host calendar checking behavior to target email reminders to likely inattentive hosts.[8] Especially

[5] For example, the number of possible consideration sets when there are 70 choices from 1000 listings equals $7.04036 * 10^{108}$.

[6] Farronato and Fradkin (2018) show that many hosts are on the margin of hosting on a given night. Consequently, there is not much room for hosts to lower prices. Between 2013 and 2015, the median number of contacts per person stays constant at 1, even as rejection rates decrease and booking rates increase. This suggests that search intensity is not responsive to modest changes in rejection rates.

relevant for this paper is that they find a monotonic relationship between a measure of host inattention to the calendar and rejections, which is consistent with the mechanism for stale vacancies described above.

Lastly, Airbnb has been promoting 'Instant Book', which allows hosts to opt-in to automatically accepting guests who specify particular criteria. The feature's major trade-off is that it gives hosts less control. Nonetheless, this feature has grown from accounting for fewer than 10% during my sample period to accounting for 60% of bookings in 2017 (Skift (2017)).

The closest papers to this one include both empirical and theoretical contributions regarding search in digital marketplaces. Horton (2016) empirically shows that the availability friction is also important on Upwork and studies how availability signaling can help. Arnosti, Johari and Kanoria (2014) present a theoretical model of matching in a peer-to-peer market use it to show that the equilibrium in the model is inefficient due to uncertain availability. They show that limiting applications can improve market efficiency. My paper confirms the relevance of their theoretical model by demonstrating that stale and congested vacancies are a major cause of rejections on Airbnb. Furthermore, I show how the platform can use the data it observes on the search and matching process to alleviate these frictions by altering the consideration sets of searchers.

Dinerstein et al. (2014) study a redesign of the search engine on eBay, and Ellison and Ellison (2009) show that sellers respond to high competition online by obfuscation, counteracting some of the benefits of search engines. Chen and Yao (Forthcoming), los Santos and Koulayev (Forthcoming), Ghose, Ipeirotis and Li (2014), and Ursu (2018) study the effects of search rankings in the hotel industry. Bronnenberg, Kim and Mela (2016) study search in the market for digital cameras and find that, like in this setting, search is highly directed and subject to recall. Relative to these contributions, my paper is unique in that it focuses on a two-sided matching market with availability frictions and that it has very rich data on the matching process. The complexity of Airbnb makes it very difficult to estimate rational expectations models of browsing and communi-

[6] For an Airbnb blog post regarding "Host Preferences" see: `http://nerds.airbnb.com/host-preferences/`.

[7] For an Airbnb blog post regarding "Contextual Calendar Reminders" see: `https://medium.com/airbnb-engineering/contextual-calendar-reminder-key-to-successful-hosting-9be89e1a32fd`

cation. Consequently, the focus of this paper is in documenting the search and matching process and exploring how the consideration set of searchers affects the likelihood of a match.

This paper necessarily focuses on a limited set of digital marketplace design choices. Other mechanisms also play an important role in these markets. Reputation systems reveal seller quality in ways that reduce both adverse selection and moral hazard on the part of sellers (Resnick et al. (2000), Klein, Lambertz and Stahl (Forthcoming), Fradkin, Grewal and Holtz (2017)). Innovative pricing and matching mechanisms such as auctions (Einav et al. (Forthcoming)), employer initiated search (Horton (2017)), and surge pricing (Hall, Kendrick and Nosko (2016)) are frequently used to clear the market. The informational structure of the market, including the ability to post photos (Lewis (2011)) and disclose quality (Tadelis and Zettelmeyer (2015)), can also affect market efficiency. Lastly, non-design features of peer-to-peer marketplaces affect their success as well. Cullen and Farronato (2016) show that Taskrabbit is more successful in cities with high geographic density and Farronato and Fradkin (2018) show that Airbnb is more successful in cities with constraints to building hotels.

Search and matching also plays a large role in the economics of labor, housing, and household formation. Theoretical results in this literature such as Burdett, Shi and Wright (2001), Albrecht, Gautier and Vroman (2006), and Kircher (2009) show that markets where sellers have limited capacity, such as Airbnb, entail higher search costs than markets with large firms. Two recent papers by Cheron and Decreuse (2017) and Albrecht, Decreuse and Vroman (2017) are particularly relevant. These papers build models of search and matching with 'phantom' vacancies, which are analogous to congested and stale vacancies in my setting, and study their implications in the labor market. One notable advantage of my setting is the availability of data on search and communication, which is missing in the above papers.[9] Consequently, I can show directly that, at least in my setting, phantom vacancies are quantitatively important and are mainly caused by sellers not updating their calendars. Furthermore, I show that without the platform's availability tracking, phantom vacancies would predominantly be caused by congestion rather than stale vacancies.

[9] Other papers with similar data include Wolthoff (2018) for the labor market, Hitsch, Hortaçsu and Ariely (2010) for the dating market, and Piazzesi, Schneider and Stroebel (2015) for the housing market.

2 Setting

Airbnb describes itself as a trusted community marketplace for people to list, discover, and book unique accommodations around the world – online or from a mobile phone. The marketplace was founded in 2008 and has intermediated over 300 million guest arrivals since then. As of 2018, Airbnb has homes in over 81,000 cities and over 4 million total listings. Airbnb created a market for a previously rare transaction: the short-term rental of an apartment or part of an apartment by a consumer.[10] In addition to a search engine, Airbnb operates a reputation and fraud detection system, customer service, a communications platform, a mobile application, an insurance policy for hosts, and a transactions processing platform. This paper investigates the role of a subset of these technologies.

Depending on the market, Airbnb's main competition is either with hotels or with traditional vacation rentals. For the top 50 US cities, Airbnb comprised 2% of available rooms by the end of 2014 and an even lower share of transactions. However, this masks a lot of heterogeneity since some cities had Airbnb supply shares exceeding 15% by the end of 2014. See Farronato and Fradkin (2018) for more details regarding the outside option (typically hotels, not taking a trip, or traveling for fewer days) and the seasonality of these markets.

A typical Airbnb transaction consists of the following steps:

1. Using the Search Engine (Figure 1) - Searchers enter the travel dates, number of guests and location into a search engine and receive a list of results. The search can then be refined using filters. Only listings that have an 'open' calendar for the trip dates are potentially loaded in the search results. Calendar dates become unavailable either when a listing is booked or when a host updates the calendar to be unavailable. Importantly, calendars are frequently not an accurate representation of true availability. This occurs because hosts do not always attend to their calendar or because hosts may be in conversation with other potential bookers

[10] Couchsurfing, a large travel based social network started in 2003, facilitates similar stays but without monetary exchange. Vacation rentals by owners in tourist destinations have also existed for a long time.

(either on or off of Airbnb the platform).

2. Investigation (Figure 3) - The searcher clicks on a listing in search. The subsequent page displays additional photos, amenities, reviews, house rules and information about the host.

3. Communication and Booking (Figure 4) - The searcher sends a message to hosts inquiring about room details and availability. This can be done in one of two ways, either by sending an inquiry or by clicking the "Book It" button. In the case of an inquiry, a host will typically reply with an acceptance or a rejection. If accepted, the guest can then click the "Book It" button to go through with the booking. A host who has received either type of request has the right to make a final decision of whether to accept or reject. There are two exceptions. First, some hosts are available to be "Instant Booked" by some guests, in which a transaction is confirmed as soon as guests click "Instant Book". Second, a host can "Pre-approve" a guest after an inquiry, which subsequently allows the guest to book without further communication. Figure 5 shows that instant booking occured for fewer than 10% of contacts during the time period studied in this paper.

4. Stay - There is frequently communication regarding the key exchange and details of the trip. Either party can cancel a booking with a pre-specified cancellation penalty (a monetary amount for a guest and an Airbnb specific punishment for the host). Following the transaction, guests and hosts can review each other.

3 Browsing Behavior

In this section, I provide a comprehensive description of the process by which searchers on Airbnb find a suitable listing. I first start by describing the data used to conduct the study. I then use the data to document three features of search behavior: limited consideration sets, redundant search, and the filtering of results.

3.1 Data Selection Procedure

The data use for this study focuses on a set of users who searched for short-term rentals in Chicago between September 2013 and September 2014. For contacting and booking behavior, I validate the representativeness of the results with data from other US cities. For browsing data, I limit the analysis to a 10% sample of these users because the quantity of data would otherwise be even more difficult to work with. I define a short-term rental as having fewer than 8 nights in duration. I limit my focus to one city due to the fact that I need to estimate supply and demand, which vary at a city level. Chicago is a typical US city for Airbnb in the sense that its supply share of the market in 2014 was close to the sample median of 2% in Farronato and Fradkin (2018). Because Airbnb was the first platform to successfully enable peer-to-peer short-term rental transactions in US cities, the US city setting is particularly interesting to study.

The search data I observe includes each search by a user fulfilling the above criteria, the filters applied, the page number of the search results, the dates of the stay, the number of results, the position the map was centered on and the zoom level (when the map was used), the time of the search, the filters loaded, and the listings shown (as well as their ranks on the page). I do not observe whether and when the searcher clicked on a particular listing. I further limit the sample to searches whose cookie or device-ID can be linked to a registered user. Furthermore, I remove anomalous searches such as searches with 0 nights, searches where the check-in date has already passed, searches with more than 6 guests, searches likely to be conducted by bots (e.g. those with more than 100 searches).[11] Lastly, searches more than 8 weeks ahead of the check-in date are removed. These comprise just 26% of inquiries in the sample but complicate the algorithm I use for keeping track of potentially visible listings.

Next, I group searches into distinct search spells so that searches by the same searcher which differ in destinations and trip date are kept separate. To create a search spell, I first link the searches to a contact (an inquiry or booking) conducted by the searcher. For those searches that can be

[11] Bots are software agents that programatically browse websites for the sole purpose of collecting information. E-commerce sites are frequently 'scraped' by bots for the purposes of competitive analysis and research.

linked to a contact, I only keep the searches which occurred within two days preceding the contact. Furthermore, I use only the searches related to the first contact by a user in the city during the sample period. For those searches that cannot be linked to a contact, I keep only the searches conducted within the last two days of search activity. This selection criteria ensures that the search results in the data reflect perceived availability of the searcher before an inquiry. The final set of searches contains 236 thousand observations.

3.2 Descriptive Statistics

I now describe the behavior of searchers in the sample. I first focus on a sample of all searchers, since the decision to send a contact to a host depends on the search effort of the searcher and the quality of the options in the market. Table 1a displays the summary statistics at a search-spell level for the sample. The median searcher submits 9 distinct search requests in the process of searching.[12] There is significant heterogeneity in the number of searches. The 25th percentile is just 3 searches while the 75th percentile results in 21 searches. Furthermore, the mean number of searches, 19, is twice as large as the median indicating significant skewness in the distribution. This search behavior is not simply a result of differences in the time spent per page of search results. I measure the time spent browsing by first creating sessions, which are sets of searches without a gap between user actions longer than five minutes. I then compute the time spent in each session as the difference between the start and end times of a session. Lastly, I sum the times spent in each session across all sessions in a search spell to compute the total total time spent browsing. The median searcher takes 17 minutes (average of 45 minutes) in search but this search time is heterogeneous with an interquartile range between 6.5 and 41 minutes. Searchers who eventually send an inquiry search spend a median of 32 minutes (average of 58 minutes) searching, reflecting their greater intent and perhaps the presence of better matching options.

Those searchers who eventually send a contact engage in more search than those who do not.

[12] A search request can result from an application of a filter, a shift of the map, a click to the next page, or a return to a previously seen set of search parameters.

Table 1b displays the summary statistics for the searchers in the sample who sent at least one contact. The typical searcher with a contact spends approximately twice as much time browsing and views twice as many search results as the typical searcher in the sample. This increased search activity can be caused by two factors. First, there is unobserved heterogeneity between searchers. Those who get more benefit from using Airbnb relative to the outside option should search more intensively and be more likely to contact. Second, there is endogenous selection into contacting based on the quality of search results. Searchers should be more likely to book when there are more relevant listings shown in search.

