Can Lower-quality Images Lead to Greater Demand on AirBnB?

Shunyuan Zhang

David A. Tepper School of Business

Carnegie Mellon University, Pittsburgh, PA

shunyuaz@andrew.cmu.edu

Nitin Mehta

Rotman School of Management

University of Toronto, Toronto, Canada

Nmehta@rotman.utoronto.ca

Param Vir Singh

David A. Tepper School of Business

Carnegie Mellon University, Pittsburgh, PA

psidhu@cmu.edu

Kannan Srinivasan

David A. Tepper School of Business

Carnegie Mellon University, Pittsburgh, PA

kannans@cmu.edu

Can Lower-quality Images Lead to Greater Demand on AirBnB?

Abstract

We investigate how AirBnB hosts make decisions on the quality of property images to post. Prior literature has shown that the images play the role of advertisements and the quality of the images have a strong impact on the present demand of the property – as compared to lower quality amateur images, high quality professional images can increase the present demand by 14.3% on matched samples (Zhang et al. 2018). However, the reality is that there exist a large number (approximately two-thirds) of amateur (low-quality) images on AirBnB. One possible explanation is that these images are costly for the hosts, as most of them are amateur photographers. However, this does not completely explain the result – in 2011, AirBnB started offering highest quality professional images *for free* to all the hosts by sending their professional photographers to the property and shoot, process and post the photos for the hosts. To AirBnB's surprise, only 30% of the hosts used the AirBnB professional photography program. We posit that the host's decision on what quality of images to post depends not only on the advertising impact of images on the present demand and on the cost of images, but also on the impact of images on the future demand. Thus, some hosts would be hesitant to post high-quality images because they can create unrealistically high expectations for the guests, especially if the actual property is not as good as what the images portray and if the hosts are unable to provide a high level of service to match those expectations. This would result in the satisfaction level of guests to decrease, who would then leave a bad review or not write any review at all; and since the number/quality of reviews is one of the key drivers in generating new bookings, this will adversely affect the future demand.

In this paper, we attempt to disentangle the aforementioned factors that influence the host's decision on the type of photographs to post and explore policies that AirBnB can employ to improve the hosts' adoption of professional photos and thereby improve the profitability of both the hosts and AirBnB. To do so, we build a structural model of demand and supply, where the demand side entails modeling of guests' decisions on which property to stay, and the supply side entails modeling of hosts' decisions on what quality of images to post and what level of service to provide in each period. We estimate our model on a unique one-year panel data consisting of a random sample of 958 AirBnB properties in Manhattan (New York City) where we observe hosts' monthly choices of the quality of images posted and their and service that they provided. Our key findings are: *First*, guests who pay more attention to images tend to care more about reviews, revealing an interesting trade-off problem for the hosts. *Second*, hosts incur considerable costs for posting above-average quality of images. *Third*, hosts are heterogenous in their abilities in investing service effort. In counterfactual analyses we simulate AirBnB properties assuming they all start with entry state and low-level images. We then compare the impact of the current policy (offering free high-level images to hosts) and of a proposed policy (offering free medium-level images to hosts) on the average property demand. We show that the proposed policy, though dominated by the current policy in the short-run (for the first four periods), outperformed the currently policy in the long-run (7.6 % vs 12.4%). The interpretation is that, medium-level images, compared to high-level images, despite forming a smaller expected utility for the consumers, has a greater effect on property demand in the long-run as they, with lower risks of creating a dissatisfactory gap, help hosts to obtain new reviews. Moreover, individual hosts who might end up using amateur (low-level) images to avoid the dissatisfactory gap under the current policy, now use free medium-level images to make more revenues under the proposed policy. In the second counterfactual, we explore an alternative policy in which AirBnB were to offer a menu of image quality choices for free. The menu includes both high- and medium- level of property images (images examples are provided) and allow the hosts to self-select which program they want. Comparing with the proposed policy in the first simulation, we find that this policy performs the best in the long-run by improving average property demand by 16.2%.

Keyword: Image Analytics; Sharing Economy; Deep Learning; Dynamic Structural Model;

1. Introduction

AirBnB, a peer-to-peer short-term lodging marketplace provider offering near 5 million listings across 81,000 cities, has hosted over 300 million host arrivals since its start in 2008[1]. Previous studies on AirBnB suggest that the determinates in property demand include property type (apartment, house etc.), property size (number of bedrooms), property price, property location, guest reviews, host service, and property image quality (Zhang et al. 2008, Li and Srinivasan 2018, Li et al. 2016). Only a few factors, i.e., image quality, service, and price—are in immediate control of the hosts. The focus of this paper is on investigating how the AirBnB hosts make decisions on the quality of photographs to post.

Property images, as a way of providing visual information about a listing to the consumers[2], can effectively increase the demand for an AirBnB property. In the context of photos, Zhang et al. (2018) investigated the impact of photo quality on AirBnB demand. Analyzing property images with a deep learning model, they classified the image quality along two levels, high-quality (professional images) and low-quality (amateur images)[3]. They found that the photographs played the role of an advertisement and found the quality of photographs had a significant impact on the property demand, with high-quality images increasing a property's present booking by 14.3%. The strong advertisement impact of images on present demand arises because of the special context of AirBnB. Specifically, AirBnB properties, relative to hotels' rooms that are fairly standardized, have a large variation in terms of styles, characteristics, and hosts. Moreover, most consumers, with a large number of alternatives on the platform, do not repeatedly choose and stay in the same property. In addition, analyzing written details of the property and comparing those across properties can be very onerous and time consuming. Consumers, in order to ease decision-making, hence rely heavily on visual information, which can be easily accessed and processed, to skim through and quickly compare lots of the lodging alternatives. Based on the images, consumers form an expectation on the quality of their stay at the property and accordingly make their decision on which host's property to stay.

Despite the importance of the quality of property images in enhancing the demand, the reality is that there exist a large number (approximately two-thirds) of amateur (low-quality) images on AirBnB. One possible explanation for the low adoption of high quality images is that high quality images are costly for

[1] https://press.atAirBnB.com/fast-facts/.

[2] In this paper we use listing and property interchangeably, guest and consumer interchangeably.

[3] The image quality is classified with a convolution neural network (CNN) that analyzes the basics very high-dimensional pixel information in the images in training set, extracting a hierarchical set of image features that have the most prediction power on the image quality label.

the hosts, as most of them would be amateur photographers. However, this does not completely explain the result—in 2011, AirBnB offered highest quality professional images to all the hosts by sending their professional photographers to the property and shoot, process and post the photos for the hosts. Not only was this program free for all hosts, it also required very little effort from the host's part. Moreover, all hosts were made aware of this program (explain how they were made aware). To AirBnB's surprise, only thirty percepts of the hosts used the AirBnB professional photography program after its launch. This result is intriguing since we would expect the demand to explode when a high quality product/service is offered for free (Shampanier et al. 2007).

In this paper, we provide an explanation for hosts' behavior. We posit that the host's decision on the quality of images to post depends not only on the advertising impact of the photos on present demand and the cost of photos, but also on the impact of the photos on the satisfaction level of the guest post consumption, which would then in turn impact the future demand. The last point follows from the reference dependence literature which suggests that the images create a reference point for the guest in terms of what quality to expect, and their satisfaction level post consumption individual's utility from consuming a product depends not only on the realized outcome, but also on her reference point—the individual's pre-consumption expectation (Kahneman and Tversky 1979, Tversky and Kahneman 1991, Koszegi and Rabin 2006). Particularly, individuals tend to react severely to a 'dissatisfaction or disappointment gap'—i.e., when the actual outcome turn out to be worse than the expectation (Genesove and Mayer 2001). Thus, some hosts would be hesitant to post high quality photographs (even if they were free) because they can create unrealistically high expectations for the consumers, especially if the actual property is not as good as what the photos portray and if the hosts are unable to provide a high level of service to match the high expectations. As a result, the consumers' satisfaction level would decrease, and they would either leave a bad review or would not write any review at all[4]; and this lack of reviews would in turn adversely impact the future demand of future guests of that property.

In summary, professional property images, although more expensive, can help generate bookings for the AirBnB hosts in the current period, since consumers with imperfect information rely heavily on images

[4] Dissatisfied guests instead of writing bad reviews, they tend to not to write a review. As previous studies on online reviews suggest, the 'silence' in online reviews actually reflects customer dissatisfaction (see Dellarocas and Wood (2008)'s work on 'the sound of silence"), because customers with bad experience tend to choose not to leave a review (Masterov et al. 2014, Nosko and Tadelis 2015). This is consistent with the observations that online reviews tend to be positive, possibly because giving a negative rating is costly for the consumers. Particularly, with a field experiment on AirBnB that involves encouraging consumers to leave reviews, Fradkin et al (2018) found that guests get more utility from leaving a positive review, and they also don't like to misrepresent their experiences. As a result, a guest is less likely to leave a review, if she is unsatisfied with her stay experience.

to make their lodging decisions. On the other hand, professional images can lead to a dissatisfaction gap if the actual property is not as good as what the professional images portray or if the hosts are unable to provide a high level of service to match the high expectations. Our goal is to disentangle the aforementioned factors that influence the host's decision on the type of photographs to post, and to explore policies that platforms such as AirBnB can employ to improve the hosts' adoption of professional photos and thereby improve the profitability of both the hosts and AirBnB. To achieve this goal, we have the following objectives, which we explain as follows.

1.1. Research Objectives and Main Findings

The *first* objective of this paper is to model hosts' periodic (monthly) decisions on the quality of property images to post, and the quality of service to provide. The image decision entails choosing between three quality levels of images—low, medium, and high[5]. The service decision entails choosing between two levels of service: high and low. The image decision affects the host's profits in the short run through the costs associated with preparing, shooting and editing the particular quality-level images, and through the impact of images on the present demand. And it affects the host's future profits via the following mechanism: professional images come with a risk of increasing the guests' dissatisfaction gap. This decreases their likelihood of writing reviews, which then negatively impacts the future demand. The service decision impacts the host's profits in the short run through the costs (good services come with a cost), and in the long run by impacting the guests' satisfaction level and their subsequent likelihood of writing reviews.

To achieve the first objective, we need to model both the guests' decisions on which property to choose (the demand side) and the hosts' decision on the quality of images to post and the level of service to provide (the supply side). The property demand is a function of property characteristics including property images, number of reviews and prices. The property supply function models how hosts make images and service decisions, taking into account the impact of their actions on the current and future utility.

Regarding the demand side, we estimate a random-coefficient logit model (Berry et al. 1995) using AirBnB properties' aggregate monthly market-share data. The guests' utilities are modeled as functions of property characteristics which include property images, number of reviews and property prices. Heterogeneous consumers form expectations on the lodging alternatives, based on their preferences of the property attributes, which include property image quality, prices and the number of reviews. Regarding the supply side, we model the hosts' image quality and service decisions as outcomes of their long-term profit maximization.

[5] This is the aggregate photo quality decision across all photos posted.

The *second* objective is to estimate the model using a panel data consisting of 900+ individual AirBnB hosts' choices of images and service over time. The data contains rich information on property characteristics, property reservation days and monthly revenues, and guest's reviews. There are two unique features in this data: 1) we observe the dynamics in property images and whether an option of AirBnB's free professional photos was available to the host, from which we know hosts' periodic image decisions and infer the associated costs[6]; 2) the periodic reviews a property received from its guests, from which we infer the dissatisfaction gap between image-induce expectation and realized property as well as the invested effort in service. Our key empirical findings are: a) guests have heterogeneous and correlated preferences on property attributes. Particularly, guests who value professional images more, also value the number of reviews more. Thus, for a property that faces such a pool of consumers, using professional images may have a high marginal effect on generating booking in current period. Yet, the 'penalty' in the future is also likely to be high, as these consumers value highly the number of reviews. b) hosts have a considerable degree of heterogeneity in their ability (cost) in investing in service and in values of their outside options. Such heterogeneity results in hosts self-selecting to choose different quality levels of images.

Our *third* objective is to explore image-related policies that a peer-to-peer platform such as AirBnB can employ to effectively improve the overall property performance and service quality. To do so, we run two counterfactuals. In the first counterfactual, we examine three image polices. The first policy is the same as AirBnB's professional photography program where it provides the highest-level professional images for free to all hosts (current policy). The second policy is an alternative policy in which AirBnB instead provides medium quality-level images for free to all hosts (proposed policy 1). In both policies A and B, we allow for hosts to self-select on whether they would adopt the free AirBnB program or not. The third policy is the baseline policy in which AirBnB were not to offer any photography service to the hosts (baseline). We find that, both policies A and B, compared to the baseline policy, substantially improve the average property demand across all properties. Interestingly, policy B, though dominated by the policy A in the short-run (for the first four periods), outperformed policy A in the long-run (12.4% vs 7.6%).

Our results indicate that, medium-level images, compared to high-level images, despite forming a smaller expected utility for the consumers, has a greater effect on property booking in the long-run as they, with lower risks of creating a dissatisfactory gap, help hosts to obtain new reviews. Moreover, individual hosts who might end up using amateur (bad-quality) images to avoid the dissatisfactory gap under current

[6] AirBnB offered professional photos during our observation window. However, AirBnB's photography program says it can be offered to the same listing for once. Hence, for roughly one-third of the properties in our sample, which were observed to already have AirBnB-offered photos by the start of our observation, if they were to change their photos with another batch of professional photos, they would incur a cost on their own.

policy, now use free medium-level images to make more revenues under proposed policy 1. In the second counterfactual, we explore an alternative policy in which AirBnB were to offer a menu of image quality choices for free. The menu includes both high- and medium- level of property images (images examples are provided) and allow the hosts to self-select which program they want. Comparing with the proposed policy in the first simulation, we find that this policy performance the best in the long-run by improving average property demand by 16.2%.

1.2. Literature Review

We start with discussing the relatively new stream of literature on the sharing economy platform AirBnB. A few recent studies looked at the reputation system, i.e., consumer reviews, on AirBnB (Zervas et al. 2015, Proserpio et al. 2017, Fradkin et al 2018). These studies mainly focused on documenting a particular aspect existing in the reputation system, e.g., reciprocity, without investigating how the reputation system affects consumers' choices of demand and drives hosts' choices of supply. Another stream of studies investigated the impact of AirBnB's entry on the incumbent lodging industry, namely hotels (Li and Srinivasan 2018, Zervas et al. 2017). There are a couple of distinctions between these papers and our paper. 1) these papers focused on quantifying the impact of AirBnB's supply on hotels, while ours aims to understand how AirBnB's supply choices (including image and service choices) are endogenously determined by the hosts. 2) these papers largely treat AirBnB properties as a homogeneous party (or categorized properties into a few sub-groups based on the property type), without taking into account the heterogeneity across the properties and hosts. An exception is the work of Farronato and Fradkin (2018), which incorporates hosts' heterogeneous variable costs in modeling their supply decisions. However, they did not consider the heterogeneity in terms of a property's quality. We argue that, given different properties at different states, hosts' supply choices could be different even they have the same variable cost. (3) none of these studies investigate the role of property images. Our paper differs from the existing study as we have access to a panel data that consist rich information on hosts' periodic image choices. An exception is Zhang et al. (2018), where property images are analyzed and scored by the image quality. However, they employed a quasi-experimental method (difference-in-difference) to make a causal link between property images and property demand, treating image choices as if they were exogenously given. In contrast, we endogenize the image decisions and explain how host heterogeneity drives the observed choices of images and service. This heterogeneity nature of peers, though unobserved, plays a significant role in driving the market equilibrium outcome. This distinction allows us to estimate a more comprehensive model of hosts' decisions, to resolve the observed puzzle—why many hosts did not use AirBnB's professional photos despite free, and to answer our research question posed in section 1.2—what else can AirBnB do with its image-related policy to improve the market equilibrium outcomes?

As discussed above, guests' post-consumption behavior, influenced by potential dissatisfaction gap, affects hosts' pre-consumption decisions. Here we briefly discuss how behavioral argument of consumer dissatisfaction gap is related to and used in our paper. The idea that utility depends on both the actual outcome and the alternative outcome that could have occurred has been posited in a large body of literature in psychology and behavior economics (Gul 1991; Kahneman and Tversky 1979). Specifically, studies on reference dependence suggest that individual's utility from consuming a product depends not only on the realized experience, but also on the reference point—the individual's pre-consumption expectation (Koszegi and Rabin 2006). Hence, with the same actual outcome, individuals may react differently, if they had different expectations and experienced different 'gaps' (Mas 2006). We conjecture that in the context of AirBnB, guests' post- consumption behavior is affected by both the realized stay experience after their arrivals and the expected experience they had when seeing the property images. Particularly, since people react more severely to a 'loss' than to a 'gain' (Genesove and Mayer 2001), if the reality did not meet the expectation, a dissatisfaction gap significantly reduces the guests' likelihood of writing a review.

2. Research Context and Descriptive Statistics

2.1. Research Context

Our research context is AirBnB—one of the largest sharing-economy platform for peers to list their spare rooms and to find short-term lodgings. AirBnB now offers near 5 million listings in over 81,000 cities. Since its foundation in 2008, AirBnB has hosted more than 300 million guest arrivals. AirBnB makes revenue from charging a service fee of 9~12% proportional to each transaction.

2.2. Data Description and Measures of Key Variables

Our sample consists of 958 randomly selected AirBnB properties in Mahanttan, New York. For each property in the sample, we collected property time-invariant characteristics, including property's location, type, size, and capacity. We constructed a panel data of the 958 properties spanning 12 months (January 2016-December 2017). For each property in each month, we obtained dynamic information about the property's demand (i.e., the number of reservation days) as well as the property's supply (i.e., whether a listing was active in a particular month). Such dynamic information also includes property nightly rate, guests review, and property images. Below we describe the definitions of key variables used in our analyses.

Property Characteristics

We obtained property characteristics, defined by the following variables: 1) *EntireHome*, which equals 1 (0) if the property is listed as an entire (shared) place, 2) *Apartment,* which equals 1 (0) if the property is an apartment (or not, e.g., condo or house), 3) *Bedrooms, Bathrooms*, and *Beds*, which indicates the number of bedrooms, bathrooms and beds, respectively, 4) *MaxGuests*, which indicates the maximum number of accommodated guests), 5) *DriveTime*, the driving commute time (in minutes) from each property's address to the downtown area). The driving time is further scaled by 1/10 in the analyses, and 6) *WalkScore,* a score 0-100 based the evaluation of the available nearby amenities such as restaurants, malls, parkings etc., and 7) the area code associated with each property (Manhattan is categorized into 10 subareas or neighborhoods)[7].

Property Reservation

We purchased listing-level reservation data from a third-party company that specializes in collecting AirBnB property booking data. One unique feature in the reservation data is that, they distinguish real booking (days when the property was booked by a guest) from blocking (days when hosts marked the property as 'unavailable). Since blocking days do not reflect the actual property demand, we used only the number of reservation days to construct our demand measurement. For each property *i* in month *t*, variable $ReserveDays_{it}$ indicates the number of days that *i* was booked in that period.

Market Share and Market Size A property's market share reflects its demand on the lodging market. It's defined as the number of property-nights sold in each period (month), divided by the market size. Market size measures the total number of nights (including AirBnB property, hotel, and other alternatives) that could be possibly sold to the travelers, approximated from combining New York City tourism trend report in 2016 and the distribution of hotels across the five boroughs in NYC[8]. In our study, we used a constant market size of 2,400,000 across the periods. The seasonality trend is captured through a series of period fixed effects we incorporated in the property demand equation (see section 4.1). Variable $MarketShare_{it} = ReserveDays_{it}/MarketSize$ indicates property i's market share in month *t*.