The typical searcher browses pages that load 54 unique listings during the search spell. This corresponds to 4.2% of all listings that match the search criteria and have not yet been booked for the set of search dates. Those searchers who send a contact typically browse pages comprising 73 unique listings, representing 5.5% of all potentially listings satisfying the search parameters. This limited search suggests an important role for marketplace technology to improve matches. If those listings which are not seen would be good matches, then the search engine could in principle design an algorithm to show these listings to the searcher at the beginning of search.

Another fact about search on the platform is that much of it is redundant. Airbnb's search engine displays a maximum of 18 results on a given page but the typical ratio of unique listings loaded in search to searches is 6. This redundant search happens for several reasons. First, searchers are frequently distracted by other tasks while looking for a room and oftentimes restart searching at a later time with generic parameters. Second, there is often overlap between filtered search results, such as when a user zooms in or slightly shifts the map. Third, users go back to generic search if they've closed the search results tab or pressed the 'back' button after investigating a listing. The eventually chosen listing is first seen early in search (typically on the fourth search) but the typical searcher conducts 7 additional searches before stopping the search process and sending a contact. Recall is also present in follow-up inquiries, as I'll discuss in the section on communication.

The presence of recall suggests that searchers either learn about product attributes sequentially

[12]The maximum number of listings per search was 21 in the earlier part of the sample.

as in (Ke, Shen and Villas-Boas (2016)), learn about the market-wide distribution of utilities during search (De Los Santos, Hortaçsu and Wildenbeest (2015)), or face an exogenous increasing marginal cost of search (Ellison and Wolitzky (2012)), which may occur in this setting due to work interruptions.

Filters offer a way for searchers to direct their search towards options which are likely to be better matches. Indeed, the estimates in section 5 show that filters are predictive of the options selected by a searcher. In this section I focus on the filters most prominently displayed by Airbnb: the map, price, and room type filters.[13] Maximum price filters and the map are used by over 50% of searchers and over 60% of those who contact. Those who use the map not only change the default location but also use the zooming features to focus on specific areas of interest. In addition to using the map to specify a geography, searchers can explicitly specify a search neighborhood from a menu, and 8% choose to do so. Lastly, guests on Airbnb can choose to rent a room within a property or an entire property. Over 63% of guests filter for a room type at least once during the search process.

Filtering behavior is heterogeneous across searchers and reflects the heterogeneity in searcher preferences. Figure 6, displays the distribution of map filtering behavior across 20 neighborhoods in the city. The most popular neighborhood is filtered in greater than 10% of searches with a map filter. The next five neighborhoods are filtered for more than 5% of the time, and the subsequent 13 neighborhoods are filtered for more than 1% of the time. Furthermore, of the 70% of searchers who use a room type filter, 87% select entire property and 35% select private room. These sum to more than one because the same searcher can filter for both. Lastly, Figure 7 displays the distribution of the maximum of the price filters each searcher uses. Similar to the other filters, there is significant heterogeneity in price filtering activity across searchers. The presence and heterogeneity of filtering behavior suggests new, and to my knowledge non-modeled, market design choices. Specifically, each marketplace must decide the types of filters offered, their prominence in the

[13] The search engine also allows explicit filters for neighborhood, which I group with the map filter. There are also offers filters for various amenities, property types, languages, and other miscellaneous options. These filters require more clicks to access and were used less than 1% of the time in the sample period.

interface, and the extent to which those filters are binding. For example, a redesign of Airbnb's search engine in 2013 expanded the map of results to occupy half of the search screen. In contrast, Booking.com, the most popular hotel search engine, makes a small map visible only after a searcher has scrolled past the initial search results.

4 Communication, Rejection, and Booking Behavior

The goal of this paper is to measure the transaction costs in this market and to study how digital market design reduces those costs. In this section, I show that even after Airbnb's initial success, transaction costs related to communication remained significant and caused searchers to leave the platform without transacting.

I use two samples to study communications and rejections. The first sample, analogous to the sample used to study browsing, consists of all contacts in Chicago between September 2013 and September 2014. The second sample, used to validate the representativeness of Chicago, consists of a 10% sample of all searchers who sent contacts to US markets. This sample is further limited to the top 50 markets in terms of contacts according to Airbnb's market definitions. For both samples, I keep only contacts regarding the first set of contact dates in a city in my sample.[14] This is done so that I measure the first interaction of a guest with Airbnb in the city of interest (as in Nosko and Tadelis (2015)).

Table 2a displays the summary statistics regarding the communication process in Chicago for first contacts by a searcher in Chicago. As a reminder of the notation, I call any communication a 'contact,' a non-binding communication by the searcher an 'inquiry,' and a binding communication, a 'booking request.' Turning first to the number of contacts, the median number of listings contacted by a searcher in Chicago for a given set of dates is 1 and the mean is 2.4. Of these inquiries, 1.4 are on average sent simultaneously, which I define as within 2 hours of the time at

[14] To do so, I find the minimum check-in date for all contacts in a city by a given guest. I then exclude any contacts from the sample for which the check-in date in the city is more than 2 days after the initial check-in. I also exclude any contacts which occur on the date of check-in and which require more than 7 nights.

which the first contact was sent by the searcher. Importantly, since 57% of contacts begin with an inquiry, searchers could send multiple inquiries without committing themselves to a purchase. Nonetheless most searchers choose not do so. Figure 8 shows the distribution of the number of simultaneous and sequential inquiries. While most searchers send just one simultaneous inquiry, 9% of users send at least 3 simultaneous inquiries.

Of all first contacts, 36% are rejected by hosts in the sample. This rejection rate is similar to the overall rejection rates for first contacts in US cities (Table 2b). The fact that searchers initially send few contacts has implications for the importance of these rejections. If a searcher sends one inquiry and it is rejected, then the searcher must conduct a new search before booking. On the other hand, a searcher who sends many initial contacts does not have to search again after one rejection.[15] Figure 8 shows the distribution of the number of sequential inquiries after a rejection. Those searchers who do return often send more than one inquiry. This is consistent with searchers learning about the possibility of being rejected and sending more inquiries as a result. The presence of searcher learning suggests that policies encouraging searchers to send more inquiries in the flow could potentially ameliorate the costly effects of rejections.

Lastly, when searchers do send additional inquiries, those inquiries are typically sent to listings that were seen prior to the first inquiry. To compute this number, I consider the set of users who send between one and five inquiries two hours after sending an initial inquiry. I then check whether the these inquiries were sent to listings present in the set of search results before the initial inquiry. Indeed, 78% of sequential inquiries were sent to listings seen in the initial search results. This provides evidence that searchers could have initially sent more inquiries to listings that were better than the outside option.

[15] Bargaining is another reason for communication in many search and matching markets. However, bargaining is rare on Airbnb. Bargaining is impossible when guests contact with the 'Book It' button. Furthermore, natural language processing of inquiries shows that bargaining related terms rarely occur for the short-term stays studied in this paper.

4.1 The Types and Frequencies of Rejection

If applicants knew, a priori, which hosts would accept and reject which inquiries, then they would not need to waste search effort looking at and sending contacts to rejecting hosts. The platform has an incentive to design the marketplace in order to prevent this wasteful search. However, the appropriate policies for preventing rejections depend on the mechanisms behind rejections. In this section, I provide a framework for categorizing rejections by hosts into three categories, congestion, "stale" vacancies, and screening, each with differing implications for market design and the sources of inefficiency. Note that my classification is more nuanced than is typical in the market design literature (e.g. Roth (2008)), which typically groups all rejections under the term 'congestion.'

In my framework, congestion occurs when a guest sends an inquiry to a host who is about to transact with someone else. The host must consequently reject the inquiry because she can only host one trip for a given set of dates. The reason that congestion happens is that potentially available listings are not removed from the search results until the transaction clears. The longer it takes the transaction to clear, the more likely it is that congestion occurs. In turn, the clearing time of transactions is determined by the time it takes for a host to respond to an inquiry and the time it takes for a guest to confirm the transaction.

I classify rejections as being caused by congestion when an inquiry is sent to a host who is subsequently booked as a result of a previous inquiry.[16] Using this methodology, congestion rejections occur for 7.8% of all inquiries in Chicago (Table 2a) and 7.1% of all inquiries in the US (Table 2b). These rejections constitute a relatively small percentage of the total rejections on this site and increase when the ratio of searchers to listings increases.

The reason for the relatively small frequency of congestion is that hosts tend to respond quickly when they accept a booking request. Figure 9 plots the distribution of response times for first in-

[16] I assume that hosts evaluate each inquiry sequentially rather than waiting to receive several inquiries and picking the best. In practice, there are cases when a host may receive inquiries in parallel, if, for example, she checks Airbnb infrequently. I abstract from this scenario because hosts are notified by text or email of an inquiry and have an incentive to respond quickly.

quiries by searchers to locations in the US. Over 50% of acceptances come within 3 hours of the initial inquiry and fewer than 10% take more than 2 days. On the other hand, rejections typically take much longer. For example, over 30% of rejections take longer than 2 days. The likely reason for this divergence in response times is that hosts have little incentive to respond quickly for inquiries that are unlikely to result in bookings.[17]

The second type of rejection in my framework is due to stale vacancies. These rejections occur when guests send inquiries regarding listings which are not actually available for a set of dates. Stale vacancies occur because hosts don't promptly block specific dates on their calendar even though they are not available. I am able to observe rejections due to stale vacancies when an inquiry is rejected for a set of dates which are subsequently marked as unavailable by the host. Hosts who never update their calendars to be unavailable may still have stale vacancies and reject all inquiries as a result. For example, of all listing/week of check-in combinations where the host did not block the calendar and received at least 3 inquiries, 5.1% rejected all the inquiries that they received. Therefore, my methodology understates the true extent of stale vacancies.

Stale vacancy rejections occur for 14.5% of first contacts in Chicago and 15.3% of first contacts to US hosts. An important contribution of my study is its ability to directly measure stale vacancies and to show how the platform mitigates the friction caused by these vacancies. Although stale vacancies may seem like an Airbnb specific phenomenon, they are common to many search and matching markets. For example, employers may not promptly remove posted vacancies even when a position has been filled or is no longer needed. Similarly, in online dating markets, people may not promptly disable a profile even when they are too busy to date or are in a relationship. Albrecht, Decreuse and Vroman (2017) calibrate a model of search and matching with 'phantom vacancies' in the labor market and use it to show that phantom vacancies contribute importantly to labor market frictions. My paper is the first to directly measure stale vacancies and demonstrate that they are important in a search and matching market.

The above methodology for identifying 'stale vacancies' could potentially conflate cases where

[17] To correct this problem, Airbnb has begun enforcing 'hosting standards', which, among other things, reward hosts in search rankings if they respond quickly. For details see: https://www.airbnb.com/hospitality.

a calendar was marked as unavailable because the listing was booked off of the Airbnb platform, either through another platform or in an informal transaction which dis-intermediates Airbnb. While this type of behavior is likely to be important in many search and matching settings, it is likely to be less important on Airbnb for the following reasons. First, with regards to multi-homing across platforms by hosts, surveys of Airbnb hosts in the US suggest that most hosts transact exclusively on Airbnb. Specifically, Airbnb ran a survey of hosts in New York and San Francisco during the timeframe studied in this paper asking them whether they rent exclusively through Airbnb. Over 70% of hosts said that, and over 20% said that they mostly used Airbnb.

Dis-intermediation without a platform is also unlikely to be a major reason for why calendar dates are marked as unavailable. First, there are large benefits to keeping transactions on the site because of the insurance, reputation and secure monetary transfer that using Airbnb offers. Second, Airbnb actively tries to prevent dis-intermediation by removing phone numbers, emails, and other contact information from messages before transactions are confirmed.

Screening, the final rejection type in my framework, occurs because hosts have preferences over trips and guests, and those preferences are not explicitly expressed by hosts to the platform. This results in hosts receiving inquiries regarding trips they are not willing to host, which are consequently rejected. 14% of first inquiries are rejected due to screening. In Section 5 I model hosts' screening rejection decisions in detail.

4.2 The Effects of Rejection on Searchers

Rejections and the related communication about a transaction are costly from both a user's and a platform's perspective. In order for searchers to make travel plans, they need to know where they'll be staying and when. This planning process is potentially delayed when communication takes time and there is a possibility of rejection. When rejection does occur, it may cause searchers to give up on using the Airbnb platform altogether and to switch into a marketplace with lower transaction costs, such as a typical hotel booking website. Lastly, to the extent that potential searchers know that rejections frequently occur Airbnb, they may not use the platform in the first place.

In this section, I document that rejection causes searchers to leave the Airbnb platform without transacting. While the direction of this effect is ex-ante obvious, the large effect size that I show is not. After all, a guest could have sent more contacts or could have returned to the platform after an initial rejection to search more. In the exercises below, I attempt to isolate *exogenous* rejections of guests, which may occur, if, for example, a guest was unlucky and tried to book a stale vacancy. In this case there were other listings in the market that would've been good enough for the guest had they contacted them instead. If there are other good matches, then market designs which steer searchers to those matches can improve market efficiency and platform profits.

Figure 10 displays summary statistics regarding potential trips where the searcher sends one initial inquiry. Of the 37% of searchers with a rejected first inquiry, 51% don't send another inquiry in the sample. Of those that do send another inquiry, 67% end up booking. On the other hand, those whose first inquiries are not rejected book at a 75% rate. In total, cases where an initial rejection is followed by the searcher leaving comprise 19% of trip attempts in this sample.

However, the association between rejection and a lack of subsequent booking may not be causal.[18] Consider the following thought exercise. Suppose that the listing whose host rejected the searcher was not shown to the searcher at all. If the effect of a rejection is causal, then there would be other suitable listings whose hosts would accept the searcher. On the other hand, if there were no such listings or if the searcher never intended to book the listing in the first place, then the association is spurious. Below, I study whether controlling for these potentially non-causal mechanisms affects the baseline estimates of the effect of a rejection.

Consider a potential guest-trip, g (a guest, market, and check-in time combination), sending a first inquiry to a host, h. Column (1) of Table 3 reports the results of the following OLS regression:

$$B_g = \beta_0 + \beta_1 r_{gh} + \varepsilon_{gh} \tag{1}$$

[18] Horton (2016) uses an instrumental variable technique to show that, in the setting of Odesk / Upwork, rejection by employees of invitations to apply by employers has a causal effect on the probability at which a job is eventually filled. In his paper, the two-stage least squared estimate is actually larger than the simple OLS estimate of the effect of a rejection.

where B_g is an indicator whether the guest-trip results in a booking and r_{gh} is an indicator for whether host, h, rejects the guest-trip, g. The simple estimate of the effect is -.41 compared to a mean booking rate of .63.

One potential confounder is that there is insufficient supply in the market for a given check-in week. In that case, even if the guest tried to find another listing, they would fail. Another potential confounder is that guest-trips which are rejected may have low intent or may be not be desirable for hosts. To control for both of these, I add a fixed effect for each week of inquiry and week of check-in combination. I also control for guest and trip characteristics (country of origin, number of guests, whether the guest had a prior booking, lead time, whether the guest is reviewed, and number of nights). In combination, these controls reduce the point estimate by .1 (Column (2)). A large part of this reduction is explained by the presence of a control for whether the guest has a prior booking. Guests with prior bookings have a 36 percentage point higher chance of booking conditional on sending an inquiry.

In column (3) I test whether the rejection effect is mitigated by past guest experience by interacting rejection with whether a guest has had a prior booking. The coefficient on the interaction is -.16, representing an additional 25% decrease in the booking rate for experienced guests. While the effect of a rejection is lower in percentage terms for experienced guests, the absolute size of the effect is bigger. This means that even guests who are ex-ante highly likely to book are deterred from doing so by a rejection.

I further explore whether the lack of available supply affects the rejection effect by limiting the sample to inquiries occurring more than 2 weeks away from the check-in date. For these inquiries, there is still a relatively large number of suitable listings visible in search results because they have not yet been booked. Column (4) displays the results of this specification, and the point estimate differs from the baseline by less than 5 percentage points.

In column (5), I limit the sample to contacts which used the 'Book It' button. In this case, the guests would have been committed to booking had the hosts accepted and therefore have a high intent to book. The coefficient remains the same in magnitude. Lastly, in column (6), I use a two-stage least squares strategy where I instrument for rejection with an indicator for whether the host blocked off at least one of the dates of the inquiry without being booked. The motivation for this instrument is that, while screening rejections may be due to undesirable guest-trip characteristics, stale vacancy rejections should be less related to undesirable guest-trip characteristics. Column (6) shows that the estimate from the 2SLS specification is -.45, even larger than from the OLS specifications.

To summarize, there is a large association between rejections and searchers leaving Airbnb without booking. This association persists and is similar in magnitude even in specifications that control for non-causal mechanisms that may result in this association. Therefore, rejection represents an important transaction cost in this market. In the next two sections, I use a model of searcher choice and host rejection to demonstrate how Airbnb's market design reduces transaction costs in the market.

5 The Role of Search Engine Design

In this section I estimate models of search and rejection and use them to study the role of search engine design. The goal of this section is to understand the relative contributions of different search engine features in determining rejections and the quantity of matches in the market.

To make the counterfactuals in this section concrete, consider the search engine design on Craigslist, which existed as a marketplace for urban short-term rentals before Airbnb. Figure 2 displays the results of a 2005 Craigslist search for vacation rentals in New York, which differ from Airbnb's among several dimensions. First, many of the listings displayed are not in New York, don't have a specified availability date, and don't have a standardized price per night. Second, Craigslist's search engine has limited functionality relative to Airbnb's. It does not include accurate filters for geography within a city or for the days of the trip.[19] A filter for price is available but is not fully functional because listed prices are not standardized per night. Furthermore, because transactions do not take place on the site, the listed prices do not necessarily reflect the actual prices which listings charge. Third, displayed listings are not automatically removed when they are booked and there is no availability tracking system on the site. Lastly, the results are displayed in chronological order rather than by relevance. My results demonstrate that transaction costs would be much larger in a Craigslist-like counterfactual, which may explain why it was not a successful marketplace for short-term urban rentals.

5.1 A Model of Searcher Choice

The goal of the searcher choice model is to predict which option a searcher will choose from the set of all options that are shown by the search engine during search. I model this choice using a random utility discrete choice model. The searcher's contact decision is a function of the property characteristics, searcher and search characteristics, and filtering choices. Conditional on these observables, the searcher chooses listings to contact. The most important difference between this choice model and standard discrete choice models is that I use the realized filter choices as proxies for otherwise unobservable idiosyncratic preferences for neighborhood, room type, and price. For example, if a searcher uses the map to filter for a particular neighborhood, I allow the searcher's choice probability to differ for listings in that neighborhood.

Denote each guest-trip (a combination of unique searcher, city, and trip dates) as g. Each g receives utility from property, h, according to a linear combination of property characteristics, interactions with idiosyncratic preferences, and a guest specific error term according to the equation below:

$$\begin{aligned} u_{gh} = \alpha_0 + p_{gh} * (FP_g'\alpha_1 + NFP_g'\alpha_1 + Z_g'\alpha_2) + f(X_{ht}, Z_g)'\beta \\ + \kappa_N + \phi_{FN} FN_{gh} + \phi_{FR} FR_{gh} + \phi_R R_{gh} + \gamma_h + \varepsilon_{gh} \end{aligned} \tag{2}$$

where X_{ht} is a vector of property characteristics including review quality, property type and whether the host controls multiple listings. Z_g is a vector of trip and guest characteristics (Nights,

[19] Although there are links for specific boroughs, these links also yield results which are not geographically limited.

Age, Guests, and a constant), p_{ght} is the nightly price of the property for the trip inclusive of platform fees and cleaning costs, FP_g is the maximum price filter used by the searcher (set to 0 if no price filter used), NFP_g is an indicator that takes the value of 1 if a price filter is used, $f(X_{ht}, Z_g)$ is a set of interactions between guest and host characteristics. κ_N is a neighborhood fixed effect, FN_{gh} is an indicator variable for whether a listing's neighborhood was specified by the searcher's filtering action, FR_{gh} is an indicator variable for whether a listing's room type was specified by the searcher's filter, R_{gh} is the lowest search rank a listing was shown at to the guest, and ε_{ght} is an unobserved component of the utility which is distributed according to the type 1 Extreme Value (EV) distribution.

The searcher can also choose the outside option and leave the platform without sending a contact. The searcher's value of the outside option is determined by the following equation:

$$u_{go} = T_g'\mu + \gamma_{FP} + \phi_H H_g + \varepsilon_{go} \tag{3}$$

where T_{go} are guest and trip characteristics, γ_{FP} is a set of fixed effects for filters used by the searcher, H_g is the number of unique listings shown by the search engine to the guest and ε_{go} is a type 1 EV error term.[20]

Before moving to the estimation, it is important to clarify the interpretation of this model. The main purpose of this model in this paper is for prediction. The estimated parameters only have a structural interpretation if additional and highly restrictive assumptions are made (see Appendix A).

The data used to estimate the searcher model is described in subsection 3.1. I further reduce this sample in several ways. First, I remove searchers whose cumulative searches load fewer than 18 unique listings, the maximum number of search results potentially displayed on the page. Second, I limit searches to those that occur within 60 days of check-in. Lastly, just for the purposes of this estimation, I include the chosen option as well as random sample of up to 20 other options in the

[20]This term can be interpreted in two ways. First, it controls for unobserved heterogeneity in a searcher's returns from searching. Second, it serves as an analogue to the procedure in Ackerberg and Rysman (2005), which mitigates the tendency of discrete choice models to overstate the benefits of variety.

estimation procedure. This sampling procedure reduces the computational time of the estimation procedure and retains consistency (see Train (2009) and Wasi and Keane (2012) for details).

First, it is worth discussing the predictive accuracy of the model. The searcher in the model faces many options that were selected into the consideration set by the platform due to having desirable characteristics and by the searcher while filtering. Furthermore, there is relatively little information in the data on searcher preferences other than the query parameters. Consequently, this is a difficult prediction problem. The model predicted probability that a given searcher chooses the listing they were observed to have chosen is 10% conditional on not choosing the outside option. The chosen listing is in the top 5 options ranked by predicted choice probability 60% of the time.

The results of the estimation procedure are displayed in Table 4. The estimates are for the most part consistent with prior intuition regarding listing quality. First, with regard to reviews, the average rating and total number of five star reviews are predictive of choice. Second, entire properties and listings with lower search ranks are more likely to be chosen. Third, the outside option is more likely to be chosen when searches have fewer guests, when searcher's don't filter for price, and when the search is further away from the check-in date. Interestingly, listings that allow instant booking are less likely to be chosen. This likely reflects the fact that, at least during the sample period, listings that allowed instant book were of a lower quality than those that did not. Furthermore, it suggests that searchers do not strongly respond to the probability of rejection when choosing listings to book. Otherwise, we would expect the coefficient on an instant bookable listing to be positive and large.

The filtering behavior of searchers is highly predictive of choice. Listings which are in a neighborhood that the searcher filtered for have a $178 additional value to searchers. Similarly, listings which are of the property type that is filtered for, are valued an additional $135 by searchers. The price filtering behavior is also predictive of choice. The higher a searcher's maximum price filter, the less sensitive they are to the prices of listings and searchers who use a price filter are less likely to pick the outside option.