Property Nightly Rate

Property i's nightly rate, $NightlyRate_{it}$ is computed as the average of daily prices over the days in period *t*. We further take the logarithm form of the average nightly rate. As in many other markets, property prices

[7] In Manhattan, the 10 areas are: Central Harlem, Chelsea and Clinton, East Harlem, Gramercy Park and Murray Hill, Greenwich Village and Soho, Lower Manhattan, Lower East Side, Upper East Side, Upper West Side, Inwood and Washington Heights. Refer to the following website for more details: https://www.health.ny.gov/statistics/cancer/registry/appendix/neighborhoods.htm

[8] 60.5 million travelers visited New York in 2016. About 48% of the market are day trips hence do not consume any lodging, and the remaining on average consumed 2.4 nights each trip. https://www.nycgo.com/assets/files/pdf/new_york_city_travel_and_tourism_trend_report_2017.pdf; http://mycrains.crainsnewyork.com/stats-and-the-city/2017/tourism/hotel-occupancy-rate-by-year.

Though we don't have tourism data on the borough level, we can approximate the market size for each borough from the hotel distribution in NYC. This is because hotels choose locations with high demand—how the hotels geographically distribute capture how the demand distribute across the areas.

may be correlated with random shocks in the demand which are unobserved to the researchers. To address the endogeneity issue, as we will describe in section 4, a set of instruments were used. These instruments are correlated with $NightlyRate_{it}$, yet should be uncorrelated with demand shocks in the current period t. We also include the number of reviews in the previous period in the instrument variable set, as the number of reviews may affect how a property's nightly rate is set, yet it is uncorrelated with the aggregate demand shocks in current period. Further, following Li and Srinivasan (2018), we collect local (i.e., in the same zip code) rental information from Zillow to serve as an instrument for the monthly average Airbnb property nightly rate[9]. Finally, we include local utility fee in the set of instruments for $NightlyRate_{it}$, as the utility cost can influence the property price however is unlikely to be correlated with unobservables that are correlated with property demand. Specifically, we obtain average residential electricity rates by zip code[10].

Property Active

A host can choose to make a property temporally 'inactive'. For example, if a host feels constantly managing a listing (e.g., checking property page and updating property availability calendar) is a little demanding while the returns from managing is little, then he may choose to "snooze/un-list" the listing for a while. Such, the property will be temporarily not viewed. If afterwards, he/she decides to activate the listing, all the current records remain the same. Variable $Active_{it}$ equals 1 (0) if property i was active (inactive) in period t.

Property Review Count

We collected data property reviews posted on property page. Specifically, we count the accumulated number of posted reviews for property i till the beginning of period t, $NumReview_{it}$. Since number of reviews may be correlated with unobserved property characteristics (as we will describe in section 3, consumers' likelihood of writing a review is influenced by the property quality), in the demand function, we take the approach of instrument variables to deal with the endogeneity issue. Specifically, two variables serve as instruments for $NumReview_{it}$. The first is BLP instruments—the sum of number of review of other properties and the second is the number of hotel rooms in each zip code. The latter serves as an instrument as the supply from hotels is unlikely to be correlated with a property's unobserved attributes, however can influence the evolution of $NumReview_{it}$ through competition that affects the number of bookings a property can receive in a month.

[9] See Zillow Rental Index for more details: http://www.zillow.com/research/data/#rental-data.

[10] National Renewable Energy Laboratory provides average residential, commercial and industrial electricity rates by zip code for both investor owned utilities (IOU) and non-investor owned utilities by combing data from ABB, the Velocity Suite and the U.S. Energy Information Administration dataset 861. For more details, see https://openei.org/datasets/dataset/u-s-electric-utility-companies-and-rates-look-up-by-zipcode-2016.

Property Ratings

We also obtained monthly data on the lodging experiences, rated along multiple dimensions, from peers who have stayed at the property. Specifically, for each property *i* in month *t*, guests can rate the property, on a 0-10 scale, on its accuracy in description, cleanliness, host communication before/during the stay, check-in smooth/convivences, overall value/experience, and location.

Host Response Time and Service Effort

Airbnb's algorithm automatically records and compute the average time (minutes) that a host responded to the consumers in the past 30 days, denoted by *ResponseTime*. The algorithm tracks the communicate between a host and his/her guest during the process of making a reservation, prior to a stay, and during the stay. Such communication could include requesting information about the property, asking details regarding check-in, or any question that a guest may ask during a stay. From the html source code of each property page, we obtain the response time (as in minutes) that the algorithm rated each host for every month. We further categorize *ResponseTime* and create dummy variable *HighEffort* equal 1 if and only if *ResponseTime* is less than 1 hour[11]. Hence, in our study, we look at two different levels of effort—high- and low- effort. Since effort is unobserved to researchers (also unobserved to a guest ex-ante), we use a metrics that measure host's responsiveness to proxy for the level of effort that she invested in a particularly period., Though responding to a guest within 1 hour does not guarantee that a problem/question is resolved within 1 hour, this at least reflect how serious (the attitude) the host is about the guests' communication.

Property Photos

AirBnB hosts post photos of their place on the property page. To capture the dynamics in the set of photos when hosts updated their property images, we measure the aesthetic quality of images as a time-variant variable. Specifically, leveraging computer vision techniques, we built a scalable deep learning model that automatically classifies any property image into one of three categories–namely, "high-quality", "medium-quality", and 'low-quality'. The set of photos posted for property i in period t was then represented by its average image quality. For example, if property i had 10 images in period t, with 8 images classified as high-quality, 1 image classified as medium-quality, and 1 image classified as low-quality, then the average image quality $ImageQuality_{it} = (8 * 1.0 + 1 * 0.5 + 1 * 0)/10 = 0.85$.

Next, in our structural model, we discretized the image quality and categorize images into low-, med, and high- 3 quality levels. Then dummy variable $HighImage_{jt}$ is 1 if and only if $ImageQuality_{it} > 0.75$. Dummy variable $MedImage_{jt}$ is 1 if and only if the $0.75 \geq ImageQuality_{it} \geq 0.5$ Lastly, dummy

[11] Our results stay qualitatively consistent when we use different criteria (i.e., setting the threshold at 'responding within a few hours').

variable $LowImage_{jt}$ is 1 if and only if $ImageQuality_{it} < 0.5$ [12]. Hence, for any property *i* in any period *t*, its image can be expressed with a tuple of binary variables $(MedImage_{jt}, HighImage_{jt})$. Property *i* had low-level of quality property images in period *t* if and only if $(MedImage_{jt}, HighImage_{jt}) = (0,0)$.

A Deep Learning Classifier to Measure Image Quality

We leverage techniques from computer vision and deep learning to build a classifier that, for any given input property image, predicts its image quality as high- versus low- quality. Specifically, we first construct a training set consisting of 3,000 (stratified) randomly selected AirBnB property images, with each the image quality is manually evaluated and labeled by five Amazon Mechanical Turkers (AMT)[13] on a 1-7 score. The image quality for each image is then computed as the mean score averaged across the scores assigned the five raters who evaluated the images. We discretize the image quality to create a quality label for each image, with label of 'low-quality', 'medium-quality', and 'high-quality' corresponds to an average score that are in the range of 1-3, 3-5, and 5-7, respectively.

Next, we build a Convolutional Neural Networks (CNN), a deep learning framework with a series of breakthroughs in vision tasks such as image classification and object recognition (Krizhevsky et al. 2012, Simonyan and Zisserman 2015). We then apply supervised learning and train the CNN classifier on the collected training set. The classifier is optimized by extracting image features that have predictive power on its label (image quality) and by learning the relationship between extracted features and the label. The classifier achieved a high accuracy of 86.7% on a hold-out set. Lastly, we apply the optimized CNN classifier to all property images in our sample to automatically predict the image quality label for each. In appendix, we provide detailed description on machine learning steps and on the architecture of the CNN classifier as well as technical notes on the training process.

Finally, Table 1 summarizes the statistics for the key variables.

[12] Alternatively, we could use mode (i.e., majority) of the image quality to represent to quality level of images associated with a property in a period. We obtained consistent results. This is because we observe that hosts post images that have quality concentrating in the same level—there is very limited behavior of mixing different image quality levels. In the given example, using mode quality would give us a high-level quality for the specific set of images, where the majority (8 out of 10) of the images are high-quality. Using average quality would also give us a high-level quality, as the average image quality is above the cutoff for high-level (i.e., 0.75). In regard to the thresholds of discretizing the average image quality, we tried other partition points (e.g., 0.7 as cutoff for high-level and 0.4 as cutoff for medium-level), the results remain qualitatively unchanged.

[13] https://www.mturk.com/.

Table 1 Summary Statistics of AirBnB Properties

Variables	Observations	Mean	Std. Dev.	Min	Max
EntireHome	11496	0.654489	0.475555	0	1
Apartment	11496	0.938413	0.240414	0	1
Bedrooms	11496	1.180585	0.733769	0	6
Bathrooms	11496	1.094468	0.365832	0	4
Beds	11496	1.655532	1.051104	0	9
MaxGuests	11496	3.187717	1.901601	1	16
MinimumStay	11496	2.766962	2.865382	1	35
HostExperienceYear	11496	3.144262	1.481851	1	8
WalkScore	11496	98.17537	4.010409	62	100
DriveTime (minutes)	11496	13.57516	11.28168	1	56
LocalUtilityRate	11496	0.2163	0.043404	0.158837	0.249082
LocalRentalIndex (Zillow)	11496	2517.88	433.3505	1395	4525
ReserveDays	8622	16.12097	10.92966	0	31
NightlyRate	8622	228.6123	261.8237	28	5000
Active	11496	0.7621	0.433032	0	1
NumReviews	11496	40.63718	44.58707	0	383
OverallRating	9245	92.05116	5.606863	50	100
CommunicationRating	9245	9.724608	.4868291	8	10
AccuracyRating	9245	9.435803	.630361	7	10
CleanlinessRating	9245	9.116387	.8229696	5	10
CheckinRating	9245	9.637642	.5566269	6	10
LocationRating	9245	9.44186	.6645177	6	10
ValueRating	9245	9.136182	.6105189	6	10
HighEffort	8622	0.4066342	0.491234	0	1
HighImg	11496	0.227561	0.425607	0	1
MedImg	11496	0.16365	0.429303	0	1
LowImg	11496	0.608789	0.499669	0	1

2.3. Reduced-form Evidence

We explore patterns in the data and presents a series of reduced-form analyses that suggest 1) both image quality and number of review have positive impact on AirBnB property's present booking (demand) and 2)

4) high-quality of images come with a risk of creating a ‘negative/dissatisfaction gap’ for consumers and could adversely affect the future demand.

2.3.1. Regressing Demand: The Impact on Images and Reviews on Property Bookings

We run the following regression, as specified in Equation (1), to see how image quality and the number of review affect a property's present demand.

$$ReserveDays_{jt} = Intercept + \beta_1 ImageQuality_{jt} + \beta_2 NumReviews_{jt-1} + \\ + \beta_3 Controls_{jt} + Property_j + Period_t \quad (1)$$

where dependent variable—property demand—$ReserveDays_{it}$ indicates the number of reserved days for property j in period t. $ImageQuality_{jt}$ refers to the property's image quality in period t, level, with low-quality serving as the baseline (i.e., its coefficient is normalized to zero). $NumReviews_{it-1}$ indicates the number of reviews that property i has accumulated till the end of period t-1 (i.e., till the beginning of period t). Lastly, $Controls_{it}$ is a vector of control variables. $Period_t$ are time fixed-effects are incorporated to capture seasonality patterns in property demand.

As can be seen in Table 2, improving the quality of posted property images will, else being equal, lead to a greater property demand for the current period. The positive coefficient of *NumReviews* suggests the number of reviews is a key driver in generating property bookings, with a greater significance than the coefficient of image quality. Hence, in the long-run, to consistently get booking, it is essential for a host to be able to grow the reviews. This is particularly crucial, when the peers of the host are accumulating more reviews.

Table 2 Regress Property Bookings on Image Quality and Number of Reviews

VARIABLES	D.V. # Reservation Days Equation (1) # +
NumReviews	0.105***
	(0.00524)
MedImage	0.942**
	(0.3106)
HighImage	1.584***
	(0.425)
NightlyRate	-0.0095***
	(0.00284)
Observations	8622
R^2	0.6903
Fixed Effect	Property
Seasonality	Monthly

Robust Standard errors in parentheses * p<0.05 ** p<0.01 *** p<0.001
property time-varying characteristics (overall rating, service effort, MaxGuests, and MinStays) are controlled. + model is regressed over samples when a property was 'active'.

2.3.2. Law of Motion: The Impact of Image Quality on Guests' Post-Consumption Review-Writing Behavior

We present evidence that higher-quality of images can reduce the guests' post-consumption satisfaction and hence reduce their probability of writing a review. As explained in section 1, on Airbnb, guests who are unsatisfied tend to walk away without writing a review. Moreover, a departure from the expected outcome influences one's post-consumption satisfaction (Koszegi and Rabin 2006). Hence, to capture how an expectation-realization gap plays a role in affecting Airbnb guests' post-consumption satisfaction, in Equation (2) we implement a logistic regression on the probability that a consumer, upon his/her stay, will write a review, as a function of the 'gap' between expectation and realization and other relevant factors.

$$WriteProb_{jt} = \alpha_0 + \alpha_1\left(PropertyQuality_j - ImageQuality_{jt}\right) + \alpha_2 Effort_{jt} + \alpha_3 NightlyRate_{jt} + \alpha_4 Control_{jt} + PropertyQuality_j + Period_t \quad (2)$$

where the dependent variable $WriteProb_{jt}$ is measured the proportion of bookings occurred in period *t* for property *j* that led to a review. Key coefficient αs captures the impact of the realization-expectation gap, the service effort, the property's nightly rate on guests' likelihood of writing reviews after their stays. $Period_t$ are time fixed-effects are incorporated to adjust the property price based on seasonality and to capture patterns that possibly correlate with the overall guests' writing review behavior. Dummy variable $Effort_{jt}$ is 1 if and only if the guests stayed in property *j* during time *t* were provided with a high-level effort of service (i.e., we normalize the coefficient for low-level service effort to 0). The control variables include the number of reviews $NumReview_{jt}$, overall review rating $OverallRating_{\mathrm{jt}}$, and time-variant property characteristics such as maximum number of accommodated guests $MaxGuests_{jt}$, and minimum number of stay nights $MinStays_{jt}$.

Two points are worth noting in the realization-expectation gap, $PropertyQuality_j - ImageQuality_{jt}$. First, it allows bi-directional departures (i.e., the realized outcome—property quality could exceed or not meet the expected outcome—image quality). This bi-directional gap is consistent with hosts and guests' beliefs or experiences that a positive gap ($PropertyQuality_j > ImageQuality_{jt}$) increases a guest's post-consumption satisfaction while a negative gap ($PropertyQuality_j < ImageQuality_{jt}$) decreases the satisfaction[14]. Second, for the same property, we assume its quality (realization) is time-invariant across the one-year panel in our study.

Re-writing Equation (2), we obtain the following property fixed-effect logistic regression:

$$\begin{aligned} WriteProb_{jt} = {} & \alpha_0 + a_1 PropertyQuality_j + PropertyQuality_j - \alpha_1 ImageQuality_{jt} \\ & + \alpha_2 Effort_{jt} + \alpha_3 NightlyRate_{jt} + \alpha_4 Control_{jt} + Period_t \\ = {} & \alpha_j + \alpha_{high} HighImage_{jt} + \alpha_{med} MedImage_{jt} + \alpha_{effort} Effort_{jt} \\ & + \alpha_{price} NightlyRate_{jt} + \alpha_4 Control_{jt} + Period_t \end{aligned} \tag{3}$$

where $\alpha_j = \alpha_0 + a_1 PropertyQuality_j + PropertyQuality_j$ captures property fixed effect. Dummy variables $HighImage_{jt}$ and $MedImage_{jt}$ equals 1 if and only if the aggregated image quality level for property *j* in period *t* is high-quality, and medium-quality, respectively. That is, the baseline image quality here is low-level, and α_{high} (α_{med}) captures the impact, for the same property, of updating images from low-level to high-level (medium-level) on guests' likelihood of writing a review.

As can be seen in Table 3, key coefficients α_{high} and α_{med} are negative, suggesting that higher-quality property images reduce the likelihood that a guest, upon his/her stay, i.e., having observed the realized quality of property, will write a review for that property. Particularly, using high-quality of images has a greater negative impact than using medium-quality images (-0.682 < -0.441) on generating new reviews from the guests. In addition, as we anticipated, the positive estimated coefficient for service effort—α_{effort}—suggests that investing a high-level effort in providing good service to the guests can effectively increase their likelihood of writing a review. Interestingly, the coefficient for property's average nightly rate—α_{price}—is insignificant at the 0.05 significance level, suggesting that once controlling for image quality and service effort, the seasonality property price does not play a role in affecting guests' post-consumption satisfaction and their likelihood of writing reviews. The explanation is that, though property price is correlated with property quality, the factors that seasonality-adjusted price captures, such as the property's size, type, amenities, and location, are listed on property page and known to the guests beforehand. Consumers' perceptions about these factors do not change before and after they have arrived in the property. As a result, they do not have impact on the guests' post-consumption satisfaction.

[14] See hosts' discussions on their strategic thinking on how a positive or negative gap would influence the satisfaction: https://airhostsforum.com/t/worth-paying-a-photographer/12724/15. In additional, we observed, from the guests' textual comments, that exceeding the guests' expectation with a property quality better than the image quality, can improve guests' satisfaction.

The results in Table 2 combined with Table 3 suggest two trade-off problems for an Airbnb host. First, if a host uses higher-quality images to make the property look nice, he/she will attract more property bookings. However, this may adversely impact the property demand in the future if he/she is unable to get new reviews from the guests as he/she cannot deliver the stay experience (in terms of property quality and/or service quality) that meets the guests' higher-expectation. Second, though a host can purposefully decrease the expectation for the guests by using low-quality images to improve the guests' post-consumption satisfaction, he/she may unable to effectively generate new reviews from guests as a result of few number of transactions (property bookings) occurred.

Table 3 Law-of-Motion: Regressing Review-Writing Probability

	Equation (3)
VARIABLES	D.V.: WriteProb (#)(+)
HighEffort	0.1794**
	(0.0618)
MedImage	-0.421**
	(0.1493)
HighImage	-0.682***
	(0.1838)
NightlyRate	-0.2201
	(0.2407)
MaxGuests	0.0378
	(0.0792)
MinStays	-0.0110
	(0.0134)
NumReviews	0.0273***
	(0.0042)
OverallRating	0.01052
	(0.0201)
Observations	7546
Log pseudolikelihood	-2910.09
Fixed Effect	Property
Seasonality	Monthly

Robust standard errors are presented in parentheses
+ regressed over samples with property received positive number of bookings
p<0.05 ** p<0.01 *** p<0.001

3. Model

The model-free evidences presented in section 2.3 indicate that image quality, affects both the present property demand (prior-consumption), and that the guests' (post-consumption) likelihood of writing reviews. Particularly, if the higher-quality images do not match the realized lower-quality properties and/or the host does not provide a good service that meet guests' expectation, unsatisfied guests are likely to choose not a write a review. In the long-run, this host may end up losing property demand and revenue as he/she is unable to grow the number of reviews, which is a key driver in generating bookings as consumers rely on the number of reviews to make decisions (particularly when the review ratings are seriously inflated)[15]. Knowing this, using high-quality image may not be the best interest of an Airbnb host, even when the images are available for free, if the host's property has low quality and/or the host is unable to deliver good service.

3.1. Model Overview and Timing of Events

One of our main objectives is to predict the dynamic choice of property images and investment in service for each AirBnB host, given his/her own 'type'. To do so, we build a structural model that incorporates hosts' ability in investing in service and the true quality of their properties. This is a dynamic game model, in which individual hosts, other lodging alternative providers including their peers, make monthly decisions about the quality of posted property images and the amount of effort for providing guest service. We assume that hosts are rational and forward-looking, with their objectives to maximize the total discounted utility flow summed over all periods forward.