5.2 A Model of Host Screening

In order to model counterfactual scenarios in which the consideration sets of searchers change, I need to be able to predict when contacts will be rejected by hosts. Subsection 4.1 describes three reasons why rejection occurs in this market: congestion, stale vacancies, and screening. Of these, congestion and stale vacancies occur in a manner that is unrelated to host preferences. On the other hand, screening rejections occur because hosts have preferences over when and whom they host. For example, a host might reject a contact because the guest is not reviewed, has a vague inquiry, or does not have enough information in his profile. Hosts also reject guests because the check-in dates of the inquiry can break up a bigger, uninterrupted time of availability for the host, preventing future inquiries. Lastly, hosts may be waiting for a better guest/trip combination or might not be willing to take a particular guest for idiosyncratic reasons.[21] In this section, I describe a simple model of the decision to reject as a function of guest, trip, and listing characteristics.

The estimating equation for the screening model is:

$$Pr(R_{gh}) = Pr(\alpha_0 + Z_h'\delta + f(X_g, Z_h)'\beta + \eta_{gh} > 0) \tag{4}$$

where η_{gh} is a logit error term, R_{gh} is an indicator for whether the response is a rejection, X_g are the number of guests, guest reviews, guest gender, weekly demand, days in advance of the trip nights, guest age, searcher census tract demographics and month of check-in. Z_h are property type, multi-listing host indicator, host age, the number of reviews and price. $f(X_g, Z_h)$ are interactions between guest and listing characteristics.

I account for the dynamic aspects of the host decision by controlling for the time in advance of the trip of inquiry and for the overall demand for each week of check-in. A more complete model of a host's decision to accept or reject would require the host to have expectations over the flow and

[21] For example, Edelman, Luca and Svirsky (2016) use an audit study to show that some hosts discriminate against non-reviewed guests with African-American names. I do not observe race in my sample and cannot consequently control for it in this regression. However, since minority applicants are a minority of site users, this omitted variable in unlikely to be driving my results. Furthermore, even though the audit study applicants had no reviews, pictures, or profile descriptions, they were still frequently accepted by hosts.

quality of potential future searchers in the market. I choose not to add this additional complexity for two reasons. First, inquiries with longer lead times are actually less likely to be rejected. This suggests that option value concerns are not the primary driver of rejection behavior. Second, screening rejections represent just one reason for rejection in this framework and the option value component of this decision is unlikely to influence the later results in a qualitatively important manner.

The dataset for estimation is described in section 4. From this dataset, I further select only contacts that were not rejected due to congestion or stale vacancies. This leaves me with 93,851 observations of contacts. Table 5 displays the estimates from the above specification without random effects in column (1), with listing random effects in column (2), and with listing fixed effects in column (3). The results are similar across all three models. Guests who send an inquiry are more likely to be rejected. This is a function of two factors. First, guests who send 'book it' requests are more committed to booking and are likely to be accepted by hosts. Second, some 'book it' requests are instant bookings, which means that they are guaranteed to be accepted.

Guest reputation affects host decisions. Guests with a prior review are less likely to be rejected.[22] The extent to which reviews are valued by hosts varies across host types. Hosts who have multiple-listings and are less likely to value the social aspect of the Airbnb transaction show no statistically significant differences in rejection behavior between reviewed and non-reviewed guests. Other types of guest information including Airbnb verification of identity and profile descriptions are also associated with lower guest rejection rates.

Trip characteristics also affect the decisions of hosts. Trips with more guests are less likely to be rejected by bigger listings. Furthermore, when the number of guests equals the capacity of the host, the host is more likely to reject. The number of nights is negatively correlated with rejections (conditional on not being rejected due to congestion or a stale vacancy). All else equal, hosts prefer longer trips to shorter trips and longer lead times.

On the listing side, hosts with entire properties are more likely to reject and hosts that have

[22] I exclude variables related to guest ratings because these are not seen by the host when a guest submits an inquiry and because hosts recommend guests 99% of the time (Fradkin, Grewal and Holtz (2017)).

enabled instant booking are less likely to reject. For qualifying guests and trips, instant booking guarantees that a booking request will be accepted. The rejection rate for instant-book enabled listings is not exactly zero for two reasons. First, not all guests are eligible for instant booking. Second, guests may choose to send an inquiry to the host first and the host may respond with a rejection. There is large heterogeneity in rejection probabilities across hosts. The standard deviation of the host random effect in column (2) is 1.093, which is larger in magnitude than any of the estimated coefficients. These heterogeneous rejection probabilities will be important in the section regarding potential improvements to the ranking algorithm. The reason is that the platform can potentially use the ranking algorithm to redirect searchers to listings with lower rejection probabilities.

5.3 The Effects of Search Engine Design

In this section, I use the previously estimated choice and rejection models (with listing random effects) to show that the probability of searchers choosing the outside option and being rejected would both drastically increase if Airbnb's search engine was degraded relative to it's state during the sample period.

I make several simplifying assumptions regarding searcher and host behavior in calculating expected outcomes. First, with regards to searchers, I keep searcher and search characteristics the same as in the data, and assume that searchers see the same number of unique listings while searching. While searchers would adjust their search intensity to changes in the search technology, this is unlikely to reverse my findings. Figure 5 displays the trends in the number of inquiries between 2013 and mid-2015 for the US. Even as the rejection rate falls and the booking rate increases, the 50th and 75th percentile of weekly inquiries by users does not substantially change. There is a decease in the number of contacts in the 95th percentile of the distribution, but this 5% of most intense searchers would not reverse my main findings.

With regards to hosts, I assume that their screening rejection function remains the same across counterfactual scenarios. While this is not fully realistic – hosts may become less selective if there are fewer contacters — it is unlikely to change the qualitative results of the counterfactuals. The are two reasons for this. First, in counterfactuals where the search engine does not track availability, rejection is not typically caused by screening. Second, even if some hosts reduce their screening intensity, heterogeneity in host rejection rates due to host preferences will still persist.

Lastly, my counterfactuals are in partial equilibrium, meaning that the number, types, and prices of market participants do not change across counterfactuals. For each entering searcher, I assume that the past and future availability of a listing, as well as its characteristics, remains as it was in the data. As explored in Fradkin (2015), there are dynamic effects of marketplace policies. If a policy change induces an additional listing to be booked today, that means that later searchers will not be able to book that listing tomorrow. Furthermore, the communication choices of searchers and hosts determine the level of congestion in the market.[23]

These equilibrium effects are unlikely to affect the baseline conclusion of these counterfactuals — that without availability tracking and filtering, searchers will see worse listings, of which a large share will be unavailable. Without information on availability, guests would contact these unavailable listings and would leave the platform when rejected.

I first calculate expected outcomes in an approximation to the market in the status quo. To do this, I use each searcher's realized set of browsed listings as a choice set, and calculate expected outcomes using the model predicted probabilities of sending a contact and being rejected by the host who received the contact.[24] Rejections in this scenario can happen for three reasons. First, the listing's availability can be stale for the contact dates. Second, the host can choose to reject the contact due to guest or trip characteristics (screening). Third, the listing may have been booked or declared unavailable prior to the search. In the "Status Quo" scenario the third reason for rejection is excluded because Airbnb automatically removes these listings from search.

[23] One possibility is that hosts will reduce prices in response to a worse marketplace design. However, Farronato and Fradkin (2018) show that many hosts are at the margin of participation giving the prevailing market prices and only transact in peak demand periods. Consequently, there is not much room for most hosts to decrease prices in this market.

[24] I make several additional simplifying adjustments. First, some listings shown in search were already booked or marked as unavailable at the time the contact was sent. Since such listings are not present in counterfactual scenarios, I remove these listings from a searcher's choice set and re-sample from the set of listings that are still marked as available on the calendar so that the total number of listings seen by each searcher is the same as in the data. Second,

Row (1) of Table 6 displays the results corresponding to this scenario. As in the data, approximately 36% of searchers choose to send at least one inquiry. Furthermore, 32% of those inquiries are rejected, which is several percentage points less than the 36.5% rejection rate observed in the data. This is expected because congestion rejections are not a part of these counterfactuals. Consequently, 76% of searchers either choose to leave the platform before sending an inquiry or are rejected in their first inquiry.

The set of browsed listings under the status quo is a function of several features. First, the platform displays search results according to a ranking algorithm. Second, the searcher used the search engine filters to find listings closer to his or her preferences. Lastly, the search engine automatically removed previously booked or unavailable listings from search. In the first counterfactual, I assume that the search results are instead randomly drawn from the set of all active listings in the market, regardless of their availability. Row (2) shows that there is a 2 percentage point decrease in the share of searchers who send a contact. This small decrease is due to the fact that while the search is less directed, the previously booked listings are typically of higher quality than those remaining in the market. In contrast to the contact rate, the rejection rate more than doubles from 32% to 78%. This increase comes primarily from the fact that 22% of searchers are now rejected because the listing was previously booked. Furthermore, there is also an increase in screening rejections because the previously booked listings also tend to be more selective.

In aggregate, the features of filtering, ranking, and availability tracking combine to increase the rate of searchers with accepted first contacts on Airbnb from 7.7% to 24%. Furthermore, if we assume that the discrete choice model provides a valid estimate of utility, the expected utility from booking for those that are accepted falls from $155 to $140 per night. This means that under random search, searchers find worse matches, even when they are successful in finding a transaction partner.

In row (3), I consider what occurs when search is still random but the previously booked and

since congestion rejections are relatively unimportant and are a property of multiple interacting searchers, I abstract away from these in the counterfactuals. Lastly, I draw each listing's minimum rank for the choice model according to the empirical distribution of minimum ranks in the data.

unavailable listings can no longer be seen in the search results. Under this scenario, there is a 5 percentage point decrease in the contact rate and an 11 percentage point increase in the rejection rate. This results in a decrease of 6.1 percentage points in the share of searchers who send a first contact which is accepted. This decrease is much smaller than the decrease in the prior scenario, suggesting that availability tracking is relatively more important in the functioning of this market than filtering. Interestingly, the expected utility for an accepted contact is even smaller in this scenario than in the scenario without availability tracking. The reason for this is that when the previously booked listings are loaded in the search results, available listings must have a high enough utility to compete with the previously booked listings. In contrast, when the previously booked listings are not visible in the search results, only lower utility listings are left for the searchers to contact.

I've shown that two search engine features greatly affect the probability of a successful search on the Airbnb marketplace. If the marketplace did not keep track of availability, then rejection rates for contacts regarding bookings would increase to 78%. If anything, this provides an underestimate of the true effects of a laissez-faire marketplace design. The Airbnb marketplace as of 2014 was already curated. Many listings with high rejection rates or low quality were either manually removed by Airbnb or endogenously left due to their lack of competitive success in the market. In the next section, I study the potential for ranking algorithms to further improve market outcomes.

6 The Potential for Improvement in Matching

Even with the marketplace design in 2014, many searchers either chose the outside option or were rejected in their communications with hosts. This outcome may be efficient from the perspective of the platform in two cases. First, if there are no listings in the marketplace suitable to the searcher, then the platform cannot improve the outcome of that searcher other than by adding more suitable listings. Second, if the platform could not predict a rejection, then the only way to discover the availability of a listing to a searcher would be through communication. However, if either of these conditions fail then the platform could potentially improve matching through search ranking or other types of marketplace curation.