Figure 1 illustrates the timing/sequence of the events for our model. At the beginning of each period (month) t, an individual host j observes her current state of 1) the current aggregated images quality of her property, categorized into low-, med-, and high- quality states, 3) the total number of guests reviews she accumulated till now, and 4) the states of her peers $-j$. Next, she makes a decision on whether or not to keep the listing active for this period. If she decides to 'snooze/de-active' the listing, then the model for the current period ends and she receives a realized value from choosing an outside option. The model will begin again for her at the beginning of period $t+1$, with her individual state remain the same and her peers' states updated according to peers' actions in period t. If she chooses to operate the listing in period t, then she pays a cost for the operation regardless of received property bookings in this period. Then she makes a decision on: 1) choosing the aggregate image quality from the three quality levels of low-, medium-, and high- for period t, and 2) choosing an effort between low- versus high- levels to invest in providing service in period t. If she updates the images to a different quality level, she pays a cost of preparing and posting images associated with that quality level. She incurs a cost associated with the level of invested service effort. After every individual host has made decisions, the consumers/guests will 1) observe the properties and their characteristics, including the image quality and number of reviews, 2) form an expectation on

[15] Airbnb reported that United States homes were rated 4.8 out of 5 stars with 26,000,000+ reviews.

each property and choose one lodging alternative (guests are allowed to choose an outside option such as a hotel room), 3) receive, after having arrived in the property, a realized loading experience by observing the property quality and being hosted with a certain quality of service, and 4) decide whether or not to leave a review, based on the post-consumption satisfaction. Next, each individual host receives a total amount of the revenue generated from renting out their property. Furthermore, for each reserved day, she pays a cost for hosting the guests associated to the amount of invested service effort. At the beginning of each period, a host chooses the action that maximizes her summed discounted profit (*V*). Lastly, each individual property's number of reviews is updated, corresponding to the number of bookings received and their guests' post-consumption (review-writing) behavior. The model then moves to period $t+1$ and the sequence of events is repeated monthly.

A host's per-period profit can be simply decomposing into revenue she makes from renting out the listing and costs of her actions. Hence, to construct individual host's objective function, we need to first estimate a property demand model from which a host can compute her property's market share and the corresponding revenue.

Figure 1 Timing of Events for Each Month

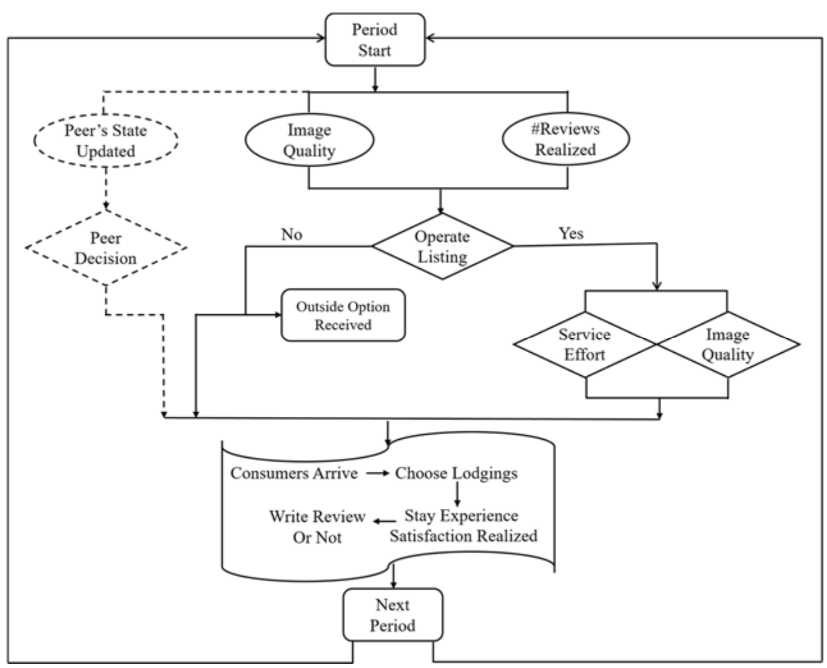

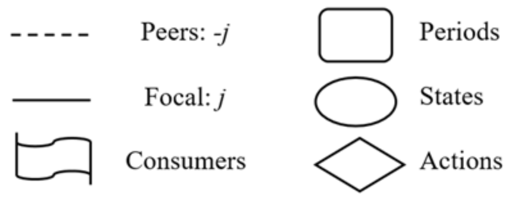

3.2. Property Market-Share Model and Hosts' Revenues

As a main component of hots' objectives, the revenues of renting out their places comes through the population of consumers/travelers choosing lodging alternatives on the lodging supply market. Notably, there are three challenges arise for our study. First, the lodging market consists of a large number of differentiated products. Besides hotels and other lodging options, AirBnB properties themselves may be quite distinct from each other—in terms of property's type, size, location etc... Second, consumers are heterogeneous in their preferences on the lodging features. Third, our data on demand (i.e., property bookings) is at an aggregated market-level. That is, we only observe property bookings as aggregated responses from individual consumers' choices of lodgings, without knowing the trajectories of who booked particular properties.

To resolve above challenges, we an aggregate-demand model introduced in the seminal work of Berry, Levinsohn, and Pakes (Berry et al. 1995, hereafter BLP). The appealing BLP framework has been widely applied in economics and marketing, as it uses readily-available aggregate level sales data, allowing for unobserved individual consumer heterogeneity and producing a more realistic product substitution pattern (Davis 2006, Houde 2006, Nevo 2001, and Sudhir 2001).

Suppose there are *J* AirBnB properties (i.e., products) on Manhattan (New York City) lodging market. In each period, in total *I* consumers choose at most 1 property from the *J* alternatives. Consumers are also allowed choose an outside option (denote as *j*=0, e.g., choosing a hotel or staying at friend's home). Each property is viewed as a set of property attributes *X*, on which consumers evaluate to make decisions on which property to choose. That is, the utility that consumer *i* choose alternative *j* in period *t* can be written as

$$u_{ijt} = \begin{cases} X_{jt}\theta^{i} + \eta_{jt} + \epsilon_{ijt} & if\ i\ chooses\ from\ the\ J\ products \\ \epsilon_{ijt} & if\ an\ outside\ option\ is\ chosen \end{cases}$$

where X_{jt} is a 1 by *K* product-attribute vector (as described in section 3.2) and θ^{i} captures consumer *i*'s preferences over the *K* attributes. Specifically, X_{jt} = {$Period_t$, $Area_j$, $EntireHome_j$,$Apartment_j$, $Bedrooms_j$, $Bthrooms_j$, $Beds_j$, $MaxGuests_{jt}$, $MinimumStay_{jt}$, $NumReviews_{jt}$, $MedImage_{jt}$, $HighImage_{jt}$, $DriveTime_j$, *WalkScore*$_j$, *NightlyRate*$_{jt}$}, where *Period*$_t$ is period (month) fixed effects included to capture seasonality and *Area*$_j$ a area (neighborhood) fixed effects to capture the geographic-related popularity. η_{jt} is a common aggregate demand shock across consumers. The idiosyncratic shock ϵ_{ijt} follows an *i.i.d.* distribution $F_\epsilon(\epsilon)$.

To incorporate possible heterogeneity in the consumers' preferences, we further model individual-specific preference θ^i as an independent draw from the preference distribution $F_\theta(\theta,\omega)$ characterized by parameter ω. Each consumer chooses the alternative that gives him/her the highest utility, and the aggregated (i.e., integration over the population) response of *I* consumers' choices gives us the market share for alternative *j* in period

$$ms_{jt} = \int_{\theta^i,\epsilon_{it}|u_{ijt}\geq u_{ilt},\forall l\neq j\}} dF_\theta(\theta,\omega)\, dF_\epsilon(\epsilon) \tag{4}$$

Following the convention in related literature (Berry et al. 1995, Dube et al. 2012), we model consumer preferences $F_\theta(\theta,\omega)$ follow a normal distribution with $\omega = (\bar{\theta},\sum)$ that characterizes the means and the covariance matrix of the *K*-dimension parameter vector. Further, idiosyncratic shock ϵ_{ijt} is assumed be drawn from type-I Extreme Value distribution. Then Equation (4) can be written as

$$ms_{jt} = \int ms_{ijt}\phi(\theta^i|\bar{\theta},\Sigma)d\theta^i = \int \frac{exp(X_{jt}\theta^i+\eta_{jt})}{1+\sum_{l=1}^{J} exp(X_{lt}\theta^i+\eta_{lt})}\phi(\theta^i|\bar{\theta},\Sigma)d\theta^i \tag{5}$$

where $ms_{ijt}\frac{exp(X_{jt}\theta^i+\eta_{jt})}{1+\sum_{l=1}^{J} exp(X_{lt}\theta^i+\eta_{lt})}$ indicates the probability that consumer *i* chooses property *j* from the choice set *j*=0, 1, …, *J* (*j*=0 indicates the outside option, with its coefficient normalized to one for identification's purpose) in period *t*. Here individual-specific preference $\theta^i = \bar{\theta} + v^i$, with v^i an independent draw from $N(\mathbf{0},\sum)$. Hence $\bar{\theta}$ reflects the average preference in the population on the *K* property attributes and v^i quantifies individual's deviation from mean preference $\bar{\theta}$. The covariance matrix of the normal distribution, $\sum$, thus captures the extent of consumer heterogeneity and correlations, if any, between the preferences.

Finally, with our market share data on *J+1* alternatives spanning *T* periods, we obtain *J*T* demand equations. For each *j=1,2,…,J* and *t=1,2,..T*, we write a market share as specified in Equation (6):

$$\begin{aligned} ms_{jt} &= \int \frac{exp(X_{jt}\bar{\theta}+\eta_{jt}+X_{jt}v^i)}{1+\sum_{l=1}^{J} exp(X_{lt}\bar{\theta}+\eta_{lt}+X_{lt}v^i)}\phi(v^i|\mathbf{0},\Sigma)dv \\ &= \int \frac{exp(\mu_{jt}+X_{jt}v^i)}{1+\sum_{l=1}^{J} exp(\mu_{lt}+X_{lt}v^i)}\phi(v^i|\mathbf{0},\Sigma)dv \end{aligned} \tag{6}$$

where $\mu_{jt} = X_{jt}\bar{\theta} + \eta_{jt}$ is the "mean utility" for alternative j common across consumers in period t.

As can be seen from Equation (6), the set of market shares, $ms_t = (ms_{1t}, ms_{2t}, \dots, ms_{Jt})'$, given the observed covariates $X_t = (X_{1t}, X_{2t}, \dots, X_{Jt})'$ and the preference distribution $N(\bar{\theta}, \Sigma)$, is a function of the demand shocks $\eta_t = (\eta_{1t}, \eta_{2t}, \dots, \eta_{Jt})'$. As we will discuss in section 4, following Jiang et al. (2009), we specify distribution for η_{jt} then evaluate the likelihood function for estimating the demand equations. Though the Bayes estimators, compared to GMM estimators, require an additional distributional assumption (i.e., the assumption on η_t) to derive the likelihood, they have a couple of prominent advantages: 1) The MCMC (Monte Carlo Markov Chain) method implemented in deriving Bayes estimators provide a natural and unified framework for conducting inference (from the stationary posterior distribution) on the functions of model parameters such as price elasticity and markups. However, in the GMM framework, one would have to implement extra computations outside the model parameter estimation procedures, e.g., through bootstrap methods (Nevo 2001) to obtain asymptotic standard errors of theses nonlinear functions of the model parameters. 2) Jiang et al. (2009) conduct simulation experiments and show that GMM estimators' asymptotic standard errors understate the true variance in the (simulated) samples and that Bayes estimators have lower MSE (Mean Squared Error) than GMM estimators. 3) The distributional assumption on the aggregate demand shocks give flexibility when conducting policy simulations or computing price elasticities. However, in a GMM framework, one would either impose a value of zero for demand shocks or use the realized demand shocks when computing price elasticities. 4) Jiang et al. (2009) show that Bayes estimators are quite robust to possible misspecifications on the i.i.d. normal distribution assumption of η_t (a departure from normality, independence, homoscedasticity etc.).

3.2.1 Addressing Endogeneity

A couple of variables in the demand equation suffer from potential endogeneity issue. Below we first describe how the endogeneity concern is addressed in a Bayesian framework and then introduce the set of Instrument Variables (IV) that we use for each of endogenous variables.

In the presence of endogenous variables in the BLP model, a conventional approach is GMM method that exploit the orthogonality conditions (Berry et al. 1995, Nevo 2000). Specifically, suppose the observed covariates X_{jt} can be decomposed as $X_{jt} = \{W_{jt}, P_{jt}\}$, where P_{jt} indicates the endogenous variable that may be correlated with demand shocks (η_{jt}) that are unobserved to the researchers and W_{jt} are all other exogenous variables. GMM method requires to find a set of instrumental variables for P_{jt}—Z_{jt} that is orthogonal to η_{jt} to construct moment conditions.

In our Bayesian-BLP framework, following Jiang et al. (2009), we use a Bayesian approach to instrumental variable (see Rossi et al. (2005)) to address the endogeneity issue. Similar to a 'first-stage' in a classic 2SLS estimation, the following linear regression relates P to Z:

$$P_{jt} = Z_{jt}\delta + \xi_{jt}$$

where ξ_{jt} is a stochastics shock. The endogeneity of variable P_{jt} arises as ξ_{jt} is correlated with the common demand shock η_{jt}. We specify the distributional assumption on ξ and η:

$$\begin{pmatrix} \xi_{jt} \\ \eta_{jt} \end{pmatrix} \sim N\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \Omega = \begin{pmatrix} \Omega_{11} & \Omega_{12} \\ \Omega_{21} & \Omega_{22} \end{pmatrix}\right)$$

Recall that $\mu_{jt} = X_{jt}\bar{\theta} + \eta_{jt} = [W_{jt}, P_{jt}]\bar{\theta} + \eta_{jt}$ and that market shares ms_t are functions of μ_{jt} and η_{jt} (see Equation (6)), the endogeneity properly captured through writing the joint distribution of (P_t, ms_t) as a function of $\begin{pmatrix} \xi_{jt} \\ \eta_{jt} \end{pmatrix}$.

As can be seen, the key step here is to find a set of instrumental variables (IV) for each potential endogenous variable. In our demand model, there are three variables that may be endogenous—property price, image quality, and number of reviews. Below we discuss the endogeneity concern and describe the set of IV used for each variable.

a) Addressing Endogeneity in Property Price

As in the previous literature, price in our study is endogenous, i.e., they may be correlated with demand shocks that are unobserved to the researchers. Following the extent literature using aggregated market share model, we first include so called "BLP instruments' in the set of IV for property price (Berry et al. 1995, Nevo 2001). These instruments use own product's characteristics and the sum of competitor products' characteristics. The logic for the former is that, the own characteristics are determined before the prices are set, hence they are unlikely to be correlated with time-variant unobserved demand shocks. The logic for the latter is that competitors' characteristics are unlikely to be correlated with the unobserved shocks in a product's demand. However, the proximity in product characteristics space between a product and its competitors influence the competition, and as a result, influence the product markup and the price. In additional, we use cost-based instruments—the factors that enter a product's cost side but not demand side, i.e., product-specific cost shifters (BLP 1999, Dube 2012). For this study, we use local (zip code level) residential utility fee obtained and rental information (collected from Zillow). The logic is that these factors serve as an indirect measure of cost and enter price through affecting the cost on the supply side. However, they're unlikely to be correlated with the unobserved factors on the demand side.

b) Addressing Endogeneity in Choice of Image Quality

Hosts' choices of image quality are endogenous in the sense that hosts' incentive to use high-quality images to attract more bookings when the overall demand level varies. Moreover, a host's decision on image quality is affected by the quality of her property. As introduced above, we use "BLP instruments'—the sum of quality of competitors' property images. In addition, we include the following variables as an instrument for the choice of image quality. Property location—some locations or neighborhoods may have higher supply of local photographers, which make it easier for the host to hire a professional photographers on her own; Listing type—it may be easier to manage a professional photography shooting for a property that is listed as entire home, compared to a shared place (especially if the host has a roommate or sublease living in the same property); The number of years of experience as an Airbnb host—a host with more years of hosting experiences may be more likely to hire a professional photographers on her own, because she has more experience or knowledge, compared to a host with less experience, in knowing how to manage the whole process.

c) Addressing Endogeneity in Number of Reviews

The number of reviews on the demand side is endogenous because they may be correlated with unobserved property characteristics (as shown in section 2.3, the guests' likelihood of writing a review is dependent on the property's quality). We use two sets of instruments to address the endogeneity concern. The instruments are uncorrelated with the focal property's unobserved quality, however can influence, through competition, the evolution of number of reviews by affecting the number of bookings a property can receive. The first set of instruments are the 'BLP instruments'—the sum of number of reviews of other Airbnb properties in the same neighborhood, as the number of reviews of local properties are unlikely to be correlated with one's unobserved property quality but is correlate with one's property quality via competition. The second set of captures coemption through the suppled lodging alternatives—the number hotel rooms in the same neighborhood. Similarly, the number of supplied hotel rooms influences a property's number of review via affecting one's received property bookings, however is uncorrelated with one's unobserved quality.

3.2.2. Hosts' Revenues

With the market share model specified in section 3.2.1., a host *j* can approximate the expected market share and hence the expected revenue in any given period *t* assuming he knows $\{X_{jt}\}_{j=1}^{J}$ or expected $\{X_{jt}\}_{j=1}^{J}$. Being able to make money is a major motive that people choose to host on AirBnB. In Equation (7), we specify the indirect- payoff from renting the property *j* in period *t* to capture the monetary motive for a host:

$$
\begin{aligned}
Revenue_{jt} &= ReservationDays_{jt} \cdot NightlyRate_{jt} \\
&= (ms_{jt} \cdot MarketSize) \cdot NightlyRate_{jt} \\
&= (\int \frac{\exp(\mu_{jt} + X_{jt} v^i)}{1 + \sum_{l=1}^{J} \exp(\mu_{lt} + X_{lt} v^i)} \phi(v^i | \mathbf{0}, \Sigma) dv) \cdot MarketSize \\
&\cdot NightlyRate_{jt}
\end{aligned} \tag{7}
$$

where ms_{jt} indicates the property j's market share in period t, as defined in Equation (6), given the property characteristics for the set of alternatives, $\{X_{lt}\}_{l=1}^{J}$, the mean utility for the set of alternatives, $\{\mu_{lt}\}_{l=1}^{J}$, and the heterogeneity in the guests' (consumers') preferences over these characteristics, $v^i \sim N(\mathbf{0}, \Sigma)$.

3.3. Individual Host's Cost

Individual-Specific Cost: Investing Effort in Guest Service

Before a guest booked a property, she does not have much interaction with the host. After she has arrived, for each day of her stay, she receives guest service provided by the host. Such service may include communicating with the guest when then check in/leave to guarantee a smooth transition, keeping the place clean and air fresh, leaving message/travel guide in the room, and answering guests' questions (e.g., regarding dining options in that city or how to use an appliance in the room). A survey on AirBnB suggests that guests care about the quality of received service[16]. In fact, unfriendly/irresponsive hosts and unpleasant conditions are major sources that could lead to serious unsatisfactory in the lodging experiences[17].

However, providing high-quality service is costly to the hosts. Particularly, some hosts, because of their occupations and/or the location of their property, may find it difficult to frequently check their messages and promptly respond to the guests. Thus, we assume that investing more effort for hosting guests cost more to the same host and that investing the same amount of effort may have different costs for different hosts (i.e., hosts have different ability in investing service effort).For property j in period t, the cost of providing service, $ServiceCost_{jt}$ is specified as below

$$
ServiceCost_{jt} = \begin{cases} \lambda_j^{low} & if\ investe\ low\ effort \\ \lambda_j^{high} & if\ investe\ high\ effort \end{cases}
$$

where λ_j^{low} and λ_j^{high} indicates the cost for j to invest low effort and high effort in service, respectively. For identification's purpose, we further normalize λ_j^{low} to 0 and instead estimate the relative marginal effort

[16] https://www.asherfergusson.com/airbnb/.

[17] https://www.asherfergusson.com/AirBnB/.

cost $\lambda_j^{effort} = \lambda_j^{high} - \lambda_j^{low}$ for j. We allow for heterogeneity in individual's marginal effort cost assume λ_j to be i.i.d. draws from a normal distribution with mean $\bar{\lambda}$ and variance σ_λ^2—$\lambda_j^{effort} \sim N(\bar{\lambda}, \sigma_\lambda^2)$ $\bar{\lambda}$ then captures the cost of investing high effort in service to an average host and σ_λ^2 reflects the variation in the host population.

Finally, Equation (8) summarize the service cost that host j incurs in period t:

$$ServiceCost_{jt} = \begin{cases} 0 & if\ HighEffort_{jt} = 0 \\ \lambda_j^{effort} & if\ HighEffort_{jt} = 1 \end{cases} \tag{8}$$

Individual-Specific Cost: Operating Listing (i.e., Opportunity Cost)

At the beginning of each period, a host can choose to make her listing 'active', i.e., to operate the listing in current period, or to 'snooze' her listing, i.e., to temporally exit from Airbnb (no operation on Airbnb) for that period. A host incurs a listing-operation cost for managing an 'active' listing on Airbnb. Listing-operation cost include the activities that a host take to remain the listing 'active', such as managing the property page, keeping the property availability calendar updated, and possible social cost as the neighbors may be unhappy with an AirBnB listing in the neighborhood/building[18], and potential opportunity cost due to listing their property on AirBnB.

We assume an individual-specific operation cost[19] and denote property j's operation cost with $\lambda_{operate}$. Further, for identification's purpose we normalize the cost of keeping an inactive listing in a month to zero (i.e., normalizing the operating cost and outside option value to zero). A host hence incurs zero cost and receives zero revenue from Airbnb in that period if she snoozes the listing.

Lastly, Equation (9) specifies the structure of the operation costs. Note that, unlike individual-specific effort cost, which the host incurs for every booked day, the operation cost is incurred as a fixed cost at the beginning of each period, regardless of the realized booking in that period.

$$OperationCost_{jt} = \begin{cases} 0 & if\ Active_{jt} = 0 \\ \lambda_j^{operate} & if\ \ Active_{jt} = 1 \end{cases} \tag{9}$$

[18] https://www.nbclosangeles.com/news/local/I-Team-Investigation-Short-Term-Rentals-Property-AirBnB-415128373.html.

[19] In reality, these costs are likely to heterogeneous across the hosts/properties. If a neighborhood has very strict policy on home-sharing platform or the neighbors are more again home-sharing, then hosting a property in this neighborhood is likely to be more costly than other hosts. Additionally, the locations of properties introduce variation in the opportunity cost, as local rental (lodging) popularity leads to different outside option values for the properties across geographic areas.

Common Cost: Photography Cost

In addition to the heterogeneous costs of investing service effort and operating listings, there is a key component in the cost that a host may incur—cost of posting property images. Image-posting cost include things such as organize/clean the place, taking photos, do post-processing, and then upload photos. We assume that the photography costs are common across the hosts for two reasons—1) The cost of hiring a professional photographer in a specific market (e.g., Manhattan) are likely to be relatively the same across subareas in that market, and 2) our sample does not observe sufficient variation, at the individual property level, in updating their aggregate-quality level of the property images[20].

Since we categorize property images' aggregate quality into a low-, med-, and high- 3 level, the cost of posting images is likely to differ across the levels. For example, it is easy for a host to take amateur images (with their smartphone phone camera). But to take a med-level image, someone may need to organize the place, ask her friend who can help or spend a whole day of taking lots of photos (trying different scene organization, camera angle, illumination etc.) to pick some good from, and then edit/post process the photos. Taking high-level photos is likely to be the costliest, as one may have to clean and prepare the place, then pay a professional photographer to take photos for her. Thus, we assume that the cost of posting property images for a host is:

$$ImageCost_{jt} = \tilde{\lambda}_{high} I(HighImage) + \tilde{\lambda}_{med} I(MedImage) + \tilde{\lambda}_{low} I(LowImage)$$

where *I(.)* is an indicator function and $\tilde{\lambda}_{high}, \tilde{\lambda}_{med}$, and $\tilde{\lambda}_{low}$ refers to the cost of posting high-, med-, and low- quality level images, respectively. We further normalize the $\tilde{\lambda}_{low}$ to 0 and instead identify $\lambda_{high} = \tilde{\lambda}_{high} - \tilde{\lambda}_{low}$ and $\lambda_{\text{med}} = \tilde{\lambda}_{med} - \tilde{\lambda}_{low}$. We do so for identification's purpose, as we do not observe that hosts post no image. Moreover, we assume that a cost is incurred only when the quality level of images is updated. For example, if a host had high-level image in $t\text{-}1$ period and decides to remain those images for period t, then there is no cost of 'posting' high-quality images.

Lastly, Equation (10) specifies the structure of the photography costs. Note that the photography cost is incurred as a fixed cost at the beginning of each period, depending on a property's current image quality and the image quality decision, before the property bookings are realized in that period. That is, a host

[20] In reality, these costs may vary across hosts. For example, if a host herself is a professional photographer, then we expect that the cost of taking a high-quality photo for her to be lower than for hosts who are amateurs. However, we think the variance in hosts' ability of shooting professional photos are likely to be small as most of them are amateurs. By 2011, very few of the photos on Airbnb were professional, which motivated the company to launch its professional photography program was launched in 2011 (see the report in 2012 from Joe Zadeh, Product Lead at Airbnb).

incurs a photography cost if an only if he updates the image quality to a different quality level that is not low-level.

$$ImageCost_{jt} = \begin{cases} \lambda^{MedImg} & if\ MedImage_{\mathrm{jt}} = 1\ and\ MedImage_{\mathrm{jt-1}} = 0 \\ \lambda^{HighImg} & if\ HighImage_{\mathrm{jt}} = 1\ and\ HighImage_{\mathrm{jt-1}} = 0 \end{cases} \quad (10)$$

3.4. Individual Host's Per-period Payoff

As discussed above, an individual host's per-period payoff can be decomposed into 1) revenue making from renting out the property, 2) effort cost of investing on service, 3) costs of updating property images, and 4) cost of operating the listing. An individual host *j* can choose from 7 possible combinations of actions (3 levels of images*2 levels of effort + operating/snoozing action), denoted by a finite set $A_j = \{1,2, \dots 7\}$. In every period *t*, every host *j* makes a choice $a_{jt} \in A_j$. We further let $a_t = (a_{1t}, a_{2t}, \dots, a_{Jt})$ denotes the set of actions of all individuals in period *t*.

The payoff of taking action *k* for host *j* in period *t* is specified in Equation (11):

$$\Pi_{jat}(a_{jt} = k) \quad (11)$$
$$= \begin{cases} Revenue_{jkt} + ServiceCost_{jkt} + ImageCost_{jkt} + OperationCost_{jkt} + \varepsilon_{jkt} & if\ Active_{jkt} = 1 \\ \varepsilon_{jkt} & if\ Active_{jkt} = 0 \end{cases}$$

where ε_{jkt} refers to action-specific random shocks for individuals and is assumed to follow a Type-I extreme value distribution with $\varepsilon_{jkt} \sim EV(\mu_\varepsilon, \sigma_\varepsilon)$. Prior to taking an action, the host can only form an expectation on the payoff received in current period as the revenue is realized only at the end of the period. Let $\widetilde{\Pi}_{jkt} = E\left(\Pi_{jkt} \middle| a_{jkt}, \left\{\lambda_j^{effort}, \lambda_j^{operate}, X_{jt}, a_{jt}\right\}_{j=1}^{J}, \lambda^{HighImg}, \lambda^{MedImg}\right)$denote the expected payoff, conditional on the set of individual-specific parameters α_j, λ_j for all hosts, the market-share relevant covariates X_{jt} for all properties, and the actions that her peers will take a_{jt} in current period. For identification's purpose, we normalized the mean payoff of 'snoozing the listing' to zero, then Equation (12)) specifies the conditional expected payoff from taking action a_{jt} for *j* in period t. Note for simplicity, we use -*j* to denote the set of individuals excluding *j*.

$$\widetilde{\Pi}_{jkt}(a_{jt}|\{\lambda_j^{effort},\lambda_j^{operate},\lambda_{-j}^{effort},\lambda_{-j}^{operate},\lambda^{HighImg},\lambda^{MedImg},X_{jt},X_{-jt},a_{-jt}\} \tag{12}$$

$$= \begin{cases} ReserveDays_{jt}(a_{jt},X_{jt},X_{-jt},a_{-jt})\cdot\left(NightlyRate_{jt}+\lambda_j^{effort}I\{a_{jt}(HighEffort)=1\}\right) \\ +\lambda^{MedImg}I\{a_{jt}(MedImage)=1\}\cdot I\{a_{jt-1}(MedImage)=0\} \\ +\lambda^{HighImg}I\{a_{jt}(HighImage)=1\}\cdot I\{a_{jt-1}(HighImage)=0\} \\ +\lambda_j^{operate} & if\ Active_{jt}=1 \\ 0 & if\ Active_{jt}=0\ (snooze\ listing) \end{cases}$$

where $ReserveDays_{jt}(a_{jt},X_{jt},X_{-jt},a_{-jt})$ denotes the number of booked days for *j* in period *t*, supposing she takes action a_{jt} with market-share relevant covariate X_{jt} and her peers take action a_{-jt} with covariate X_{-jt}. Recall that a property's market share is a function of her own and her peers' state (see Equation 6 and 7). $I\{\cdot\}$ is an indicator function and $a_{jt}(\cdot)$ refers to a specific activity in this action. For example, $I\{a_{jt}(MedImage)=1\}$ is 1 if the action of a_{jt} will result in a med-level property image for property *j* in period *t*, and is 0 if otherwise.

3.5. State Variables

This section defines state variables that affect an individual host's payoff over time and discusses the dynamics in the transition of individual states driven by hosts' actions.

The set of state variables for individual *j* at time *t* is $s_{jt}=(MedImage_{jt-1},HighImage_{jt-1},NumReviews_{jt-1})$. Similarly, let $s_{-jt}=(MedImage_{-jt-1},HighImage_{-jt-1},NumReviews_{-jt-1})$ denote the set of state variables for *j*'s peers at time *t*. Note we write the *t-1* in the subscript for *MedImage*, *HighImage,* and *NumReviews* to emphasize that the state of image and reviews in current period comes from the action and the outcome in the previous period. Lastly, the combination of the set of states for all individuals constitute state. Then, for *j* at time *t*, her strategy profile depends on $s_t=(s_{jt,}s_{-jt})$, as the not just her own state, but also the states of her peers affect her decisions and state transition.

State Transitions

For each individual *j*, the evolution of her own state at time *t*, s_{jt}, depends on s_{jt} itself and her choice a_{jt}. The individual type $(\lambda_j^{effort},\lambda_j^{operate})$ does not evolve over time. Below, we formalize the evolution of the dynamic states—namely images and number of reviews.

Images. The transition of images is governed by individual's action (choice of images) in every period. For example, if individual *j* chose med-level images in period *t*, i.e., $a_{jt}(MedImage_{jt}) = 1$, then the outcome of this action leads to the state in next period. As a result, in period $t+1$, med-level images will be the individual's state, i.e., $s_{jt+1}(MedImage_{jt+1}) = 1$. If a host chooses to snooze the listing for current period, then next period the image state remains the same. That is:

$$s_{jt+1}(Image_{jt+1}) = \begin{cases} a_{jt}(Image_{jt}) & if\ a_{jt}(Active_{jt}) = 1 \\ s_{jt}(Image_{jt}) & if\ a_{jt}(Active_{jt}) = 0 \end{cases}$$

Note that the evolution of image state of *j*'s peers depends on peers' choices of images in current period, regardless of *j*'s choice.

Number of Reviews. The evolution of number of reviews is straightforward. For each individual *j*, we have following transition rule for state *NumReviews*:

$$NumReview_{jt+1}|NumReview_{jt} = \begin{cases} NumReview_{jt} + NewReview_{jt} & if\ a_{jt}(Active_{jt}) = 1 \\ NumReview_{jt} & if\ a_{jt}(Active_{jt}) = 0 \end{cases}$$

where $NumReview_{jt}$ denotes the number of reviews that *j* has accumulated till the beginning of period *t* and $NewReview_{jt}$ indicates the number of new (added) review in period *t*. If a host chooses to snooze the listing for period *t*, then review state also 'snoozes' for time *t*, as there will be no booking at al.

In above equation, $NewReview_{jt}$ is generated through guests in period *t* who, conditional on the transactions and their stays, are willing to leave a review. Hence, $NewReview_{jt}$ is a function of $\{X_{jt}, X_{-jt}, a_{jt}(MedImage), a_{-jt}(MedImage), a_{jt}(HighImage), a_{-jt}(HighImage)\}$, which affects the current market share for property *j*, s_{jt} . It is also a function of $a_{jt}(MedImage), a_{jt}(HighImage), a_{jt}(HighEffort)$which , combined with property <u>*j*'s quality, affects</u> the likelihood that guests, conditional on their stays, leave a review.

A host, before the property bookings and stay experiences are realized, he can only form an expectation on the number of generated new review, given the states and actions of her and her peers. We assume that hosts have learned the relationship between consumers' likelihood of writing reviews and the realized gap (i.e., departure of property quality from image quality) and the quality of provided guest services. Specifically, we use the empirical relationship (see section 2.3.2 and Equation (3)) to compute the review-writing likelihood (conditional on one's property quality, improving image quality would reduce guest's post-consumption likelihood of writing reviews, see Table 3). To do so, we empirically identify the relationships from our data, i.e., the Law of Motion as shown in Equation (3):

$$WriteProb_{jt} = \alpha_j + \alpha_{high}HighImage_{jt} + \alpha_{med}MedImage_{jt} + \alpha_{effort}Effort_{jt} + \alpha_{price}NightlyRate_{jt} + \alpha_4 Control_{jt} + Period_t$$

where α_j is property fixed effect term, which is also a function of property quality. Using the Law-of-Motion as reported in Table 3, a host can approximate the likelihood that her guests will write reviews after their stays, given her α_j as well as her choice of image quality and service effort.

Hence, the choice of images not only affect one's current market share, also impact the transition of reviews, as guests' post-consumption are affected by the 'gap' that arise when images create a high expectation. The choice of effort, though unobserved to consumers ex-ante and hence do not affect current market share, will actually impact one's long-term utilities through its control on the review evolution. At last, Equation (13) summarizes the review transition:

$$E\left(NumReview_{jt+1} \middle| NumReview_{jt}, a_{jt}, a_{-jt}, s_{jt}, s_{-jt}, X_{jt}, X_{-jt}\right) = \tag{13}$$

$$\begin{aligned} NumReview_{jt} &+ NumReservations_{jt} \cdot WriteProb_{jt} \\ &= NumReview_{jt} + (\frac{ResearvationDays_{jt}}{MinSatys_{jt}}) \cdot WriteProb_{jt} \\ &= NumReview_{jt} + (MarketShare_{jt} \cdot \frac{Marketsize}{MinSatys_{jt}}) \cdot WriteProb_{jt} \\ &= NumReview_{jt} + s_{jt}(X_{jt}, X_{-jt}, \mu_{jt}, \mu_{-jt}, s_j, s_{-jt}, \Sigma, \{v^i\}_{i=1}^{I}) \\ &\cdot MarketSize/MinSatys_{jt} \\ &\cdot WriteProb_{jt}(a_{jt}(Image_{jt}), a_{jt}(HighEffort_{jt}), \alpha_j) \end{aligned}$$

where α_j is property fixed effect from Equation (3).

In summary, the whole state space for any individual j is at time t, $s_{jt} = \{\alpha_j, MedImage_{jt}, HighImage_{jt}, NumReview_{jt}, MedImage_{-jt}, HighImage_{-jt}, NumReview_{-jt}\}$ all observed.

3.6. Individual's Optimization Problem

We model an individual host's choice of images and service effort as a dynamic optimization problem. On an infinite-time horizon, each individual j chooses an infinite sequence of actions $a_{jt} = \{Image_{jt}, Effort_{jt}, Active_{jt}\}_{t=1}^{\infty}$ to maximize the sum of expected life-time payoff:

$$\max_{\{Image_{jt}, Effort_{jt}, Active_{jt}\}_{t=1}^{\infty}} E_{\{s'_{jt}, s'_{-jt}\}}\{\sum_{t=0}^{\infty} \beta^t \cdot (\widetilde{\Pi}_{jt}\left(Image_{jt}, Effort_{jt}, Active_{jt} \middle| s_{jt}, s_{-jt}\right)\}$$

where s'_{jt}, s'_{-jt} denotes the transitioned individual state in the next period for j and for her peers, respectively. $\widetilde{\Pi}_{jt}$ is the expected per-period payoff (expectation over the idiosyncratic payoff shocks ε_{jkt}, see Equation 11):

$$\widetilde{\Pi}_{jt} = (Active_{jt} == 0) * 0 + (Active_{jt} == 1) * \{(\int \frac{\exp(\mu_{jt} + X_{jt}v^i)}{1 + \sum_{l=1}^{J} \exp(\mu_{lt} + X_{lt}v^i)} \phi(v^i|\mathbf{0}, \Sigma) dv) \cdot \frac{MarketSize}{MinSatys_{jt}} \cdot NightlyRate_{jt} + \lambda_j^{effort} \cdot (Effort_{jt} == HighEffort) + \lambda_j^{MedImg} \cdot (Image_{jt} == MedImage) \cdot (s_{jt}(MedImage) == 0) + \lambda_j^{HighImg} \cdot (Image_{jt} == HighImage) \cdot (s_{jt}(HighImage) == 0) + \lambda_j^{operate}\}$$

The specification of individual's per-period payoff and the state transition rule reveal the strategic interactions across peers—an individual's per-period payoff and decision is a function of her peers' states. That is, an individual need to approximate her peers' action in every period, given their states, as the peers' action affects one's current payoff and the evolution of the states. Moreover, this is a dynamic model in the sense that an individual's optimal decision change over time, as her and the peers' state evolve over time. Lastly, there are two interesting intertemporal trade-offs that worth emphasized. First, a host faces trade-off between posting high-level images to improve present property demand versus forging temporary revenue to grow the reviews and to improve future demand. Posting high-quality images improves the expected payoff for the consumers and thus improves a property's temporal market share. However, a high expectation may also induce a greater dissatisfaction as consumers will be happy about the 'negative gap' between the expectation and the realized property quality. As a result, their likelihood of writing review is reduced. In the long-run, this will hurt the host, as number of reviews plays a significant role in generating demand, particularly when one's peers are growing their review. Second, when choosing the amount of service effort, the host compares the effort cost for her (given her ability of investing effort) with the expected gain from an increased review (due to a better service to the guests) in the future. Providing a good service can improve the guests' post-consumption satisfaction and hence increase their likelihood of writing reviews. As a result, the host can effectively grow the reviews and improve the future demand (and payoff). However, to do so, she must incur a present service cost.

3.7. A Dynamic Game and Equilibrium Concept

As mentioned in previous section, each individual host's decision is dependent on her own state and her peers' state. The hosts (properties), given the current states and each's private shock, make their decisions simultaneously and compete with each other in each period. A proper equilibrium concept for this dynamic game is Markov Perfect Equilibrium (Ericson and Pakes 1995, hereafter MPE).