In this section I study the potential for improved matching. To do this, I use my models of search and rejection to derive rankings of listings and I compute what happens in the market when searchers see alternative consideration sets based on these rankings. I consider three types of rankings calculated according to the following equations:

1. $w_{h,a} = \sum_h \bar{\mu}_{gh}$ (Average Quality)

2. $w_{gh,p} = \bar{\mu}_{gh}$ (Personalized Quality)

3. $w_{gh,t} = \bar{\mu}_{gh} * (1 - Pr(R_{gh}))$ (Rejection Weighted)

Ranking 1 is a measure of the average utility a listing provides to searchers in the sample. This would be the easiest ranking to implement since it requires no particular information about a specific searcher's preferences.[25] Rows (4) and (5) of Table 6 display the market outcomes if searchers saw the same sized consideration set as they do in the data, but that consideration set was picked according to the ranking. Row (4) shows the results if the platform did not keep track of availability. First, the average quality of the listings that the searchers see greatly improves. Consequently, the share of searchers who send a contact increases by 21 percentage points and the expected utility of an accepted contact increases to $372 per night. However, because availability is not tracked, 86% of contacts are rejected and the total share of searchers with accepted first contracts is just 8%. This demonstrates that without availability tracking, better ranking only has a limited effect on market outcomes. With availability tracking, the market outcomes, shown in row (5), do improve. The share of searchers with an accepted contact increases by 3.3 percentage points. However, the probability of rejection is still higher relative to the status quo. This occurs because better listings tend to be more selective and consequently reject more inquiries. This is likely to be an underestimate of the true effect on rejections because the highly correlated rankings among searchers should also lead to more congestion rejections.

[25]In contrast, a ranking algorithm based on the filters used in the process of search requires a more sophisticated technical infrastructure and set of algorithms.

Ranking 2 uses the realized searcher characteristics and filters to calculate a personalized utility estimate for each searcher and listing combination. This likely represents an upper bound on the benefits of personalization because information on which filters the searcher applies are not available until after the search. Columns (6) and (7) display the results when this ranking is used to form each searcher's consideration set. Row (6) shows that, as in the average quality ranking, without availability tracking the increase in rejection rates overwhelms the benefits of personalization.

Row (7) shows that when availability is tracked, the personalized algorithm decreases the rate at which searchers choose the outside option from 64% to 45%. While the rejection rate does increase relative the status quo, the increase is not as large as in the case of the average utility ranking. The personalized ranking yields a 3 percentage point increase in searchers who send a contact and are accepted by their top choice. Lastly, the expected utility from an acceptance increases by $79 relative to the average quality algorithm and $164 relative to the status quo.

Neither ranking 1 nor 2 uses the information on the screening propensities of hosts. Ranking 3 explores the possibility of weighting the expected utility from a listing by the host's probability of rejecting a guest due to screening. Such a ranking trades off listing quality for a lower chance of rejection. The results of this counterfactual are displayed in row (8). The share of searchers who choose the outside option increases by 1.6 percentage points relative to the personalized ranking that maximizes expected utility conditional on an acceptance. At the same time, the expected rejection rate falls from 44% to 36%. The cumulative effect of this policy is that the share of searchers with an accepted first contact increases by 2.9 percentage points relative to the prior ranking. This confirms that many searchers who are rejected leave because of incurred or predicted future transaction costs rather than a lack of attractive listings.

Lastly, row (9) displays the results from an alternative search policy where any option with an expected screening rejection probability greater than .45 is removed from the results. This policy slightly reduces the rejection probability but does lead to worse options for searchers — the expected utility from an accepted contact drops from $282 to $206. Therefore the policy of removing high rejection results from search preforms worse than the rejection weighted ranking.

The results from these counterfactuals suggest there is a large opportunity for ranking to improve market efficiency relative to the status quo in 2014. Rankings based on the expected quality of the match between the searcher and the listing preform especially well in these counterfactuals. Indeed, informed by earlier versions of this analysis, subsequent policy by the platform has focused on matching in ranking algorithm design.

More generally, an optimal ranking algorithm would consider searcher and host match utilities, the disutility to searchers from rejection, the benefits of screening for hosts, and the costs to hosts of maintaining calendar accuracy. An important and interesting question for the platform is how to choose the relative utility weights across these market outcomes. Algorithmic design choices become even more complex when considering their equilibrium effects. For example, hosts who value the ability to select guests may be relatively disadvantaged by algorithms which redirect search effort to less picky hosts. An algorithm that penalizes rejection may encourage hosts to keep more accurate calendars, to be more precise about the desired guest and trip characteristics, and to become less selective. If the platform reduces average searcher transaction costs sufficiently through a new algorithm, this may improve outcomes for all hosts on the platform through an increase in aggregate demand.

7 Discussion and Implications for Platform Design

Decentralized search and matching markets suffer from a variety of frictions which result in transaction costs. Dating back to Coase (1937), social scientists have emphasized how these costs affect the optimal structure of production and exchange throughout the economy. In this paper, I've shown how the combination of digital technology and marketplace design helps previously high transaction cost modes of exchange to compete with the more centralized forms of exchange which comprise the majority of transactions in the accommodation industry.

In the context of Airbnb, the process of transacting is complicated by the presence of large and heterogeneous choice sets as well as uncertainty regarding the availability of an option. Airbnb's marketplace design, which tracks availability and offers precise filters, greatly reduces these costs. Without these features, searchers would need to expend much more effort to find suitable matches in this market. This reduction in transaction costs is especially important given the fact that searchers have the outside option of booking a hotel room while incurring much lower transaction costs.

I've also shown how the marketplace can use data on historical user behavior to predict match quality and generate better rankings. Subsequent to this research, Airbnb announced a change in the ranking algorithm called 'Host Preferences,' which uses a model predicting host rejections to rank listings. More broadly, the marketplace can design interfaces which elicit useful information on both buyer and seller preferences. The data generated by these interfaces can be used to create better matches.

For example, Airbnb has advocated for the increased usage of 'Instant Book' by hosts throughout the platform. Instant booking reduces the rejection problem by making listings visible in search available by default. However, it entails a trade-off since hosts give up some of their ability to choose between guests. In order to increase host adoption, Airbnb has used both incentives and product features. With regards to incentives, hosts who use instant book receive advantageous search engine placement and a boost towards 'Superhost' status.[26] Furthermore, Airbnb has added controls for hosts to specify which guests and trips can instant book. For example, hosts can choose to let only experienced and well-reviewed guests instant book. Instant booking has grown from accounting for fewer than 10% of bookings in my sample to 60% of bookings in 2017. Similar mechanisms can likely be used to improve the efficiency of other peer-to-peer digital marketplaces and to create marketplaces in new verticals.

I also showed that stale vacancies account for a large share of rejections on Airbnb. These stale vacancies are possible even for hosts who have enabled instant book because these hosts may cancel due to a lack of availability. In 2018, Airbnb announced a successful market design intervention called 'contextual calendar reminders,' that used data on host calendar checking behavior

[26] https://www.airbnb.com/help/article/523/what-is-instant-book

to target email reminders to likely inattentive hosts. This intervention highlights are more general point — that by tracking usage, a digital platform can infer whether particular users are currently active and available. This principle has applications to a variety of markets including the labor, dating, and housing markets.

One important issue that this paper does not discuss is the dynamic mis-allocation of matches. For example, earlier searchers may book highly desirable listings, which would have been better matched with later searchers. Economists typically assume that the price mechanism will solve these issues, but in a complex marketplace where demand is volatile and the market structure is evolving, the assumption of optimal pricing by sellers is unlikely to hold. Instead, the marketplace possesses much more information than individual sellers and this information can be used for pricing or other matching mechanisms. Indeed, marketplaces such as Airbnb, eBay, and Uber are experimenting with novel price-setting mechanisms that rely on large scale data regarding demand, supply, and user behavior. New theories and empirical methods are needed to understand the implications of these algorithmic policies on platform profits, market efficiency, and equity.

An increasing share of activity in the labor, housing, and dating markets is being conducted through digital platforms. As a result, I anticipate that the issues raised in this paper will continue to grow in importance.

References

Ackerberg, Daniel A., and Marc Rysman. 2005. "Unobserved Product Differentiation in Discrete-Choice Models: Estimating Price Elasticities and Welfare Effects." *RAND Journal of Economics*, 36(4): 771–788.

Albrecht, James, Bruno Decreuse, and Susan Vroman. 2017. "Directed Search with Phantom Vacancies."

Albrecht, James, Pieter A. Gautier, and Susan Vroman. 2006. "Equilibrium Directed Search with Multiple Applications." *The Review of Economic Studies*, 73(4): 869–891.

Arnosti, Nick, Ramesh Johari, and Yash Kanoria. 2014. "Managing Congestion in Decentralized Matching Markets." *EC '14*, 451–451. New York, NY, USA:ACM.

Bronnenberg, Bart, Jun B. Kim, and Carl Mela. 2016. "Zooming in on Choice." *Marketing Science*.

Burdett, Kenneth, Shouyong Shi, and Randall Wright. 2001. "Pricing and Matching with Frictions." *Journal of Political Economy*, 109(5): 1060–1085.

Chen, Yuxin, and Song Yao. Forthcoming. "Sequential Search with Refinement: Model and Application with Click-Stream Data." *Management Science*.

Cheron, Arnaud, and Bruno Decreuse. 2017. "Matching with Phantoms*." *The Review of Economic Studies*.

Coase, Ronald H. 1937. "The Nature of the Firm." *Economica*, 4(16): 386–405.

Cullen, Zoë, and Chiara Farronato. 2016. "Outsourcing Tasks Online: Matching Supply and Demand on Peer-to-Peer Internet Platforms." *Manuscript*.

De Los Santos, Babur, Ali Hortaçsu, and Matthijs R. Wildenbeest. 2015. "Search with Learning for Differentiated Products: Evidence from E-Commerce." *Journal of Business & Economic Statistics*, 00–00.

De los Santos, Babur, and Sergei Koulayev. 2017. "Optimizing Click-Through in Online Rankings with Endogenous Search Refinement." *Marketing Science*.

Dinerstein, Michael, Liran Einav, Jonathan Levin, and Neel Sundaresan. 2014. "Consumer Price Search and Platform Design in Internet Commerce." National Bureau of Economic Research Working Paper 20415.

Edelman, Benjamin G., Michael Luca, and Dan Svirsky. 2016. "Racial Discrimination in the Sharing Economy: Evidence from a Field Experiment." *Manuscript*.

Einav, Liran, Chiara Farronato, and Jonathan Levin. 2016. "Peer-to-Peer Markets." *The Annual Review of Economics*.

Einav, Liran, Chiara Farronato, Jonathan Levin, and Neel Sundaresan. Forthcoming. "Auctions versus Posted Prices in Online Markets."

Ellison, Glenn, and Alexander Wolitzky. 2012. "A Search Cost Model of Obfuscation." *The RAND Journal of Economics*, 43(3): 417–441.

Ellison, Glenn, and Sara Fisher Ellison. 2009. "Search, Obfuscation, and Price Elasticities on the Internet." *Econometrica*, 77(2): 427–452.

Farronato, Chiara, and Andrey Fradkin. 2018. "The Welfare Effects of Peer Entry in the Accommodations Market: The Case of Airbnb." *NBER Working Paper*.

Fradkin, Andrey. 2015. "Search Frictions and the Design of Online Marketplaces."

Fradkin, Andrey, Elena Grewal, and David Holtz. 2017. "The Determinants of Online Review Informativeness: Evidence from Field Experiments on Airbnb."

Fraiberger, Samuel P., and Arun Sundararajan. 2015. "Peer-to-Peer Rental Markets in the Sharing Economy." Social Science Research Network SSRN Scholarly Paper ID 2574337, Rochester, NY.

Gavazza, A. 2016. "An Empirical Equilibrium Model of a Decentralized Asset Market." *Econometrica*.

Ghose, Anindya, Panagiotis G. Ipeirotis, and Beibei Li. 2014. "Examining the Impact of Ranking on Consumer Behavior and Search Engine Revenue." *Management Science*, 60(7): 1632–1654.

Hall, Jonathan, Cory Kendrick, and Chris Nosko. 2016. "The Effects of Uber's Surge Pricing: A Case Study."