An MPE is described as a profile of Markov strategies for each individual: $\boldsymbol{\sigma} = (\sigma_1, \sigma_2, \ldots, \sigma_J)$. For any individual *j*, her Markov strategy is operated as a function that maps the current state and *j*'s action-specific private shock into an action[21]:

$$\sigma_j: S \times \varepsilon_j \rightarrow A_j$$

where S denotes the state of all individuals and ε denotes the action-specific private shock that *j* received before making decision. Let $V(s_{jt})$ denote the value function for individual *j* at time *t*:

$$V(S_{jt}) = \max_{\{Image_{jt}, Effort_{jt}, Active_{jt}\}_{t=1}^{\infty}} E_{\{s'_{jt}, s'_{-jt}, \varepsilon_t\}} \left\{ \sum_{t=0}^{\infty} \beta^t \cdot (\Pi_{jt}(Image_{jt}, Effort_{jt}, Active_{jt} | s_{jt}, s_{-jt}) \right\}$$

In an infinite-horizon optimization problem, the above equation can be solved through Bellman Equation (Bellman 1957):

$$\begin{aligned} V_j(s_j, s_{-j}; \boldsymbol{\sigma}) = {} & E_\varepsilon[\Pi_j(\boldsymbol{\sigma}(s_j, s_{-j}, \varepsilon), s_j, s_{-j}, \varepsilon_j) \\ & + \beta \int V_j(s'_j, s'_{-j}; \boldsymbol{\sigma}) dP(s'_j, s'_{-j} | \boldsymbol{\sigma}(s_j, s_{-j}), s_j, s_{-j}) | s_j, s_{-j}] \end{aligned} \tag{14}$$

Note we dropped time subscript *t* as the Markov strategy does not dependent on time. In Equation (14), β is the common discount factor with $0 \leq \beta < 1$. s_j, s_{-j} denotes to the state of *j* and her peers -*j*, respectively. (s'_j, s'_{-j}) denote to the states in the next period, conditional on the current state (s_j, s_{-j}) and the actions that all individuals take, assuming their actions are governed by the Markov strategy.

In an MPE, every host *j*, given the Markov strategy profile $\boldsymbol{\sigma}$, will choose the action that maximized the discounted life-time payoff. That is, for a profile $\boldsymbol{\sigma}$ to be an MPE, an individual would not choose an alternative strategy σ'_j, given that her peers follow $\boldsymbol{\sigma_{-j}}$, that is,

$$V_j(s_j, s_{-j}; \sigma_j, \boldsymbol{\sigma_{-j}}) \geq V_j(s_j, s_{-j}; \sigma'_j, \boldsymbol{\sigma_{-j}})$$

However, one challenge arises for computing MPE in our context. Specifically, because of the large number of individuals and the huge state space, solving for MPE is computationally intractable[22]. To resolve the issue of curse of dimensionality, Weintraub et al. (2008, 2010) proposed an approximation of MPE—Oblivious Equilibrium (OE). OE is developed for a market with large number of players (for example, see

[21] To be specific, we consider symmetric and anonymous pure strategy.

[22] Consider the number of individual's state first. Each property at each period, has 3 possible states of image and 301 possible states of reviews (we truncate number of reviews at 300, as the observation with *NumReviws*>300 is less than 1%, hence reviews can vary between 0-300). Hence, for each individual, the number of her own state is 3*301=903. Then the whole state space (including one and her peers) has a dimension of $(903)^{958}$.

Huang et al. 2015 for an application of OE in the context of enterprise social media). The key notion is that in such a market, the simultaneous changes in each individual's moves can be averaged out. As a result, the average industry state either remains stationary over time or can be tracked as a deterministic trajectory changing with a stationary (steady) pace (Weintraub et al. 2008, 2010). Thus, each individual does not need to track everyone's state over time. Instead, it is sufficient for one to make a near-optimal decision by considering only her own state and the average industry state. As Farias et al. (2012) demonstrated, OE can approximate MPE very well, particularly if the market is not too concentrated and the number of individual players is not too few. Therefore, OE fits our setting and should give us a sufficiently good approximation to MPE for the following three reasons. First, we have a large number of individual Airbnb host, none of which is likely to dominate the market. Second, given the large number of hosts and the small role each host plays on the market, in the reality it is difficult for hosts to track all hosts' states in every period. Third, in our data, we observe that the average state of images stay relatively constant over time, and the average number of reviews grow steadily over time (each period increases approximately 2 reviews). Thus, in our study we use OE to approximate MPE.

3.8. Unobserved Heterogeneity

In our data, individual hosts exhibit various responses/actions over time. For example, some hosts tend to invest more effort and provide a good service, while others frequently provide relative poor service. Some hosts choose high quality images to post on property page, while other tend to stay with relative poor images, even AirBnB was offering professional images for free. We hypothesize such different responses come through the heterogeneity in the consumers' preference on revenues and in their ability of investing service effort. Following the stream of literature on hierarchical Bayesian framework (Ching et al. 2012, Rossi et al. 2005), we incorporate individual heterogeneity into our structural model by imposing a distributional assumption on the individual-specific parameters $(\lambda_j^{effort}, \lambda_j^{operate})$. Here λ_j^{effort} is individual j's marginal cost of investing high service effort and $\lambda_j^{operate}$ is j's cost of operating an active Airbnb listing. Specifically, $(\lambda_j^{effort}, \lambda_j^{operate})$. are assumed to be independent draws from a multivariate normal distribution (MVN) with mean ρ and covariance matrix Σ_ρ, i.e., $(\lambda_j^{effort}, \lambda_j^{operate}) \sim MVN(\bar{\lambda}, \Sigma_\lambda)$.

The individual heterogeneities are time-persistent and unobserved to researchers. We assume that individuals know each other's $(\lambda_j^{effort}, \lambda_j^{operate})$ and hence we estimate a complete information game with unobserved heterogeneity. Though $(\lambda_j^{effort}, \lambda_j^{operate})$is not explicitly specified by each host, individual hosts could learn or infer the distribution of $(\lambda_{-j}^{effort}, \lambda_{-j}^{operate})$through their own experience and relevant information of their peers such as host experiences, property locations, and property types etc.

4. Estimation Strategy and Identification

The model primitives (unknown parameters) include $\{\{\theta_k^i\}_{k=1}^{K}, \{\mu_t, \eta_t, \xi_t\}_{t=1}^{12}, \delta, \Omega\}$ from the property market-share model (demand side) and $\left\{\gamma_j^{effort}, \gamma_j^{operate}\right\}_{j=1}^{J=958}, \lambda_j^{MedImg}, \lambda_j^{HighImg}, \sigma_\varepsilon\}$ from the dynamic game model (supply side).

On the demand side, θ_k^i refers to individual-level coefficient (preference) for simulated consumer *i* on the *k*th characteristics, where *k*=1...12 indicates a series of dummy variables for the 12 months from January to December in a year (i.e., monthly fixed effects). For *k*=13…*K*, the corresponding characteristics for property *j* in period *t* are {$EntireHome_j$, $Apartment_j$, $Bedrooms_j$, $Bthrooms_j$, $Beds_j$, $MaxGuests_{jt}$, $MinimumStay_{jt}$, $NumReviews_{jt}$, $MedImage_{jt}$, $HighImage_{jt}$, $DriveTime_j$, $WalkScore_j$, $NightlyRate_{jt}$},. $\mu_t = \left\{\mu_{jt}\right\}_{\mathrm{j=1}}^{J}$ indicates a set of 'mean utilities' for properties *j*=1…*J* (mean utility is normalized to zero for outside option *j*=0) in period *t*. Mean utility μ_{jt} capture the overall preference in the population on property *j*. $\eta_t = \left\{\eta_{jt}\right\}_{\mathrm{j=1}}^{J}$ indicates a set of aggregate demand shocks to properties *j*=1…*J* in period *t*. Note that for each η_{jt} this demand shock is common across all consumers *i=1…I*. $\xi_t = \left\{\xi_{jt}^{price}, \xi_{jt}^{image}, \xi_{jt}^{review}\right\}_{\mathrm{j=1}}^{J}$ indicates a set of stochastic shock, ξ_{jt}, that is correlated with the both the aggregate demand common shock η_{jt} and the endogenous variables (i.e., property price, choice image quality, and number of reviews in our study). $\delta = \{\delta^{price}, \delta^{image}, \delta^{review}\}$ relates endogenous price, image quality, and number of reviews, to their instruments , respectively, each through a linear regression. $\Omega = \{\Omega^{price}, \Omega^{image}, \Omega^{review}\}$ specifies that endogeneity source—how demand shock η_t correlates with the stochastic shocks ξ_t–through a bivariate linear regression: $\begin{pmatrix}\xi_{tj}\\ \eta_{jt}\end{pmatrix} \sim N\left(\begin{pmatrix}0\\0\end{pmatrix}, \Omega = \begin{pmatrix}\Omega_{11} & \Omega_{12}\\ \Omega_{21} & \Omega_{22}\end{pmatrix}\right)$.

Following Jiang et al. (2009), we employ a hierarchical Bayesian framework and impose distributional assumptions on the individual-level parameters. Specifically, we assume that $\theta^i \sim MVN(\bar{\theta}, \Sigma)$, where we allow that consumers preferences to be correlated, captured by the off-diagonal elements of *K* by *K* variance-covariance matrix Σ. Mean utility $\mu_{jt} = X_{jt}\bar{\theta}$. Equation (15) address the endogeneity concerns regarding property price, choice of image quality, and number of reviews, on the demand side, by describing the relationships among ξ_t, η_t, and δ with a set of bivariate normal distributions and a set of linear regressions where Z are the set of instruments for each endogenous variable:

$$\begin{cases} Price_{jt} = \delta^{price} Z^{price} \\ ImageQuality_{jt} = \delta^{image} Z^{image} \\ NumReview_{jt} = \delta^{review} Z^{review} \end{cases} \quad (15)$$

$$\begin{cases} \begin{pmatrix} \xi_{jt}^{price} \\ \eta_{jt} \end{pmatrix} \sim N\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \Omega^{price} = \begin{pmatrix} \Omega_{11}^{price} & \Omega_{12}^{price} \\ \Omega_{21}^{price} & \Omega_{22}^{price} \end{pmatrix} \right) \\ \begin{pmatrix} \xi_{jt}^{image} \\ \eta_{jt} \end{pmatrix} \sim N\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \Omega^{image} = \begin{pmatrix} \Omega_{11}^{image} & \Omega_{12}^{image} \\ \Omega_{21}^{image} & \Omega_{22}^{image} \end{pmatrix} \right) \\ \begin{pmatrix} \xi_{jt}^{review} \\ \eta_{jt} \end{pmatrix} \sim N\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \Omega^{review} = \begin{pmatrix} \Omega_{11}^{review} & \Omega_{12}^{review} \\ \Omega_{21}^{review} & \Omega_{22}^{review} \end{pmatrix} \right) \end{cases}$$

Hence, the estimated parameters on the property market-share model are: $\{\bar{\theta}, \Sigma, \Omega, \delta\}$.

On the supply side, λ_j^{effort} indicates that cost of investing a high effort, relative to a low effort, in providing service to the guests for host *j*. $\lambda_j^{operate}$ captures the cost, of operating an active AirBnB listing for host *j*. The operating cost includes costs related to listing page managing, social cost, and the potential oppournitiy cost (e.g., a property, if not listed on AirBnB, could be rented out via another home-sharing platform or long-term rental market). We assume individual-specific coefficients ($\lambda_j^{effort}, \lambda_j^{operate}$) to be independent draws from a multivariate normal distribution i.e., $\left(\lambda_j^{effort}, \lambda_j^{operate}\right) \sim MVN(\bar{\lambda}, \Sigma_\lambda) = \begin{pmatrix} \Sigma_\lambda^{11} & \Sigma_\lambda^{12} \\ \Sigma_\lambda^{21} & \Sigma_\lambda^{22} \end{pmatrix}$. γ_{med} and γ_{high} refers to the cost, common across hots, of using (including preparing, making, and posting) medium-level and high-level quality of images, respectively. Lastly, we estimate the standard deviation in the distribution of the action-specific idiosyncratic shocks, where we assume $\varepsilon_{\mathrm{jkt}} \sim EV(\mu_\varepsilon, \sigma_\varepsilon)$. Hence, the estimated parameters on the dynamic supply model are: $\{\bar{\lambda}, \Sigma_\lambda, \lambda^{MedImg}, \lambda^{HighImg}, \sigma_\varepsilon\}$.

4.1. Identification

On the demand side, the market-share model is identified by the variations across time. The intuition is that, different combinations of the parameter values would give different market equilibrium outcomes (observed market-share). For example, since the mean utility of outside option is normalized to be zero and constant over time, then the mean utility for property *j* at time *t* is identified through *j*'s market share at time *t*. Since a property *j* in period *t* is viewed as a bundle of property attributes $X_{jt} = \{X^k\}_{k=1}^K$, the coefficients for each X^k is then identified through the variations in market shares across different values of X^k. Similarly, the monthly fixed effects (i.e., X^k for *k*=1,2, …,12) is identified from changes in market shares across each month. For the endogenous variables price, the identification of property substitution patterns (i.e., the price self- and cross- elasticities) relies on the variation in the instrumental variables (lagged price and lagged number of reviews) that are assumed to be exogeneous to the aggregate demand shocks.

On the supply side, we first fix discount factor β to a constant between 0 and 1, as it cannot be jointly identified with the model primitives (Rust 1994, Magnac and Thesmar 2002). Since our model does not satisfy exclusion restriction and we are more interested in knowing how one's heterogenous coefficients would affect her choice of images and investment in service than in identifying β, we chose to fix β to 0.95 and identify other model primitives[23].

We start with discussing the identification of photography cost and the service effort cost. Identifying the cost associated with medium-level images is straightforward. All properties must incur a cost of λ^{MedImg} if they update to a state of medium-level images, this is because Airbnb doesn't provide medium-quality images for free and hence hosts must pay on their own to have medium-quality images. As a result, λ^{MedImg} can be identified from the overall frequency that hosts if they transitioned from a quality level other than medium to level medium-level. Identifying $\lambda^{HighImg}$ and λ_j^{effort} is a more complex. This is because there is an observational equivalence for two scenarios: 1) a host has a high cost of posting high-quality images and a low cost of investing service effort, and 2) a host has a low cost of posting high-quality images and a high cost of investing service effort. Both of the scenarios lead to the same observation: the host does not use high-quality images and as a result, he does not need to invest high-level effort in providing service. fact that helps us to separate the two scenarios is that Airbnb's professional photography program provides high-quality images to the same property for only once. For the case 1), the host would take the high-quality images for free (for the first time) and invest high-effort in service. When they do not qualify for the free service, they would not use high-quality images. For the case 2), we would not observe that the host's choice of image quality to vary a lot when they qualify versus not qualify for the free service. In the reality, we do not observe sufficient temporal variation in the image quality choices to help us to identify individual-specific photography cost (at least for the one-year panel of data). Hence, we identify a common $\lambda^{HighImg}$ instead. However, the logic of separating $\lambda^{HighImg}$ and λ_j^{effort} is the same: recall that approximately 30% of the properties in our sample had used the program by the time our observation started (i.e., by January 2016). Then for these properties if they updated to a state of high-level images, they must incur a cost of $\lambda^{HighImg}$ as they cannot request a free photography service again. For other properties, the cost of posting high-level (professional) images are free, as they still qualify for requesting a free professional photography service from Airbnb. As a result, $\lambda^{HighImg}$ can be identified from the overall

[23] The main findings are insensitive to alternative discount factors we tested (0.9, 0.975, 0.995).

frequency that hosts, who Airbnb will not provide images at quality level q for free, if they transitioned from a quality level q' to level q (where $q' \neq q$ and q is high-level). Conditional on identified λ^{MedImg} and $\lambda^{HighImg}$, the variation in one's choice of service effort across periods with the same expected number of reviews in the next period helps us to identify another heterogeneous parameter λ_j^{effort}—the marginal cost of investing high service effort. To illustrate, recall that the probability that consumers write reviews depend on current number of reviews, expected 'gap' between chosen image quality and property quality, and the service effort. Hence, if two hosts have the same state and expected gap in a particular period, however one chose to invest high-effort and another invested low effort, then likely the latter has a high cost of investing service effort.

Lastly, conditional on identified λ^{MedImg}, $\lambda^{HighImg}$, and λ_j^{effort}, the operating cost $\lambda_j^{operate}$ is identified through the frequency that hosts observed to operate versus 'snooze' their listings, conditional on the expected revenue and costs For example, with the same expected revenue and costs, if one host is observed to snooze the listing more often than another, then the former is likely to have a higher operation cost (or a higher value of outside option). The standard deviation of idiosyncratic shocks can be identified because we normalize the coefficient for revenue to 1.

4.2. Estimating Demand-Side Model

Jiang et al. (2009) proposed a Bayesian approach of estimating an aggregated market share (BLP) model. The model is estimated using MCMC (Monte Carlo Markov Chain) algorithm.

Given the distributional assumptions on endogenous variable P and demand shock η, using Change-of-Variable Theorem, we derive the joint distribution of market share ms_t and P_t :

$$\pi(P_t, s_t|\bar{\theta}, \Sigma, \Omega, \delta) = \pi(\xi_t, \eta_t|\bar{\theta}, \Sigma, \Omega, \delta) J_{(\xi_t,\eta_t \to P_t,s_t)} = \pi(\xi_t, \eta_t|\bar{\theta}, \Sigma, \Omega, \delta)(J_{(P_t,s_t \to \xi_t,\eta_t)})^{-1}$$

where $J_{(P_t,s_t \to \xi_t,\eta_t)} = \begin{Vmatrix} \nabla_{\xi_t} P_t & \nabla_{\eta_t} P_t \\ \nabla_{\xi_t} s_t & \nabla_{\eta_t} s_t \end{Vmatrix}$ is the Jacobian matrix $J(s_t \to \eta_t) = \begin{Vmatrix} \boldsymbol{I} & \boldsymbol{0} \\ \nabla_{\xi_t} s_t & \nabla_{\eta_t} s_t \end{Vmatrix} = \|\nabla_{\eta_t} ms_t\|$.

Furthermore, to ensure that the estimated covariance variance-matrix Σ is positive-definite, following the re-parameterization method used in Jiang et al. (2009), we use Cholesky decomposition and write:

$$\Sigma = U'U;\ U = \begin{bmatrix} e^{r_{11}} & r_{12} & \cdots & r_{1K} \\ 0 & e^{r_{22}} & \ddots & \vdots \\ \vdots & \ddots & \ddots & r_{K-1,K} \\ 0 & \cdots & 0 & e^{r_{KK}} \end{bmatrix}$$

Lastly, given the priors on the parameters and likelihood function, the joint posterior distribution of the parameters is[24]:

$$\pi(\bar{\theta},\Sigma,\Omega,\delta\big|\{P_t,s_t,X_t\}_{t=1}^{T}) \propto L(\theta,r,\delta,\Omega)\times\ \pi(\bar{\theta},r,\Omega,\delta) \tag{16}$$
$$= \prod_{t=1}^{T}\{\left(\left\|\begin{matrix}\nabla_{\xi_t}P_t & \nabla_{\eta_t}P_t\\ \nabla_{\xi_t}s_t & \nabla_{\eta_t}s_t\end{matrix}\right\|\right)^{-1}\times\phi(\begin{pmatrix}\xi_{jt}=P_{jt}-Z_{jt}\delta\\ \eta_{jt}=\mu_{jt}-X_{jt}\bar{\theta}\end{pmatrix}|\begin{pmatrix}0\\0\end{pmatrix},\Omega)\}$$
$$\times|V_{\bar{\theta}}|^{-\frac{1}{2}}exp\{-\frac{1}{2}(\bar{\theta}-\theta_0)'V_{\bar{\theta}}^{-1}(\bar{\theta}-\theta_0)\}\times\prod_{l=1}^{K}exp\{-\frac{(r_{ll})^2}{2\sigma_{r_{ll}}^2}\}$$
$$\times\prod_{l=1}^{K-1}\prod_{k=l+1}^{K}exp\{-\frac{(r_{lk})^2}{2\sigma_{r_{off}}^2}\}$$

where $\pi(\bar{\theta},r,\Omega,\delta)$ is specified priors on the parameters. Specifically, for variance-covariance matrix, we specify the priors on $r=\{r_{lk}\}_{l,k=1\ldots K,l\leq k}$ with $r_{ll}\sim N(0,\sigma_{r_{ll}}^2), r_{lk}\sim N(0,\sigma_{r_{off}}^2)$ for the diagonal, and off-diagonal elements in matrix U, respectively. For the population mean for characteristics coefficients $\bar{\theta}$,as written in Equation (16), we specify a multivariate normal distribution prior: $\bar{\theta}\sim MVN(\bar{\theta}_0,V_{\bar{\theta}})$. We specify the following priors for δ and Ω: $\delta\sim MVN(\bar{\delta},V_\delta),\Omega\sim IW(\upsilon_0,V_\Omega)$, where IW indicates an inverse Wishart distribution.