Hitsch, Gunter J., Ali Hortaçsu, and Dan Ariely. 2010. "Matching and Sorting in Online Dating." *American Economic Review*, 100(1): 130–163.

Horton, John J. 2016. "Misdirected Search Effort in a Matching Market: Causes, Consequences and a Partial Solution." *Working Paper*.

Horton, John J. 2017. "The Effects of Algorithmic Labor Market Recommendations: Evidence from a Field Experiment." *Journal of Labor Economics*, 35(2).

Horton, John J., and Richard J. Zeckhauser. 2016. "Owning, Using and Renting: Some Simple Economics of the" Sharing Economy"." National Bureau of Economic Research.

Ke, T. Tony, Zuo-Jun Max Shen, and J. Miguel Villas-Boas. 2016. "Search for Information on Multiple Products." *Management Science*.

Kircher, Philipp Albert Theodor. 2009. "Efficiency of Simultaneous Search." *Journal of Political Economy*, 117(5): 861–913.

Klein, Tobias J., Christian Lambertz, and Konrad O. Stahl. Forthcoming. "Market Transparency, Adverse Selection, and Moral Hazard." *Journal of Political Economy*.

Lewis, Gregory. 2011. "Asymmetric Information, Adverse Selection and Online Disclosure: The Case of eBay Motors." *The American Economic Review*, 101(4): 1535–1546.

los Santos, Babur De, and Sergei Koulayev. Forthcoming. "Optimizing Click-through in Online Rankings for Partially Anonymous Consumers." *Marketing Science*.

Nosko, Chris, and Steven Tadelis. 2015. "The Limits of Reputation in Platform Markets: An Empirical Analysis and Field Experiment." National Bureau of Economic Research Working Paper 20830.

Piazzesi, Monika, Martin Schneider, and Johannes Stroebel. 2015. "Segmented Housing Search." National Bureau of Economic Research Working Paper 20823.

Resnick, Paul, Ko Kuwabara, Richard Zeckhauser, and Eric Friedman. 2000. "Reputation Systems." *Commun. ACM*, 43(12): 45–48.

Roth, Alvin E. 2008. "What Have We Learned from Market Design?" *The Economic Journal*, 118(527): 285–310.

Skift. 2017. "Airbnb Ramps Up Push to Get More Hosts to Choose Instant Booking." *https://skift.com/2017/06/14/airbnb-ramps-up-push-to-get-more-hosts-to-choose-instant-booking/*.

Tadelis, Steven, and Florian Zettelmeyer. 2015. "Information Disclosure as a Matching Mechanism: Theory and Evidence from a Field Experiment." *American Economic Review*, 105(2): 886–905.

Train, Kenneth E. 2009. *Discrete Choice Methods with Simulation.* Cambridge; New York:Cambridge University Press.

Ursu, Raluca Mihaela. 2018. "The Power of Rankings: Quantifying the Effects of Rankings on Online Consumer Search and Choice." *Marketing Science*, Forthcoming.

Varian, Hal R. 2010. "Computer Mediated Transactions." *The American Economic Review*, 100(2): 1–10.

Wasi, Nada, and Michael P. Keane. 2012. "Estimation of Discrete Choice Models with Many Alternatives Using Random Subsets of the Full Choice Set: With an Application to Demand for Frozen Pizza." Economics Group, Nuffield College, University of Oxford Economics Paper 2012-W13.

Wolthoff, Ronald P. 2018. "Applications and Interviews: Firms' Recruiting Decisions in a Frictional Labor Market." *The Review of Economic Studies*, 85(2).

Figures

Figure 1: Search View

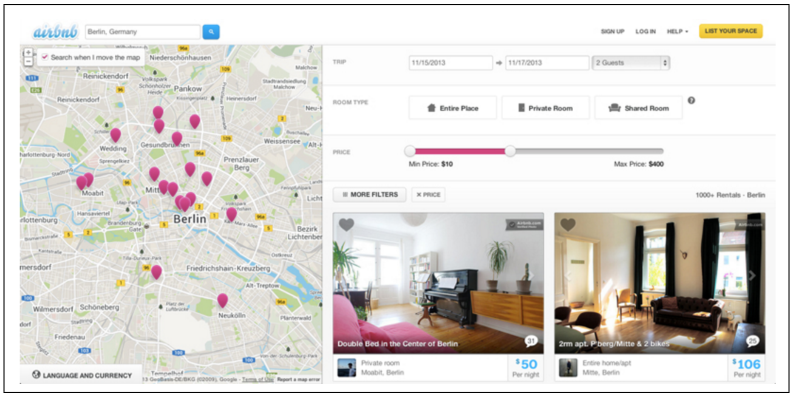

This figure displays the results of a search in Berlin for November 15, 2013 to November 17, 2013. Selecting filters or moving the map changes the set of displayed results. The searcher can scroll the page to see 21 listings before she is prompted to go to the next page of results. The results are displayed according to a ranking algorithm which is common across all searchers who input those search parameters.

Figure 2: Craiglist Vacation Rentals Search - New York, 2005

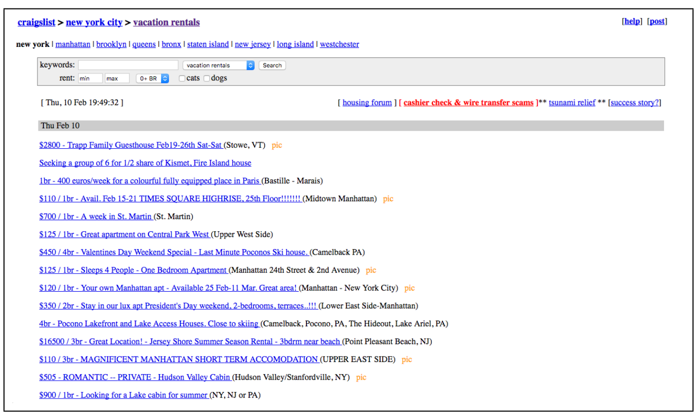

This figure displays the search results from a Craigslist search for vacation rentals in New York City on February 10th, 2005. Note that the first result is not located in New York, that prices are not standardized, and that listed availability dates vary across search results. "The Internet Archive" (https://archive.org/web/) was used to obtain these results.

Figure 3: Listing View

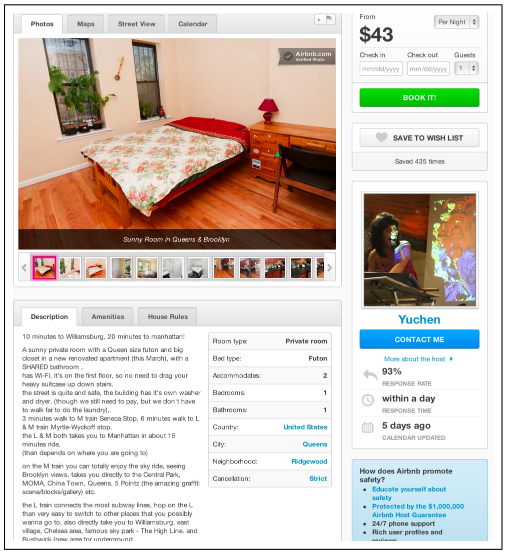

A searcher who clicks on a listing in the search results sees the following view. The ratings and text of reviews for the listing are shown lower on the page.

Figure 4: Inquiry Submission Form

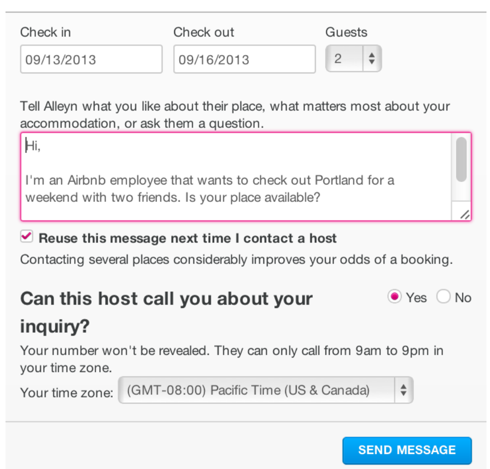

The figure above displays the prompt that searchers see when they click the "Contact Me" button.

Figure 5: Long Run Trends in Search Outcomes

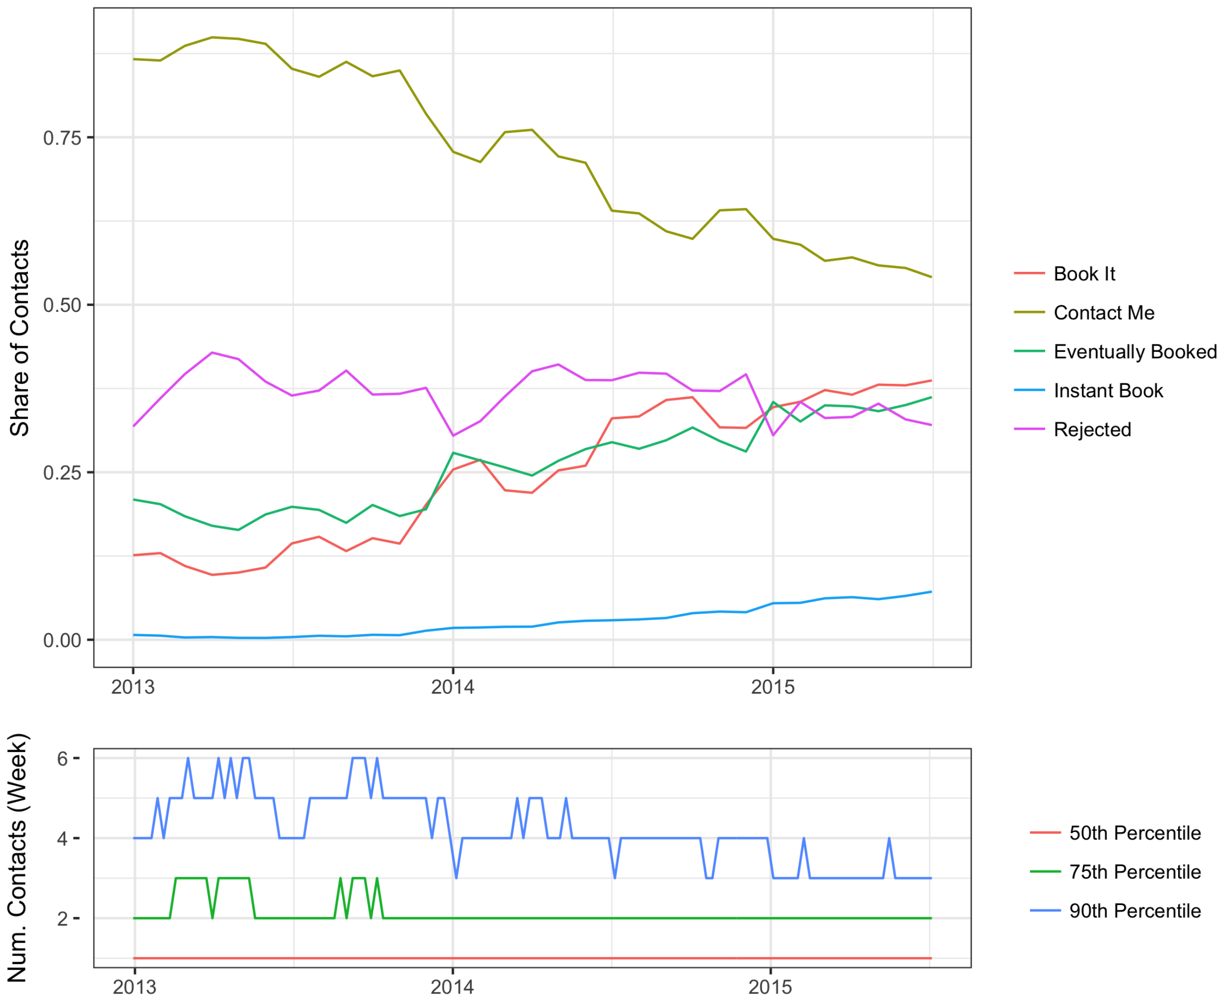

This figure displays the booking outcomes for users who sent one initial contact to the US between 2013 and July 8, 2015. The top figure displays the method of contact ('Contact Me', 'Book It', and 'Instant Book'), as well as the share of inquiries booked ('Eventually Booked') and rejected ('Rejected'). The lower figure displays the 50th, 75th, and 90th percentile of weekly contacts across users.