MCMC Estimation Steps

The MCMC estimation steps follows a strategy of Gibbs sampling combined Metropolis steps (Jiang et al. 2009, Rossi et al. 2005). Briefly speaking, in each iteration of the MCMC, we first use Gibbs Sampler to draw the conditionals of $\bar{\theta},\delta,\Omega|r,\{ms_t,P_t,W_t,Z_t\}_{t=1}^{T},\bar{\theta}_0,V_{\bar{\theta}},\bar{\delta},V_\delta,\upsilon_0,V_\Omega$ in a sequence. Then, conditional on updated $\{\bar{\theta},\delta,\Omega\}$, data $\{ms_t,P_t,W_t,Z_t\}_{t=1}^{T}$, and priors $(\sigma_{r_{ll}}^2,\sigma_{r_{off}}^2)$,we update the variance-covariance matrix, Σ, by making draws of r through a Random-Walk (RW) Metropolis chain. Specifically, we draw a proposal of r, given the accepted r in the previous iteration: $r_{new}=r_{old}+MVN(\mathbf{0},\sigma^2D_r)$, where σ^2 is one of $(\sigma_{r_{ll}}^2,\sigma_{r_{off}}^2)$ depending on whether we're drawing a diagonal or off-diagonal element of Σ. D_r is a candidate covariance matrix. r_{new} is either accepted or rejected, based on ratio computed using Equation (16). The intuition is that, if conditional on data, priors, and other parameters updated in the Gibbs sampling step, $\Sigma_{new}(r_{new})$, relative to $\Sigma_{old}(r_{old})$, is closer the true posterior of Σ, then we should have

[24] For the setup of hyper-parameters, we used diffuse priors. In appendix, we describe details on the choices of priors.

$(\bar{\theta}, \Sigma_{new}, \Omega, \delta \big| \{P_t, ms_t, X_t\}_{t=1}^{T}) > (\bar{\theta}, \Sigma_{old}, \Omega, \delta \big| \{P_t, ms_t, X_t\}_{t=1}^{T})$. In appendix we provide detailed technical notes of our estimation steps.

4.3. Estimating Supply-Side Model

Conditional on one's current state $s = (s_{jt}, s_{-jt})$, her mage and effort decisions can be described as sequentially solving a DP problem:

$$\{a_{jkt}\}_{t=0}^{\infty} = \underset{\{a_{jkt}\}_{t=0}^{\infty}}{argmax}\, E_{\varepsilon_{jkt}}\{\sum_{t=0}^{\infty} \cdot (\widetilde{\Pi}_{jkt}(a_{jkt} \big| s_{jt}, s_{-jt}) + \varepsilon_{jkt})\} \tag{17}$$

where $\widetilde{\Pi}_{jkt}$ is property *j*'s expected payoff from choosing action *k* in period *t* and ε_{jkt} is the random shock associated to action *k* that is received before *j* makes a decision.

As discussed in section 3.8., in such dynamic game with many players, computing an MPE is computationally infeasible, hence we use OE to approximate MPE. In a OE, the individual's conditional choice probability is a function of her own state s_{jt} only. The set of states of her peers, s_{-jt}, is captured by tracking an average industry sate $\bar{s}_t$, which reflects the distribution of the number of the reviews across the properties. It can be seen one's action and payoff is influenced by her peers' state—s_{-jt}, as it is the action of *j* and her peers and the subsequent state transitions that determine the average state in the next period. Then solving for an OE provides substantial computational advantage, as it converts a many-agent game problem into a problem similar to single-agent optimization, treating $\bar{s}_t$ as a single state variable that is common across all individuals at time *t*. Thus, one can use any existing estimation method that can be applied to a single-agent discrete-choice dynamic programming (DDP) model to solve for an OE. Widely-used estimation strategy includes the nested fixed-point (NFXP) algorithm (Rust 1987) and conditional choice probability (CCP) based estimation (Hotz and Miller 1993, Aguirregabiria and Mira 2007).

In this paper, we use a Bayesian estimation strategy as this way we can flexibly incorporate individual heterogeneity—a key element in our model—in a hierarchical Bayesian framework (developed by Imai, Jain and Ching (2009), hereafter IJC). IJC algorithm allows estimating a heterogeneous model with a relatively low computational burden. In addition, it overcomes the problem of "curse of dimensionality[25] when approximating the DP solution and avoids the complexity of searching for a global optimum in the space of the data likelihood function (IJC provides DP approximation that is comparable to state-of-the-art

[25] The state space grows exponentially with the dimensionality of sate variables, causing evaluating Bellman operator at every point in the state space infeasible.

likelihood-based approaches, e.g., Keane and Wolpin (1994), Ackerberg (2009). See Ching et al. (2012) for detailed discussions). The advantage of avoiding of searching in the parameter space, which usually requires the use of an optimization tool, is another reason we choose IJC algorithm. As we will discuss in section 4.4., Bayesian estimation approach can be easily combined with parallel computing and GPU computing techniques, without which it would be computationally infeasible given the large number of individuals and state space in our study.

IJC Algorithm

We briefly introduce the logics and estimation procedure in IJC algorithm. In appendix we provide technical notes and details of implementing IJC.

IJC algorithm combines MCMC with DDP approximation, solving for the DP problem and making draws of structural parameters from the posterior distribution simultaneously. At each iteration m in the MCMC, IJC saves the simulated parameter vector θ_{IJC}^{*m} and computes a corresponding pseudo-value function $\widetilde{W}^{m}\left(\theta_{IJC}^{*m}\right)$[26]. A total of the most recent N iterations of $\{\theta_{IJC}^{*m}, \widetilde{W}^{m}\left(\theta_{IJC}^{*m}\right)\}$ are saved. When at new iteration m', the simulated vector $\theta_{IJC}^{*m'}$ is rejected or accepted by comparing the pseudo- posterior likelihood evaluated at the accepted parameters from the previous iteration, $\theta_{IJC}^{m'-1}$, and at the proposed parameters at current iteration, $\theta_{IJC}^{*m'}$. When computing the pseudo-likelihood function, one needs to calculate the choice probability for each choice alternative. Recall that one solves for the DP problem by taking into account the value function (see Equation 17)), hence the likelihood function is also 'pseudo-' because the conditional choice probabilities are computed based on pseudo-value functions $\{\widetilde{W}^{n}\left(\theta_{IJC}^{*n}\right)\}_{n=m'-1}^{n=m'-N}\}$ saved in the past N iterations. Specifically, $\widetilde{W}^{m'}\left(\theta_{IJC}^{*m'}\right)$ is approximated by computing a (kernel-based) weighted average of the past N history draws of $\{\widetilde{W}^{n}\left(\theta_{IJC}^{*n}\right), \theta_{IJC}^{*n}\}_{n=m'-1}^{n=m'-N}\}$, with the more weights attributed to history that have θ_{IJC}^{*n} closer to current draw $\theta_{IJC}^{*m'}$. As Imai et al. (2009) proved, such an interactive steps of simulating parameter vector through the pseud0-Markov chain can effectively approach a steady state (after burn-in), where most of the structural parameters will be drawn from a distribution close to the true posterior distribution of the parameter vector.

[26] The pseudo-value function is obtained by applying the Bellman operator (i.e., solving for the value function) at the trial parameter vector. It is called 'pseudo' as the functions are evaluated at the simulated parameter vector not at the true parameter vector. Here the * denotes that this is proposed parameter (regardless of whether it was accepted or rejected) at that iteration.

In summary, we a Gibbs Sample to sequentially simulate parameters of $\left(\left\{\lambda_j^{effort}, \lambda_j^{operate}\right\}_{j=1}^{J}, \lambda, \Sigma_\lambda, \lambda^{MedImg}, \lambda^{HighImg}, \sigma_\varepsilon\right)$, with $\left(\lambda_j^{effort}, \lambda_j^{operate}\right) \sim MVN(\lambda, \Sigma_\lambda)$. At each iteration *m*, we have the history of the drawn parameters and the associated pseudo-value functions:

$$\{\widetilde{W}^n\left(\cdot; \left\{\lambda_j^{effort*n}, \lambda_j^{operate*n}\right\}_{j=1}^{J}, \gamma_{med}^{*n}, \gamma_{high}^{*n}, \sigma_\varepsilon^{*n})\right), \left\{\lambda_j^{effort*n}, \lambda_j^{operate*n}\right\}_{j=1}^{J}, \gamma_{med}^{*n}, \gamma_{high}^{*n}, \sigma_\varepsilon^{*n}, \bar{s}^{*n}\}_{n=m-N}^{n=m-1}\}$$

where $\widetilde{W}^n(\cdot ;)$ indicates the pseudo-value functions at all possible state and $\bar{s}^{*n}$ is proposed industry average state. Then each iteration *m* we simulate parameters of $(\left\{\lambda_j^{effort}, \lambda_j^{operate}\right\}_{j=1}^{J}$ using Metropolis-Hasting to draw a proposal, where we evaluate the pseudo-likelihood function given the observed hosts' choices and the choice-probabilities computed using $\{\widetilde{W}^n(\cdot ;), \left\{\lambda_j^{effort}, \lambda_j^{operate}\right\}_{j=1}^{J}, \lambda, \Sigma_\lambda, \lambda^{MedImg}, \lambda^{HighImg}, \sigma_\varepsilon\}$. In appendix, we provided technical details on the IJC estimation steps.

4.4. Computation Challenges

Due to the large number of properties (900+) in our sample, estimating the demand and the supply model is computationally challenging. We solve this issue leveraging parallel GPU (Graphical Processing Unit) computing. Specifically, as Bayesian estimation does not involve searching optima in the parameter space, we easily implement the estimation (MCMC) leveraging GPU computing and distribute the computation for (independent) individuals to multiple cores for parallel processing. As GPU is specialized in vector/matrix (floating) operations, we were able to accelerate the estimation by as much as 4 to 60 times, with the more number of products, the more computation advantage of GPU computing.

5. Estimation Results

We report and discuss the estimation results, starting with the estimated coefficients of the demand side (property market-share estimation).

5.1. Property Market Share Estimates (Airbnb Demand)

Table 4 presents the estimation results of the property market share model (see Equation 6). The model is estimated in a hierarchical Bayesian framework where Markov Chain Monte Carlo (MCMC) is used to make draws from the posterior distribution of the parameters $\{\bar{\theta}, \Sigma, \Omega, \delta\}$. Thus, we report the mean and the standard deviation of the posterior draws[27]. We performed MCMC diagnostics, by inspection of time-series plots with different starting points (Gilbride and Allenby 2004). The chains reached a common stationary convergence, suggesting that the chains have 'forgotten' the initial points and are drawn from the posterior distribution of $\{\bar{\theta}, \Sigma, \Omega, \delta\}$.

Several interesting findings in Table 4 note highlighting here. First, the mean coefficients on the month and area (neighborhood) fixed effect terms suggest significant dynamic patterns in the seasonality. In general, summer season (May-October) is peak season—February on average has the least demand (mean month effect=-3.803) with September has the highest demand (mean month effect=-1.106), and the area of East Harlem attracts the most demand (mean area effect=-1.516) while area of *Lower Manhattan* on average has the least popularity (mean area effect=-3.7675). Second, the positive and significant coefficient of number of reviews (mean=0.641, std.= 0.093) indicates that number of reviews helps to generate more bookings (demand). Third, the negative population mean (-0.623) of the coefficient for the driving (commute) to downtown area suggests that location plays a role when travelers choose lodging alternatives. Specifically, an average traveler to Manhattan prefers to stay somewhere close to the downtown area, possibly due a convenient public transit or concentration of attractions/restaurants/etc.—consistent with the positive coefficient of walk score (mean=0.486). Lastly, the positive coefficients of two image dummy variables confirm that good images, compared to low-quality images, can generate more bookings in current month. Specifically, as expected, the impact of high-level images (mean=0.959) is greater than the med-level images (mean=0.714).

Table 4 Estimated Property Market Share Parameters

VARIABLES #	**Estimates** +	
	Population Mean	**Population Std. Dev.**
Preferences on Property Characteristics: $\boldsymbol{\bar{\theta}}, diag(\boldsymbol{\Sigma})$		
EntireHome	0.5048	0.2707
	(0.1024)	(0.0220)
Apartment	-0.8723	0.3305
	(0.1630)	(0.0531)
Bedrooms	0.1945	0.3977
	(0.0969)	(0.1166)
Bathrooms	0.0441	0.3112
	(0.1545)	(0.0118)
Beds	-0.2037	0.1118
	(0.0704)	(0.0204)
MaxGuests	-0.1062	0.0497

[27] We run a total of 30000 MCMC iterations and drop the first 25000 for burn-in. H=1000 individual 'consumers' were simulated.

	(0.0451)	(0.0054)
MinStays	-0.7346	0.3227
	(0.0567)	(0.0148)
NumReviews (scaled by 1/10)	0.6407	0.2113
	(0.0934)	(0.0317)
MedImage	0.7137	0.3857
	(0.1035)	(0.0349)
HighImage	0.9592	0.3164
	(0.0997)	(0.0289)
DriveTime (100 mins)	-0.62267	0.2731
	(0.2808)	(0.0245)
WalkScore (1/100)	0.4855	0.3468
	(0.2971)	(0.0132)
log(NightlyRate)	-4.004	0.5015
	(0.1043)	(0.0140)
Preferences on Monthly Effects (i.e., Seasonality): $\bar{\boldsymbol{\theta}}$, $diag(\boldsymbol{\Sigma})$		
January	-2.9598	0.1397
	(0.2296)	(0.0240)
February	-3.8027	0.1382
	(0.2276)	(0.0762)
March	-2.4117	0.1712
	(0.2297)	(0.0294)
April	-1.7310	0.2008
	(0.2303)	(0.0067)
May	-1.3084	0.1868
	(0.2313)	(0.054)
June	-1.7538	0.1623
	(0.23077)	(0.0198)
July	-2.1646	0.2048
	(0.2316)	(0.0226)
August	-1.9537	0.1873
	(0.2317)	(0.0128)
September	-1.1059	0.2146
	(0.2373)	(0.0061)
October	-1.5274	0.2018
	(0.2357)	(0.0381)
November	-2.5145	0.2077
	(0.2347)	(0.0091)
December	-2.0075	0.2549
	(0.23437)	(0.0064))
Preferences on Area Effects (i.e., Neighborhood Popularity): $\bar{\boldsymbol{\theta}}$, $diag(\boldsymbol{\Sigma})$		
Central Harlem	-3.0489	0.4122
	(0.33926)	(0.0581)
Chelsea and Clinton	-2.1104	0.5739
	(0.3478)	(0.0830)
East Harlem	-1.5159	0.3784
	(0.2479)	(0.0826)
Gramercy Park and Murray Hill	-2.5231	0.5767
	(0.3438)	(0.1592)
Greenwich Village and Soho	-2.2557	0.6069

	(0.3494)	(0.1083)
Inwood and Washington Heights	-3.7201	0.5028
	(0.2389)	(0.0872)
Lower East Side	-2.7561	0.6121
	(0.3449)	(0.1034)
Lower Manhattan	-3.7675	0.8407
	(0.4508)	(0.1226)
Upper East Side	-2.5014	0.4312
	(0.2480)	(0.0843)
Upper West Side	-2.6499	0.5957
	(0.2454)	(0.1183)
Covariance Matrix of $\boldsymbol{\xi_{jt}}$ and $\boldsymbol{\eta_{jt}}$: $\boldsymbol{\Omega}$		
Ω_{11}^{price}	0.2193	
	(0.0029)	
Ω_{12}^{price} (Ω_{21}^{price})	0.0421	
	(0.0126)	
Ω_{22}^{price}	13.6129	
	(0.1792)	
Ω_{11}^{image}	0.1158	
	(0.0015)	
Ω_{12}^{image} (Ω_{21}^{image})	0.0352	
	(0.0158)	
Ω_{22}^{image}	11.6165	
	(0.2793)	
Ω_{11}^{review}	0.3350	
	(0.0281)	
Ω_{12}^{review} (Ω_{21}^{review})	0.0394	
	(0.0281)	
Ω_{22}^{review}	17.2284	
	(0.2532)	
Observations 8862		

+ Posterior means of estimates are computed on the 25000th~30000th draws.
+ Table presents the posterior mean of each estimate, with the posterior standard deviation reported in the parenthesis below each posterior mean estimate.
# Note due to limited space, for $\boldsymbol{\Sigma}$ only diagonal elements (i.e., population variance of the coefficients) are presented.

Traveler Preferences' Heterogeneity and Correlation

We explore heterogeneity and correlations in the identified individual preferences. Individual heterogeneity can be obtained from the population variance of the coefficients for each property characteristics, i.e., the diagonal of $\boldsymbol{\Sigma}$ (recall that individual-specific parameter $\theta_i \sim N(\overline{\boldsymbol{\theta}}, \boldsymbol{\Sigma})$). The correlation between every two population preferences is reflected through the off-diagonal of $\boldsymbol{\Sigma}$. Since in our estimation, we simulated H=1000 individual consumers over which we integrate to compute the market shares, we can simply use the identified coefficients of the H individuals to plot and exhibit the heterogeneity and correlation. Here we focused on the correlation between consumers' preferences (i.e., coefficients for each property attribute) on property image quality and number of guest reviews, as this correlation highlights an interesting trade-off problem for the hosts.

To examine the preference correlations, we look at how guests' taste sensitivity in property images is correlated with the sensitivity in number of reviews. In Figure 2 we provide a scatter plot between guests' review sensitivity (horizontal axis) and image sensitivity (vertical axis)[28]. The correlation coefficient of the two sets of individual parameters is 0.302 (p-value <0.001), suggesting that consumers who give more weightages to image quality tend to be those who value the number reviews more. This positive correlation reveals an interesting 'trade-off' problem for AirBnB hosts—it is the same set of consumers who reward the hosts in the short-run that will punish him in the long-run. To see why this is happening, suppose a host is facing a pool of consumers who are 'sensitive' to the quality of property images (i.e., elasticity of image quality is high). Then posting high-quality images would be more effective in generating present bookings for this host then for other hosts. However, her consumers (potential guests) are those who also care about the number of reviews. As a result, the 'penalty' of unable to get more reviews due to images-induce high expectations would also be greater for her than for others if she is unable to meet guests' expectations. Thus, the incentive of posting high-quality images may be reduced by the risk of penalty. As we discuss in the next section, it is the real quality of the property and host's ability in investing service that determines what is the optimal choice of images.