Figure 6: Map Filtering Frequencies

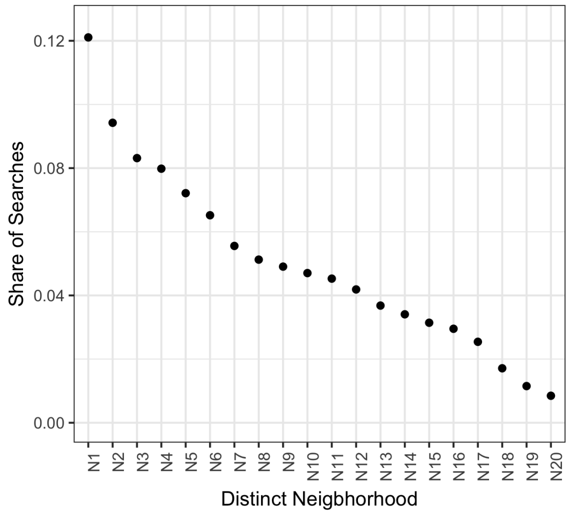

The figure displays the share of searches that included a non-default map location for a given neighborhood. Each point represents a distinct neighborhood.

Figure 7: Maximum Price Filter Distribution

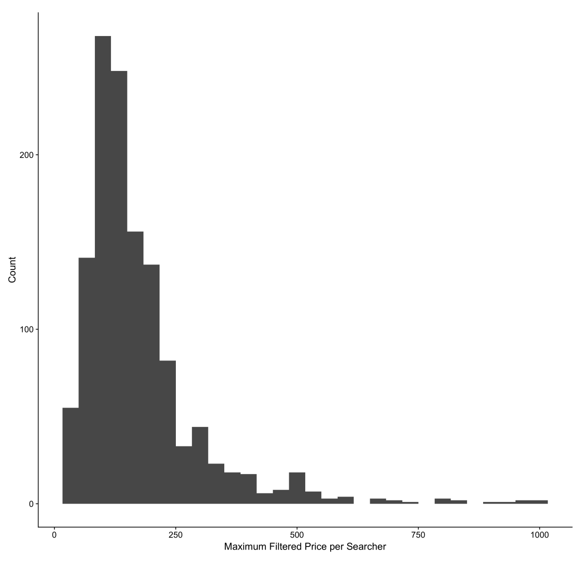

The figure displays the distribution of the maximum of the applied price filters per user for those users who used the filter at least once in the sample.

Figure 8: Number of Inquiries Sent

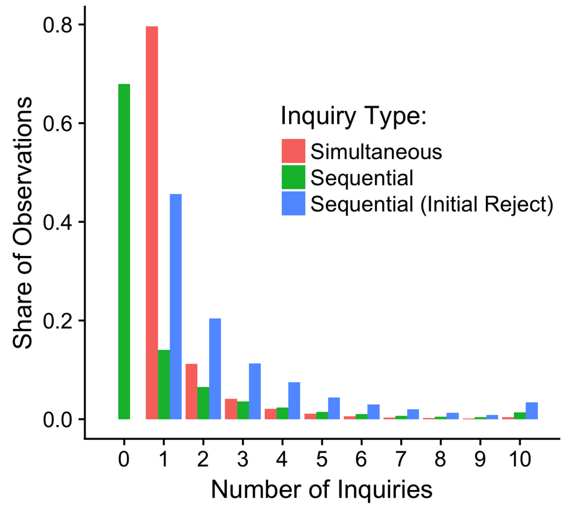

The figure displays the censored distribution of inquiry counts for the searchers in the sample. 'Simultaneous' refers to the number of inquiries sent within the first 2 hours of sending an inquiry. 'Sequential' refers to the number of inquiries sent after the first 2 hours. 'Sequential (Initial Reject)' refers to the subsample of guests who sent one simultaneous inquiry and were rejected.

Figure 9: Response Time Distribution (Hours)

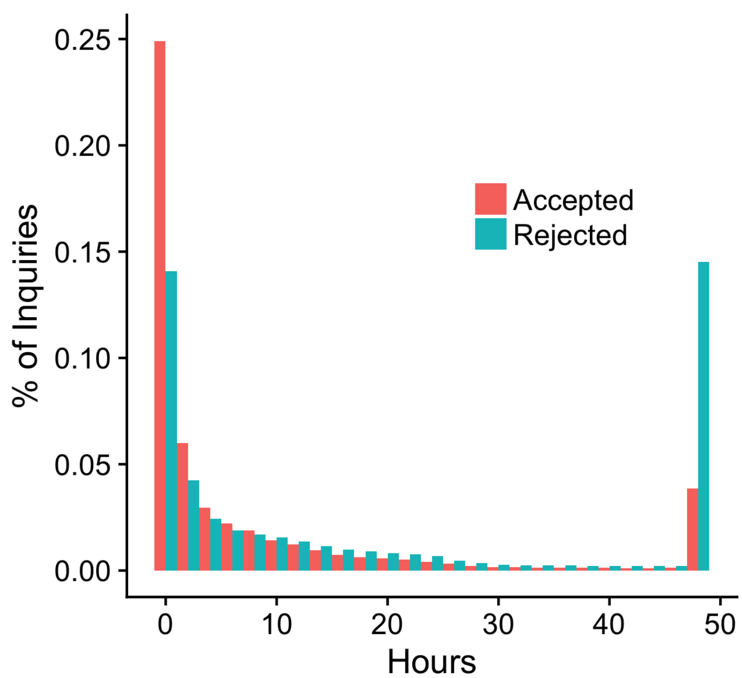

The figure displays the distribution of the response times for inquiries sent to listings in the United States during the sample period. Non-responses and responses which took longer than 48 hours are censored at 48 hours. The distributions are plotted separately for inquiries which were accepted and rejected by the host.

Figure 10: Booking and Rejection Outcomes

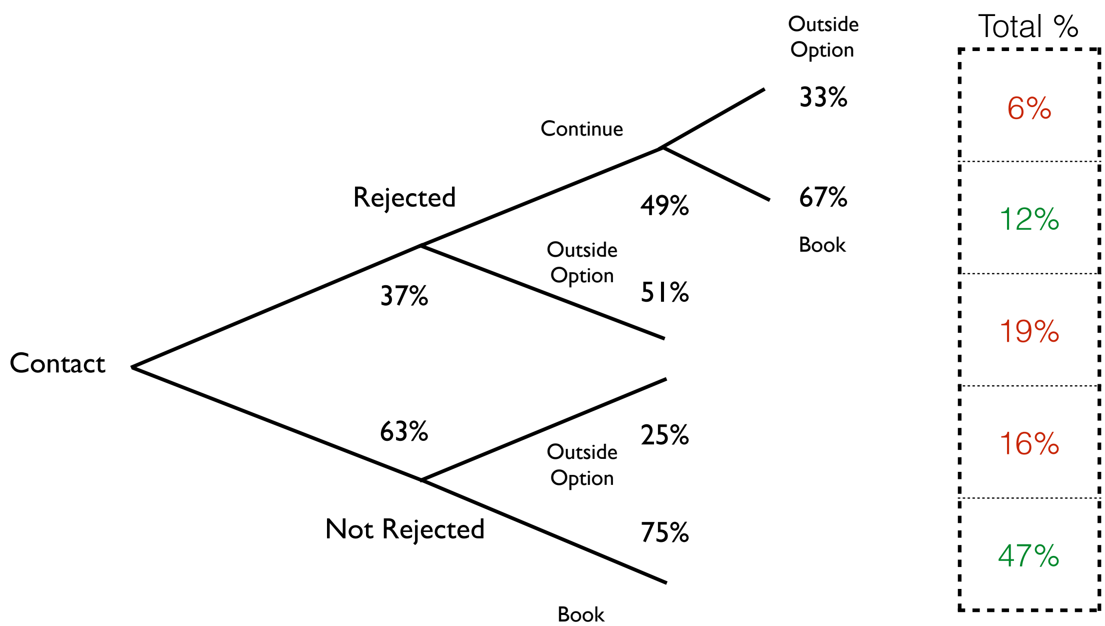

This figure displays the booking outcomes for users who sent one initial contact in the sample. Each number on the tree represents the probability conditional on reaching the prior step. The final column displays the unconditional probabilities of each outcome. ‘Outside Option’ occurs when the searcher does not make a booking of any listing in the market for the date including and close to the dates of the initial inquiry.

Tables

Table 1: Descriptive Statistics per Searcher

(a) All Searchers

Statistic	N	Mean	Pctl(25)	Median	Pctl(75)
Number of Searches	12,241	19.177	3	9	21
Unique Listings Seen	12,241	68.530	31	54	88
Share of Listings Seen	12,241	0.097	0.020	0.042	0.089
Num. Nights Requested	12,241	2.836	2	3	4
Num. Guests Requested	12,241	2.174	1	2	2
Time Spent Browsing	12,241	35.771	6.497	16.789	41.192
Changes Default Map Location?	12,241	0.531	0	1	1
Uses Map Zoom Feature?	12,241	0.383	0	0	1
Uses Maximum Price Filter?	12,241	0.531	0	1	1
Uses Room Type Filter?	12,241	0.636	0	1	1

(b) Searchers Sending a Contact

Statistic	N	Mean	Pctl(25)	Median	Pctl(75)
Number of Searches	4,426	31.174	7	17	36
Unique Listings Seen	4,426	87.812	42	73	114
Share of Listings Seen	4,426	0.119	0.028	0.055	0.117
Num. Nights Requested	4,426	2.880	2	3	4
Num. Guests Requested	4,426	2.226	1	2	2
Time Spent Browsing	4,426	57.874	14.192	32.519	70.097
Changes Default Map Location?	4,426	0.642	0	1	1
Uses Map Zoom Feature?	4,426	0.501	0	1	1
Uses Maximum Price Filter?	4,426	0.654	0	1	1
Uses Room Type Filter?	4,426	0.695	0	1	1

The above table displays summary statistics for searchers who sent a contact in the sample. 'Number of Searches' is the number of distinct searches in the two days leading up to either the first contact or the last search in the city. 'Unique listings' is the number of unique listings loaded on pages browsed by the searcher. 'Share of Listings Seen' is the number of unique listings shown to the searcher divided by the number of listings that fulfilled the parameters of the search and were still marked as available. 'Nights' and 'guests' are the modal trip parameters for the searcher. 'Time Spent Browsing' is calculated, in minutes, as the sum across all search sessions of the length of that session. Sessions are sets of searches without a five minute gap between the search actions.

Table 2: Descriptive Statistics - Contacts

(a) Chicago

Statistic	Mean	Median	Pctl(75)
Number of Contacts	2.367	1	3
Number of Simultaneous Contacts	1.435	1	1
Inquiry First	0.566	1	1
Rejection	0.365	0	1
Stale Vacancy Rejection	0.146	0	0
Congestion Rejection	0.078	0	0
Booked First Contact	0.378	0	1
Booked Any	0.623	1	1

(b) Top 50 US Markets

Statistic	Mean	Median	Pctl(75)
Number of Contacts	2.383	1	2
Number of Simultaneous Contacts	1.440	1	1
Inquiry First	0.601	1	1
Rejection	0.345	0	1
Stale Vacancy Rejection	0.151	0	0
Congestion Rejection	0.078	0	0
Booked First Contact	0.384	0	1
Booked Any	0.596	1	1

The above table displays summary statistics for searchers who sent a contact to listings in the top 50 US Markets. Each observation is a searcher who sent at least one contact in the city. ‘Simultaneous Contacts’ refers to the number of contacts that occur within one hour of the first contact (including the first contact). ‘Inquiry First’ takes the value of one when the first contact was an inquiry rather than a booking request or instant booking. ‘Rejection’ equals one when the first contact was rejection. ‘Stale Vacancy’ refers to rejections that were followed by the host setting the inquired for dates to be unavailable. ‘Congestion’ refers to rejections that occurred because a prior contact to the same host for overlapping dates resulted in a booking. ‘Booked First Contact’ refers to the first contact resulting in a booking. ‘Booked Any’ equals one if the searcher booked at least one listing in the market for the inquired set of dates.

Table 3: The Effect of Rejections on Booking

	Dependent Variable					
	Guest Books Any Listing for the Trip Date / Market					
	(1)	(2)	(3)	(4)	(5)	(6)
Rejection	−0.411	−0.302	−0.229	−0.279	−0.288	−0.451
	(0.005)	(0.004)	(0.006)	(0.006)	(0.006)	(0.016)
Has Prior Booking		0.362	0.408	0.342	0.281	0.313
		(0.005)	(0.006)	(0.007)	(0.006)	(0.007)
Rejection * Has Prior Booking			−0.160			
			(0.009)			
Mean Booked:	0.631	0.631	0.631	0.644	0.804	0.631
Mean Booked (No Prior Booking):	0.407	0.407	0.407	0.439	0.548	0.407
Send Week - Check-in Week FE	No	Yes	Yes	Yes	Yes	Yes
Guest and Trip Characteristics	No	Yes	Yes	Yes	Yes	Yes
Booking Requests Only	No	No	No	No	Yes	No
2SLS	No	No	No	No	No	Yes
Observations	43,724	43,691	43,691	26,061	22,562	43,691
Adjusted R^2	0.160	0.295	0.300	0.259	0.300	0.276

The table displays the results of regressions where the outcome variable is whether a searcher booked a listing in Chicago for a given trip date. 'Rejection' is an indicator variable for whether the first contact of the searcher was rejected and 'Has Prior Booking' is an indicator for whether the guest had previously booked on Airbnb. Only searchers with one contact within the first 2 hours of contacting are included in the sample. Column (4) limits the observations where the first contact was sent more than 2 weeks before the check-in date. Guest and trip characteristics include a cubic lead time polynomial (in days), an indicator for whether the guest was reviewed prior to the contact, number of nights, number of guests, and the guest's country of origin. 'Booking requests' occur when the searcher uses the 'Book It' or 'Instant Book' options and commits a searcher to booking if the host accepts. The two stage least squares regression in column (6) uses an indicator for whether the rejection was due to a stale vacancy as the instrument in the first stage.

Table 4: Choice Model Estimates

Variable	Estimate	Std. Error
Price	-0.005	0.0004
Avg. Rating	0.037	0.034
Total Reviews	-0.014	0.004
Total Five Star Reviews	0.021	0.005
Has Prof. Photo	-0.016	0.036
Weird Property Type	-0.460	0.202
Is Instant Bookable	-0.405	0.050
No Reviews	0.419	0.162
In Neigbhorhood Filter	0.940	0.046
Entire Prop.	0.312	0.054
In Room Type Filter	0.713	0.056
Minimum Search Rank	-0.107	0.004
Outside Option (OO)	5.711	4.581
OO x Guests	-0.187	0.033
OO x Checkin Date	-0.00004	0.0003
OO x Num. Listings Seen	-0.001	0.001
OO x Lead Time	0.007	0.002
OO x Has Max Price	-0.705	0.098
Price x Max Price Filter	0.00001	0.00000
Price x Has Max Price	-0.008	0.001

The above table displays results from a conditional logistic regression of choice probabilities among searchers. “Weird Property Type” refers to properties that are not apartments, houses, or condos. “Is Instant Bookable” refers to whether a listing is open to being instant booked by at least some sample of users. “In Neighborhood Filter” equals one when the listing’s neighborhood is filtered for by the searcher in a least one search. “In Room Type Filter” equals one when the listing’s room type (private or entire home) is filtered for by the searcher in a least one search. “Entire Prop.” equals one when the listing is for the rental of an entire property rather than a room within a property. “OO” equals one when the option is the outside option. “Max Price Filter” is the maximum price filtered for by the searcher across all searches. “Minimum Search Rank” is the minimum rank at which the listing appeared in a searchers results. “Num. Listings Seen” is the number of unique listings loaded by Airbnb for the search pages browsed by the searcher. “Lead Time” is the time (in days) between the search and the check-in date. “Hotel Price” is the average hotel price in Chicago for the days of the stay.

Table 5: Rejection Model Estimates

	Dependent variable:		
	logistic	*generalized linear mixed-effects*	*conditional logistic*
	(1)	(2)	(3)
Inquiry First	0.786	0.833	0.792
	(0.017)	(0.019)	(0.020)
Guest Reviewed	−0.066	−0.090	−0.088
	(0.023)	(0.025)	(0.026)
Guest Has About Description	−0.034	−0.037	−0.039
	(0.018)	(0.020)	(0.021)
Guest Verified	−0.109	−0.111	−0.102
	(0.018)	(0.020)	(0.021)
Guest Has Profile Photo	−0.006	−0.022	−0.020
	(0.018)	(0.020)	(0.021)
New Guest	−0.027	−0.021	−0.020
	(0.019)	(0.021)	(0.022)
Num. Guests	0.043	0.017	0.004
	(0.015)	(0.019)	(0.020)
Num. Nights	−0.086	−0.144	−0.156
	(0.005)	(0.006)	(0.006)
Entire Property	0.432	0.408	
	(0.034)	(0.055)	
Multi-listing Host	−0.005	−0.031	
	(0.021)	(0.045)	
Instant Book - Experienced	−0.942	−0.582	
	(0.030)	(0.049)	
Instant Book - All	−1.496	−0.890	
	(0.043)	(0.064)	
Instant Book - Social	−0.668	0.129	
	(0.192)	(0.359)	
Full Guest Capacity	0.104	0.110	0.086
	(0.016)	(0.022)	(0.026)
Reviewed Guest * Multi-Listing Host	0.057	0.123	0.130
	(0.039)	(0.043)	(0.044)
Num Guests * Entire Property	−0.066	−0.044	−0.015
	(0.016)	(0.020)	(0.020)
Check-in Month and Lead Time Fixed Effects	Yes	Yes	Yes
Demographic Controls	Yes	Yes	Yes
Listing RE	No	Yes	No
Listing FE	No	No	Yes
Observations	93,851	93,851	93,851
Log Likelihood	−54,053.190	−49,167.270	−38,390.240

The above table displays results from three logistic regression models predicting whether a contact was rejected, with column (2) including listing random effects, and column (3) including listing fixed effects. "Inquiry First" is an indicator for whether an inquiry was sent rather than a booking request. "Foreign Guest" refers to guests outside of the United States, "Guest Reviewed" refers to whether the guest had at least one review prior to inquiry, "Guest Has About Description" refers to whether the guest had a profile description, "Guest Verified" refers to whether the guest's identity was verified by Airbnb, "New Guest" refers to guests who signed up within 31 days of the inquiry. "Multi-listing Host" refers to a host who has more than 2 active listings. "Instant Book" refers to hosts who allow guests to book without the possibility of rejection. "Experienced" requires that guests have had a prior stay, "All" is open to all potential guests, and "Social" is open only to guests with a social connection to the host. "Full Guest Capacity" refers to inquiries in which the number of guests equals the capacity of the listing. Demographic controls for age (guest and host), gender (guest and host), and whether the guest is traveling from the US are included in the above models.

Table 6: Counterfactual Outcomes

	Scenario	Accepted	Outside Option	Reject (Contact)	Reject (All)	Rej. Prior (All)	Rej. Other (All)	E(U) - No Reject
(1)	Status Quo	0.243	0.635	0.318	0.122	0	0.122	155.073
(2)	Random Order	0.077	0.651	0.775	0.272	0.218	0.054	139.813
(3)	Random - Available	0.182	0.682	0.426	0.135	0	0.135	118.285
(4)	Average Utility	0.080	0.422	0.857	0.499	0.439	0.059	371.651
(5)	Average Utility - Available	0.278	0.501	0.444	0.221	0	0.221	239.960
(6)	Personalized	0.102	0.379	0.833	0.519	0.444	0.075	466.126
(7)	Personalized - Available	0.306	0.454	0.441	0.240	0	0.240	319.103
(8)	Rejection Weighted	0.336	0.471	0.362	0.193	0	0.193	282.366
(9)	Excluding Reject Probability Over 45%	0.295	0.544	0.360	0.161	0	0.161	206.220

This table displays model predictions from counterfactual scenarios in which the results seen by the searcher are altered. "Status Quo" shows predicted outcomes in which searchers see the listings which they actually saw in the data. "Random Order" shows outcomes where the searcher sees a random sample of listings. "Average Utility" shows predicted outcomes in a scenario where the searcher sees the top listings ranked by the average expected utility conditional on booking across all searchers. "Personalized Order" shows outcomes where the searcher sees the top listings as ranked by expected utility conditional on booking. "Rejection Weighed" displays the results from determined by the expected utility weighed by the probability of acceptance. "Exclude Reject Probability Over 45%" shows the outcomes when ranking is personalized but options with a predicted screening rejection probability of over 45% are excluded search.

Outcomes with "Available" denote that the search results are limited to listings which have not been booked or marked as unavailable by the time of the search. "Accepted" is the probability that the searcher either send a first contact that is accepted, "Outside Option" is the probability that the searcher sends no contacts, "Reject (Contact)" is the probability that a searcher's contact is rejected conditional on a contact being sent. "Reject (All)" is the unconditional probability of a rejection. "Rej. Prior (All)" is the unconditional probability that a contact is rejected because it was sent to a previously unavailable listing. "Rej. Other (All)" is the probability that a searcher sends a contact which is rejected due to stale vacancy or screening. "E(U) — No Reject" is the expected utility gain relative to the outside option from a contact if there were no rejection.

A Assumptions Required to Make the Choice Model 'Structural'

The choice model estimated in section 3 can be interepreted as structural under highly restrictive assumptions. These assumptions include (i) that the price is exogenous conditional on observed characteristics, (ii) that searchers do not factor the idiosyncratic probability of rejection by a given listing in their contact decisions, (iii) that, conditional on the observed characteristics, the set of filtering decisions is independent of the listings shown on any given search, and (iv) that the minimum search rank term is treated as a utility relevant parameter.

With regards to assumption (i), there are clearly omitted listing characteristics which are not captured in my model, such as the whether a listing's photo looks good. With regards to assumption (ii), it is hard for searchers to know rejection probabilities, as this information is not displayed in the search results. Therefore, I view this assumption as reasonable. Assumption (iii) is surely violated because the filter choice of a searcher for the next search action is likely affected by the set of listings shown by the search engine on the current search. However, there is no suitable model of this process in the literature that captures the facts discussed in section 3 and estimating such a model is not the purpose of this paper.[27]

One mitigating feature of the Airbnb setting is that because listings are removed from results when booked or blocked, there is a lot of variation in the consideration sets of searchers which is unrelated to searcher characteristics and endogenous filtering. Lastly, the estimated value of the outside option includes the effective search cost because those who choose the outside option avoid the cost of sending a contact and even those who do send a contact have some probability of leaving without a transaction.

[27] Chen and Yao (Forthcoming) and los Santos and Koulayev (Forthcoming) model the decisions of searchers to use refinements in a rational expectations framework.