Figure 2 Scatter Plots of Individual Consumers' Coefficients

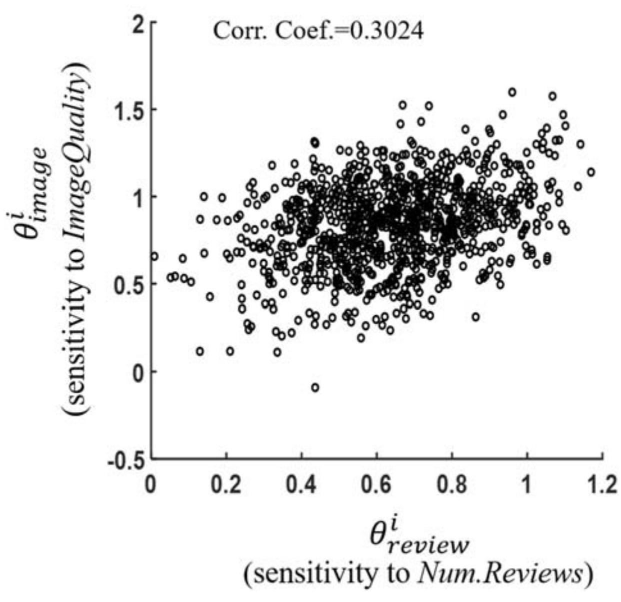

[28] In the model we have two dummy variables for property images: *MedImage*—indicating a med-level image, and *HighImage*—indicating a high-level image. The scatter plot uses an average of the coefficients for the two variables.

5.2. Host Choice of Images and Effort Estimates (AirBnB Supply)

In Table 5 we report the estimation results of the dynamic game model on the supply side (hosts' optimization problem on choice of images and service effort, see Equation 12) and 14). Similar to the property market share estimation (section 5.1), we employ a hierarchical Bayesian framework and draw Markov chains from the posterior distribution of the parameters ($\left\{\lambda_j^{operate}, \lambda_j^{effort}\right\}_{j=1}^{J}, \lambda_{med}, \lambda_{high}$). Further, we assume that individual's operation cost, $\lambda_j^{operate}$, and on marginal service effort cost, λ_j^{effort}, are drawn from $MVN(\bar{\lambda}, \Sigma_\lambda)$. Thus, in the table we report the population mean and population standard deviation of ($\lambda_j^{operate}, \lambda_j^{effort}$). We performed MCMC diagnostics and confirmed that common stationary convergence was achieved. We use the draws after having achieved convergence to compute the sample posterior means and standard deviations of the parameters[29].

The negative coefficient of effort costs confirms our intuition that investing in service effort is costly for hosts. Particularly, if evaluating its impact in monetary-term, we obtain a marginal effort cost equivalent to be $0.3244*100=$32.4/day[30]. That is, for an average host, if consider per-day return, she would be willing to invest a high effort in providing service to the guests only if she could charge an extra $32.4 per day. Furthermore, the population standard deviation in the estimated effort cost suggests there exists heterogeneity in the consumers' ability of investing effort in providing guest service.

The estimated common photography costs suggest that posting good images, relative to posting low-quality images are costly. As expected, having high-level images cost more than having med-level images. The big cost may explain the observation that hosts, once not qualified for AirBnB's free photography program, rarely use high-quality photos on their own. Also, we should note that, the cost includes but is not limited to the fee for hiring a professional photographer. It includes all the costs associated with the preparation work. For example, one may need to spend time on searching photographers, communicating with them back-and-forth, scheduling a day, preparing the house (organizing and cleaning the place) etc. The estimated $\lambda_j^{operate}$ suggests that the outside option value combined with the operation cost on Airbnb

[29] We run a total of 20000 MCMC iterations and drop the first 10000 for burn-in (the chains started to converge after about 10000 draws). We use a kernel bandwidth of 0.05 for the parameters and a bandwidth of 1 for the industry average state. We store N=500 of past pseudo-value functions for approximating *Emax* functions during the MCMC iterations.

[30] The property revenue in host's per-period payoff function is computed with *ReservationDays* x *NightlyRate*/100, where property nightly rate was scaled by 1/100(i.e., the unit is100 USD). So, the effect of effort cost to an average host equals 0.3244 *100 USD= $32.44 USD.

for an average AirBnB host in Manhattan is $10.2715/*100=$1027.2/month in monetary term[31]. If we consider an average nightly rate of $228, the result suggests that unless renting out her property for 6 days/month, one may prefer to leave AirBnB and choose an outside option (e.g., for long-term lease). The estimated $\lambda_j^{operate}$ suggests that costly listing-managing and low property occupancy are may explain the high dropout rate on AirBnB. This is particularly a concern for AirBnB in areas where the hosts can receive high-value outside options (e.g., Manhattan has strong housing market).

Table 5 Estimated Host Supply Model Parameters

VARIABLES #	Estimates +	
	Posterior Mean	**Population Std. Dev.**
Individual Parameters: $\boldsymbol{\rho}, \boldsymbol{\Sigma_\rho}$		
λ_j^{effort}	-0.3244	0.1488
	(0.0722)	(0.0441)
$\lambda_j^{operate}$	-10.2715	1.2866
	(3.2015)	(0.5877)
Common Parameters		
λ^{MedImg}	-3.018	--
$\lambda^{HighImg}$	-5.975	--
σ_ε	1.6271	--
Observations 11496		

+ Posterior means and standard deviations of estimates are computed on the 10000th~20000th draws.
+ Table presents the posterior mean of each estimate, with the posterior standard deviation reported in the parenthesis below each posterior mean estimate.

6. Policy Simulation

Our structural model on hosts' choice of images and effort, combined with property market share estimations, allows to assess the impact of AirBnB's photography policy on hosts' supply of lodgings as well as on its revenue. In this section, we conduct counterfactual analyses to analyze hosts' choice and their returns under simulated policies. We use the individual hosts in our sample and the use their estimated parameters to solve for the DP problem and 'observe' their behaviors over time, as they were making choices following the solutions of DP. Each policy was implemented for T=24 periods (i.e., two years), starting with their initial state assuming everyone just joined the platform. We implement each policy for 1000 runs and report the averaged outcomes over the 1000 runs.

[31] Since $\gamma_{operate}$ captures two components: one is the cost of managing an active listing, the other is the outside option value (or opportunity cost)

6.1. Should AirBnB Provide Medium-Quality instead of High-Quality Images for Free?

The first policy experiment is motivated by the observation that a large number of hosts did not utilize AirBnB's professional photography program, which offered hosts high-quality property images for free. Our structural model suggests that hosts face temporal trade-off between short-term image-induced bookings and long-term review-induced market share. As a result, hosts may end up using low-quality images as high-quality images may hurt the host in the long-run (hosts may not use medium-quality images due to the cost and relative-low return on property bookings). Hence, we want to see can AirBnB do better by providing medium-quality images for free to all hosts. Specifically, we consider two options: 1) providing high-level images to all hosts for free—we refer this policy option as "current policy", and 2) providing med-level images to all hosts for free—we refer this policy option as "proposed policy 1". The baseline is where AirBnB does not offer any free photography-related program.

In practice, a home-sharing platform such as AirBnB can control the quality level of offered property images by guiding/training their photographers. In the simulation, the policies are implemented by reducing ones' cost of posting high-level (med-level) images to 0 under the current policy (proposed policy 1). In Figure 3 we plot the average outcomes of the 958 properties (vertical axis) over time (horizontal axis) under the two alternative policies as well as the baseline. As expected, AirBnB would do better under both of the two image policies than under the baseline policy (where AirBnB does not implement any image policy).

Figure 3 (a) reports the average property demand (i.e., number of reserved days) across all properties in our sample. As can be seen, the current policy dominated proposed policy 1 in the short-run (for the first four periods). Interestingly, the advantage of current policy, relative the proposed policy 1, vanishes quickly over time. After seven to nine periods, the average property booking under proposed policy 1 is greater. In the long-run, proposed policy 1 outperformed the currently policy (1.3 additional reservation days/month vs 0.8 additional reservation days/month). The interpretation is that, medium-level images, compared to high-level images, despite forming a smaller expected utility for the consumers, has a greater effect on property demand in the long-run as they, with lower risks of creating a dissatisfactory gap, help hosts to obtain new reviews. Moreover, individual hosts who might end up using amateur (low-level) images to avoid the dissatisfactory gap under the current policy, now use free medium-level images to make more revenues under the proposed policy.

Such interpretation is supported by Figure 3 (b), where we plot the evolution of the average number of reviews across all properties. Clearly, the average number of review experienced a steady greater growth under proposed policy 1 than under the current policy. As properties accumulated consumer reviews, the significance of reviews started 'cancelling out' the relative negative effect of medium-quality images compared with high-quality images on property's bookings. In the end, the properties are rented out more since the number of reviews is a key driver in consumers' decisions on lodging options.

Figure 3 Policy Simulation: The Impact of AirBnB's Image Policies on Property Demand and Guest Reviews (*)

(a) (b)

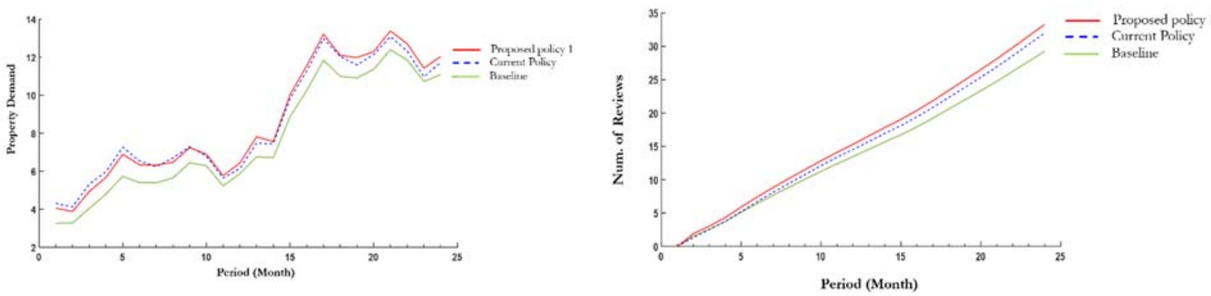

* We plot starting the periods following the initial (entry) period.

Understanding the Mechanism: Host Adoption of Current Policy versus Proposed Policy 1

To better understand the effect of proposed policy 1, here we discuss the adoption rates of Airbnb's current policy versus the proposed policy 1 across the properties with different qualities. To do so, we categorize properties into three quantiles based on their property fixed effects, α_j, identified from Law-of-Motion analysis (see Table 3 and Equation (3)). Note that though the quality of each property j, $PropertyQuality_j$, is unobserved, the property fixed effect α_j is an increasing function of $PropertyQuality_j$. As a result, the categorization gives us three quantile groups of properties whose quality fall into three levels—namely, high-level, medium-level, and low-level.

In Figure 4, we present the proportion of properties that adopted Airbnb's current policy (in blue bar) and that adopted our proposed policy 1 (in red bar) across the three quantiles of property qualities. From the left to the right along the horizontal axis indicates the property quality quantiles vary from low-, medium-, to high- levels. The bars indicate that once medium-quality images are available for free (proposed policy 1), properties that fall into the low- and the medium- quantiles have significantly greater incentive to adopt, than adopting high-quality images (current policy). In contrast, the adoption rates among the properties in the high- quantile do not vary much when the policy shifts from current policy to proposed policy 1. This finding supports our interpretation above—that (some of) the properties in the low- and medium- quantiles would end up using low-quality images under the current policy, because they would face a risk of creating a big negative gap if they used the free high-quality images. Under the proposed policy, they have the motivation to adopt the medium-quality images as the risk of dissatisfaction gap would reduce. Hence, by making medium-quality images for free, Airbnb is able to improve the utilized capacity of the large number of low- and medium- tier properties.

Figure 4 Comparing Adoption Rates of Policies Across Property Quality Quantiles

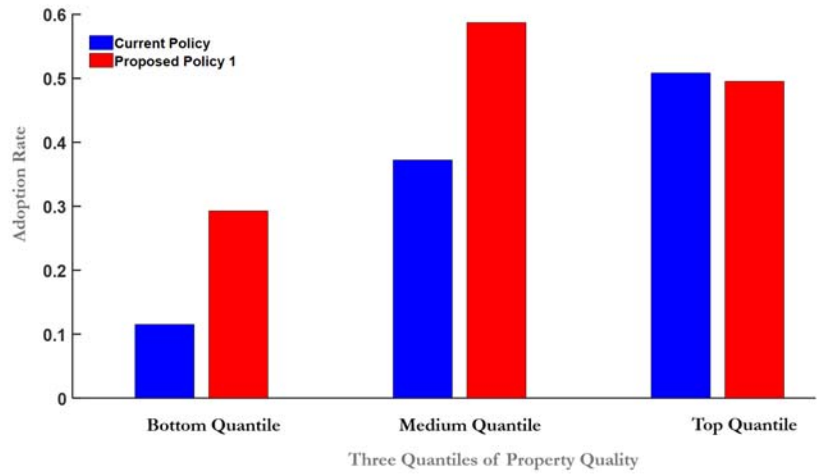

6.2. Should AirBnB Discriminate Hosts on Offering Photography Policy

As shown in Figure 4, properties of which the qualities in low- and medium- quantile have strong incentive to adopt proposed policy 1, as medium-quality images, compared to high-quality images, are better representations (better matched to) of their properties. Since there exists a large degree of heterogeneity in the quality of Airbnb properties, and in hosts' ability of providing good service (see Table 5), intuitively a better strategy is likely to be the one provides different qualities of images to different hosts. Since the true property quality and host's type are unobserved, Airbnb in reality could implement a (second-degree) discrimination on providing the photography program.

Hence, we propose that AirBnB can offer a menu of image quality choices for free. The menu includes both high- and medium- level of property images (images examples are provided) and allow the hosts to self-select which program they want. We hypothesize that such a discrimination would incentive hosts with small λ_j^{effort} and/or high property quality to use high-level images as its relatively less costly for them to provide a good service that matches the high expectation. Hosts with big λ_j^{effort} and/or with lower-quality of properties, on the other hand, will self-select to use the medium-level images as providing a good service that matches high-level images is costly for them. We refer this alternative policy as 'proposed policy 2' and compare its impact to the proposed policy in simulation 1 (i.e., proposed policy 1' in section 6.1).

As Figure 5 (a) shows, the proposed policy 2 (discriminative policy) improves the average property demand (vertical axis) in the long-run, by improving average property demand by 1.7 additional reservation days in a month (versus the long-term impact of 1.3 additional reservation days under the proposed policy 1). Hence, a photography program that provide both medium- and high- quality images for free can most effectively improve the property performance on AirBnB.

Figure 5 Policy Simulation: The Impact of AirBnB's Discriminative Image Policies on Property Demand and Guest Reviews (*)

(a)

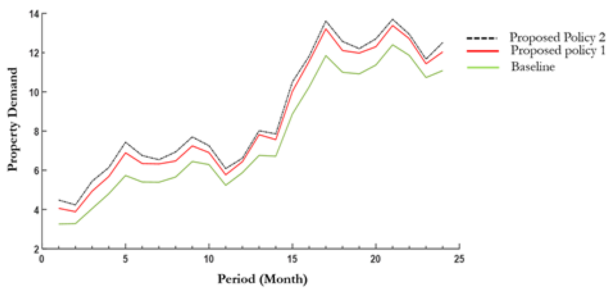

(b)

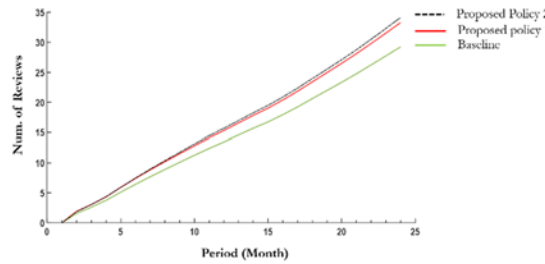

* We plot starting the periods following the initial (entry) period.

7. Discussion and Conclusion

This paper focuses on investigating how the AirBnB hosts make decisions on the quality of photographs to post. Unlike property attributes that cannot be controlled by the hosts such as the type, size, and location of the property, images are in immediate control of the hosts. Previous study found a strong advertisement impact of property images on property booking, with professional images (high-quality), relative to amateur (low-quality) images, increase the present demand by 14.3%. The advertisement effect arises in the context of AirBnB because (1) there exists a large variation in offered properties and (2) consumers rely heavily on visual information in order to ease decision-making. Recognizing the importance of images, AirBnB in 2011 started offering highest quality professional images to all the hosts for free. To AirBnB's surprise, only 30% of the hosts used the AirBnB professional photography program after its launch.

This study provides an explanation for hosts' behavior. We posit that the host's decision on the quality of images to post depends on the following factor: (1) the advertising impact of the photos on present demand, (2) the cost of photos, (3) the impact of the photos on the satisfaction level of the guest post consumption, and (4) the host's cost (ability) of investing effort in providing god service to the guest. As suggested by the reference dependence literature, consumers' satisfaction level in the post consumption depends not only on the actual outcome from consuming the product, but also on her reference point—the individual's pre-consumption expectation. A decrease in the consumers' satisfaction level would in turn adversely impact the future demand through consumer reviews. Since high quality photographs can create unrealistically high expectations, a host would be hesitant to post high quality photographs (even if they were free) if the actual property is not as good as portrayed in the images, especially if the hosts are unable to provide a high level of service to match the high expectations.

Our goal is to disentangle the aforementioned factors that influence the host's decision on the type of photographs to post, and to explore policies that platforms such as AirBnB can employ to improve the property performance. To achieve this goal, we model hosts' periodic (monthly) decisions on the quality of property images to post, and the quality of service to provide. The image decision affects the host's profits in the short run through the cost of posting (preparing, shooting etc.) of images and advertisement effect on present demand. And it affects the host's future profits through its impact on guests' dissatisfaction, which affects consumers' likelihood of writing reviews. The service decision impacts the host's profits in the short run through the service costs and in the long run through impacting the guests' satisfaction level.

We estimate our model exploring a unique panel data spanning one-year for 958 properties in Manhattan (New York City). We observe the dynamics in hosts' choice property images of and provided service quality (proxied with host responsiveness, see section 2.2). We find that guests who value professional images more, also value the number of reviews more, revealing an interesting trade-off problem for the hosts. Further, the estimation results highlight that hosts are heterogeneous in their ability (cost) in investing in service and that it is costly for hosts to post, on their own, images with above-average quality. Policy simulations suggests two proposed photography policies for AirBnB that outperforms its current policy (providing high-level images for free to the hosts) in the long-run. The first proposed policy provides medium-level of images for free to the hosts. Compared to the baseline where no policy is offered, the first proposed policy and the current policy improves the average property demand by 1.3 reservation days per month and 0.8 reservation days in the long-run, respectively. Interestingly, the proposed policy was dominated by the current policy through the first four periods. The interpretation is that, medium-level images, compared to high-level images, despite forming a smaller expected utility for the consumers, has a greater effect on property demand in the long-run as they, with lower risks of creating a dissatisfactory gap, help hosts to obtain new reviews. The second proposed policy offers both high- and medium- level images for free and allow the hosts to self-select to choose which program they want. We show that the second proposed policy performed the best in the long-run, by improving average property demand by 1.7 reservation days per month.

There are a few limitations to this study and directions for future research. First, we do not model hosts' decisions on pricing for model tractability. In addition, individual AirBnB hosts have very limited information to optimally set prices. In a context where prices are likely to be optimally determined (e.g., modeling firms' decisions), modeling pricing decision will help to identify the marginal (production) cost and allows for further implications. Second, due to the computational tractability, we look at only one market (Manhattan) across multiple period. With the future advances in computation power, one may conduct a study across different markets (e.g., Chicago, Miami etc.), which may give interesting insights on potential impact of city and travelers' demographics on hosts' choices of posting images

Reference

Ackerberg, D. A. (2009). A new use of importance sampling to reduce computational burden in simulation estimation". ***Quantitative Marketing and Economics*** 7(4): 343-376.

Aguirregabiria, V. & Mira, P. (2007). Sequential Estimation of Dynamic Discrete Games. ***Econometrica*** 75(1): 1-53.

Berry, S., Levinsohn, J., & Pakes, A. (1995). Automobile Prices in Market Equilibrium. ***Econometrica*** 63(4), 841-890.

Bellman, R. (1957). Dynamic Programming. Princeton University Press.

Ching, A., Imai, S., Ishihara, M., & Jain, N. (2012). A practitioner's guide to Bayesian estimation of discrete choice dynamic programming models. ***Quantitative Marketing and Economics*** 10 (2):151–196.

Davis, P. 2006. Spatial competition in retail markets: Movie theaters. ***RAND Journal of Economics*** 37 (4): 964-982.

Dubé, J., Fox, J., & Su, C. (2012). Improving the Numerical Performance of BLP Static and Dynamic Discrete Choice Random Coefficients Demand Estimation. ***Econometrica*** 80(5), 2231-2267.

Ericson, R., & Pakes, A. (1995). Markov-Perfect Industry Dynamics: A Framework for Empirical Work. ***The Review of Economic Studies*** 62(1), 53-82.

Farias, V., Saure, D., Weintraub, G.Y. (2012). An approximate dynamic programming approach to solving dynamic oligopoly models. ***The RAND Journal of Economics*** 43(2):253–282.

Fradkin, A. & Holtz, D. (2018). The Determinants of Online Review Informativeness: Evidence from Field Experiments on Airbnb ***Working Paper***.

Genesove, D. & Mayer, C. (2001). Loss Aversion And Seller Behavior: Evidence From The Housing Market. ***Quarterly Journal of Economics*** 116(4): 233-1260.

Gilbride, T., & Allenby, G. (2004). A Choice Model with Conjunctive, Disjunctive, and Compensatory Screening Rules. ***Marketing Science*** 23(3), 391-406.

Gul, F. (1991). A Theory of Disappointment Aversion. ***Econometrica*** *59*(3): 667-686.

Huang, Y, Singh, P., Ghose, A. (2015). A Structural Model of Employee Behavioral Dynamics in Enterprise Social Media. ***Management Science*** 61(12):2825-2844.

Hotz, V. & Miller, R. (1993). Conditional Choice Probabilities and the Estimation of Dynamic Models. ***The Review of Economic Studies*** 63(3): 497-529.

Houde, JF. 2012. Spatial Differentiation and Vertical Mergers in Retail Markets for Gasoline. ***American Economic Review*** 102 (5): 2147-82.

Imai, S., Jain, N., & Ching, A. (2009). Bayesian Estimation of Dynamic Discrete Choice Models. ***Econometrica*** 77 (6): 1865-1899.

Jiang, R., Manchanda, P., & Rossi, P. (2009). Bayesian analysis of random coefficient logit models using aggregate data. ***Journal of Econometrics*** 149 (2): 136-148.

Kahneman, D., & Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk. ***Econometrica*** 47(2) 263-291.

Keane, M. P., & Wolpin, K. I. (1994). The solution and estimation of discrete choice dynamic programming models by simulation and interpolation: Monte Carlo evidence. ***Review of Economics and Statistics*** *76*(4), 648–672.

Kőszegi, B., & Rabin, M. (2006). A Model of Reference-Dependent Preferences. ***The Quarterly Journal of Economics*** 121(4), 1133-1165.

Krizhevsky, A., Sutskever, I., & Hinton, G. (2012). ImageNet Classification with Deep Convolutional Neural Networks. *Neural Information Processing Systems* ***(NIPS)***

Li, H. & Srinivasan, K. (2018). Competitive Dynamics in the Sharing Economy: An Analysis in the Context of Airbnb and Hotels. ***CMU working paper***.

Li, J., Moreno, A. & Zhang, D. (2016). Pros vs Joes: Agent Pricing Behavior in the Sharing Economy. ***Ross School of Business Paper*** No. 1298. Available at SSRN: https://ssrn.com/abstract=2708279 .

Mas, A. (2006). Pay, Reference Points, and Police Performance. ***The Quarterly Journal of Economics*** CXXI (3): 787-821.

Nevo, A. 2000. Mergers with differentiated products: The case of the ready-to-eat cereal industry***. RAND Journal of Economics*** 31(3): 395-421.

Proserpio, D., Xu, W., & Zervas, G. (2017). You Get What You Give: Theory and Evidence of Reciprocity in the Sharing Economy. ***Working Paper***

Rossi, P., Allenby, G., & McCulloch, R. (2005). Bayesian Statistics and Marketing. ***John Wiley & Sons***, Ltd.

Rust, J. (1994). Structural Estimation of Markov Decision Processes. ***in Hanclbook of Econometrics*** Vol. 4, ed. by R. Engle and D. McFadden. Amsterda.

Simonyan, K., A. Zisserman. 2015. Very deep convolutional networks for large-scale image recognition. *In International Conference on Learning Representations* ***(ICLR)***

Shampanier et al. (2007). Zero as a Special Price: The True Value of Free Products. ***Marketing Science*** 26 (6): 742-757.

Sudhir, K., 2001. Competitive pricing behavior in the auto market: A structural analysis. ***Marketing Science*** 20(1): 42-60.

Tversky, A., & Kahneman, D. (1991). Loss Aversion in Riskless Choice: A Reference-Dependent Model. ***The Quarterly Journal of Economics*** 106(4), 1039-1061.

Weintraub G., Benkard, C., & Van Roy B. (2008). Markov perfect industry dynamics with many firms. ***Econometrica*** 76(6):1375–1411.

Weintraub, G., Bendard, C., & Van Roy, B. (2010). Computational Methods for Oblivious Equilibrium. ***Operations Research*** 58(4), pp. 1247–1265.

Zhang, S., Lee, D., Singh, P., & Srinivasan, k. (2018). How Much Is an Image Worth? Airbnb Property Demand Estimation Leveraging Large Scale Image Analytics. ***Working Paper*** Available at SSRN: https://ssrn.com/abstract=2976021.

Zervas, G., Proserpio, D., & Byers, J. (2015). A First Look at Online Reputation on Airbnb, Where Every Stay is Above Average. ***Working Paper***.

Online Appendix

A1. Technical Notes on Deep Learning-based Image Quality Classification

With a training data consists of images and the labels (in our study, the label is image quality), the task here is to build an image quality classifier that predicts, for any given input image. Given the large number of hosts in our sample, we leverage the advances in Convolutional Neural Networks (CNNs, an emerging deep learning framework, see Krizhevsky et al. 2012, Simonyan and Zisserman 2015) to build our classifier. Specifically, we use VGG16 model (Simonyan and Zisserman 2015).

Convolutional Neural Networks and VGG16

A CNNs is a special kind of a deep learning model. As shown in Figure 6, a deep learning model consists of a sequence of layers, with each layer containing multiple neurons. Each layer is basically a multidimensional matrix, with each neuro 'carrying' a weight that represents the numeric value of each element. The number of layers that carry weight define the 'depth' of a deep learning model.

In a deep learning framework, high dimensional data such as images and texts is expressed as multidimensional matrices/arrays. Then the model processes the data through the neuron layers implementing matrix multiplication on the data. What defines a CNNs is a special layer—convolution layer, which operates dot productions on the input data (below we will describe operation of convolution layers). The CNN model processes data through matrix multiplication between the input image and the first layer of neurons. This operation generates an intermediate output (also represented by a multi-dimensional matrix), which can be viewed as 'useful information' extracted from the image and serve as the input for the next layer. Such implementations continue till the last layer of the model, i.e., the output layer that computes the probability distribution over the multiple labels. The probability distribution is then converted to labels.

Figure 6 Description of Architecture and Layer Description of the VGG Model

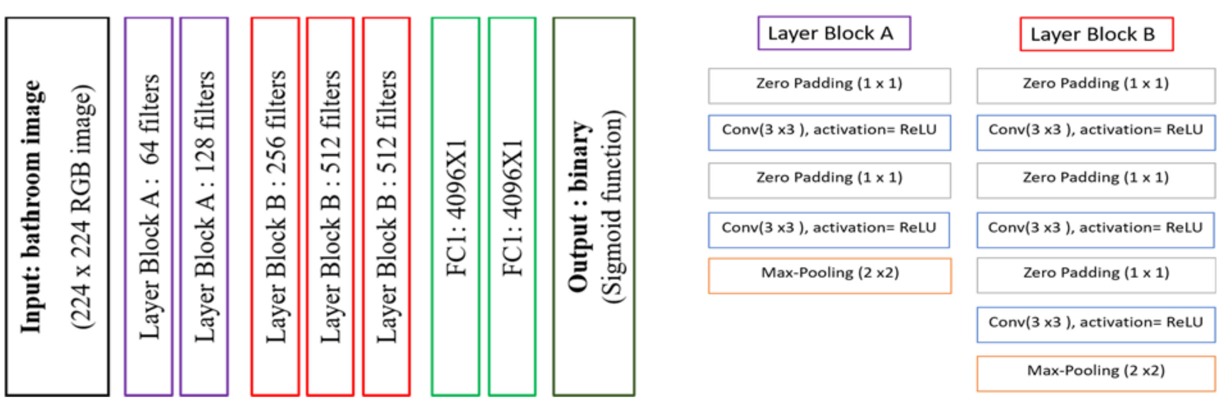

Filters: Indicates the number of convolution windows (i.e., # feature maps) on each convolution layer.

Zero Padding: pad the input with zeros on the edges. To control the spatial size of the output. Zero padding has no impact on the predicted output.

Max-pooling: subsampling method. A 2x2 window slides through (without overlap) each feature map at that layer, and then the maximum value in the window is picked as representation of the window. Reduces computation and provides translation invariance.

Operations of Key Layers

We describe convolution layer and pooling layer, which are the key layers in a CNNs.

Convolution Layer

The convolution layer is the most important and unique layer in the CNN. A convolution layer consists of a stack of so called convolution filter or convolution kernel. A convolution filter is simply a matrix with each element representing a numeric value. For example, in a convolution block, a convolution layer with a size of 3X3 and hence consists of 9 such numeric values[32]. Such a matrix, treating an image or an intermediate input as a matrix, operates a dot production by 'sliding' through the input. Therefore, for an input with relative large size (e.g., 224X224), a 3X3convolution filter operates dot production for every 3X3 patch on that input matrix. The nice features of convolution operation are that: 1) it reduces the dimensionality of parameters, and 2) it well explores and reserves the (local) spatial relationships of the input. Particularly, an intuitive example of the second feature is that: if a convolution kernel extracts a particular oriented edge of an object, then operating this kernel on every small square (e.g., 3X3 and 1X1) on an image would extract all edges in that direct from the image. Many of such kernels that extract edges would extract edges in all directions—potentially constructing the contour of an object. As can been seen in ***Figure 6***, each of the blocks consist of varying numbers of convolutional filters (e.g., 64, 128, 256, 512, 1024, and 2048filters). Hence, these kernels extract features from an input data, which represents the extracted features from the preceding layers. Towards the output layer in the CNN, the filters combined extract higher- and higher- level features. That is, the CNN is able to extract a hierarchical structure of features, that are related to predict the output labels.

Pooling Layer

It's a common practice in CNN to insert a pooling layer in-between the successive convolution layers. A pooling layer is a small square filter. In our model, the pooling filter is a 3X3 matrix. Similar to the operation of convolution filter, an average-pooling layer applies to every 3X3 square patch on an input data. The

[32] The size of a convolution layer is a choice of the model architecture. 3X3 is a widely-used choice. Another common choice is 5X5.

function of a pooling layer is to pick and using the average value in that 3X3 square. Adding pooling layers can reduce the spatial size of the intermediate features and the dimension of the trained parameters in the model. Particularly, it helps to efficiently prevent the problem of over-fitting.

Training Technical Notes

To effectively learn image features that have predicative power on image aesthetic quality labels, we leverage transfer learning and build our model on top of an existing deep learning model that was well-trained for a related task (Zhang et al. 2015). Specifically, we adopted the model of VGG16, with the output layer in the original VGG16 removed as it was specific to the original task (object classification). We then add three fully connected layers on top of that (dimension of 1), where the last layer is output layer.

To improve the training process, we initialize the model weights with the pre-trained weights of the original VGG16 and then fine-tune the parameters. For images, the extracted information is generic, to some extent, across various tasks (e.g., early layers in CNNs serve as edge and contour detectors). Hence, we were able to optimize our model starting from a point where it was already close to 'optimum'. Hence, we efficiently improved the learning process of our model, with the initialized weights able to extract relevant features from the images. The added fully connected layers, without pre-trained weights available, were initialized with LeCunn's uniform scaled initiation method (LeCun et al. 1998, LeCun et al. 1998).

To improve the generalization power of the trained model, we employed a real-time data augmentation method, by randomly flipping, rescaling, and rotating the training samples during the training process (Krizhevsky et al. 2012). Specifically, we implement a real-time (i.e., during training) image transformation over each image in the training sample, by randomly 1) flipping input image horizontally, 2) rescaling input image within a scale of 1.2, 3) rotating the image within $20°$. This method introduces random variation in the training sample, increasing the training set size and reducing the overfitting.

A2. Technical Notes on Estimation Strategy

Estimating Demand Model

Jiang et al. (2009) proposed a Bayesian approach of estimating an aggregated market share (BLP) model. The model is estimated using MCMC (Monte Carlo Markov Chain) algorithm.

As described above, a key A key step is to specify the distributional assumptions on endogenous price P and demand shock η, where endogenous variable P and exogenous variables Z satisfies:

$$\begin{cases} P_{jt} = Z_{jt}\delta + \varepsilon_{jt} \\ \mu_{jt} = X_{jt}\bar{\theta} + \eta_{jt} \\ \begin{pmatrix} \xi_{jt} \\ \eta_{jt} \end{pmatrix} N\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \Omega\right) \end{cases}$$

Using Change-of-Variable Theorem, we derive the joint distribution of market share s_t and price P_t :

$$\pi(P_t, s_t \middle| \bar{\theta}, \Sigma, \Omega, \delta) = \pi(\xi_t, \eta_t \middle| \bar{\theta}, \Sigma, \Omega, \delta) J_{(\xi_t, \eta_t \rightarrow P_t, s_t)} = \pi(\xi_t, \eta_t \middle| \bar{\theta}, \Sigma, \Omega, \delta)(J_{(P_t, s_t \rightarrow \xi_t, \eta_t)})^{-1}$$

where $J_{(P_t, s_t \rightarrow \xi_t, \eta_t)} = \begin{Vmatrix} \nabla_{\xi_t} P_t & \nabla_{\eta_t} P_t \\ \nabla_{\xi_t} s_t & \nabla_{\eta_t} s_t \end{Vmatrix}$ is the Jacobian matrix $J(s_t \rightarrow \eta_t) = \begin{Vmatrix} \boldsymbol{I} & \boldsymbol{0} \\ \nabla_{\xi_t} s_t & \nabla_{\eta_t} s_t \end{Vmatrix} = \lVert \nabla_{\eta_t} s_t \rVert$.

Next, we write the likelihood in Equation (18):

$$\begin{aligned} L(\bar{\theta}, \Sigma, \delta, \Omega) &= \prod_{t=1}^{T} \pi(P_t, ms_t \middle| \bar{\theta}, \Sigma, \Omega, \delta) \\ &= \prod_{t=1}^{T} \{ \left(\begin{Vmatrix} \nabla_{\xi_t} P_t & \nabla_{\eta_t} P_t \\ \nabla_{\xi_t} s_t & \nabla_{\eta_t} s_t \end{Vmatrix} \right)^{-1} \times \phi(\begin{pmatrix} \xi_{jt} = P_{jt} - Z_{jt}\delta \\ \eta_{jt} = \mu_{jt} - X_{jt}\bar{\theta} \end{pmatrix} | \begin{pmatrix} 0 \\ 0 \end{pmatrix}, \Omega) \} \end{aligned} \tag{18}$$

where in the normal distribution $\phi(\cdot)$, η_{jt} relies on mean utility μ_{jt}, which is numerically computed by the contraction mapping method proposed by Berry et al. (1995).

Furthermore, to ensure that the estimated covariance variance-matrix Σ is positive-definite, following the re-parameterization method used in Jiang et al. (2009), we use Cholesky decomposition and write:

$$\Sigma = U'U; \; U = \begin{bmatrix} e^{r_{11}} & r_{12} & \cdots & r_{1K} \\ 0 & e^{r_{22}} & \ddots & \vdots \\ \vdots & \ddots & \ddots & r_{K-1,K} \\ 0 & \cdots & 0 & e^{r_{KK}} \end{bmatrix}$$

Hence, instead draw a whole variance-covariance matrix directly in each MCMC iteration, we draw parameters $r = \{r_{lk}\}_{l,k=1\ldots K, l\leq k}$. We rewrite Equation (18)

$$L(\bar{\theta}, r, \delta, \Omega) = \pi(\bar{\theta}, r, \Omega, \delta) \prod_{t=1}^{T} \{ \left(\begin{Vmatrix} \nabla_{\xi_t} P_t & \nabla_{\eta_t} P_t \\ \nabla_{\xi_t} s_t & \nabla_{\eta_t} s_t \end{Vmatrix} \right)^{-1} \times \phi(\begin{pmatrix} \xi_{jt} = P_{jt} - Z_{jt}\delta \\ \eta_{jt} = \mu_{jt} - X_{jt}\bar{\theta} \end{pmatrix} | \begin{pmatrix} 0 \\ 0 \end{pmatrix}, \Omega) \}$$

Lastly, given the priors on the parameters and likelihood function, the joint posterior distribution of the parameters is[33]:

$$\pi(\bar{\theta},\Sigma,\Omega,\delta\big|\{P_t,s_t,X_t\}_{t=1}^{T}) \propto L(\theta,r,\delta,\Omega)\times\ \pi(\bar{\theta},r,\Omega,\delta) \tag{19}$$

$$= \prod_{t=1}^{T}\{\left(\left\|\begin{matrix}\nabla_{\xi_t}P_t & \nabla_{\eta_t}P_t\\ \nabla_{\xi_t}s_t & \nabla_{\eta_t}s_t\end{matrix}\right\|\right)^{-1}\times\phi(\begin{pmatrix}\xi_{jt}=P_{jt}-Z_{jt}\delta\\ \eta_{jt}=\mu_{jt}-X_{jt}\bar{\theta}\end{pmatrix}|\begin{pmatrix}0\\0\end{pmatrix},\Omega)\}$$

$$\times|V_{\bar{\theta}}|^{-\frac{1}{2}}exp\{-\frac{1}{2}(\bar{\theta}-\theta_0)'V_{\bar{\theta}}^{-1}(\bar{\theta}-\theta_0)\}\times\prod_{l=1}^{K}exp\{-\frac{(r_{ll})^2}{2\sigma_{r_{ll}}^2}\}$$

$$\times\prod_{l=1}^{K-1}\prod_{k=l+1}^{K}exp\{-\frac{(r_{lk})^2}{2\sigma_{r_{off}}^2}\}$$

where $\pi(\bar{\theta},r,\Omega,\delta)$ is specified priors on the parameters. Specifically, for variance-covariance matrix, we specify the priors on $r=\{r_{lk}\}_{l,k=1\ldots K,l\le k}$ with $r_{ll}\sim N(0,\sigma_{r_{ll}}^2), r_{lk}\sim N(0,\sigma_{r_{off}}^2)$ for the diagonal, and off-diagonal elements in matrix U, respectively. For the population mean for characteristics coefficients $\bar{\theta}$,as written in Equation (16), we specify a multivariate normal distribution prior: $\bar{\theta}\sim MVN(\bar{\theta}_0,V_{\bar{\theta}})$. Furthermore, for the instrumental variables related parameters δ and Ω, we specify the following priors:

$$\delta\sim MVN(\bar{\delta},V_{\delta})$$
$$\Omega\sim IW(\upsilon_0,V_{\Omega})$$

where IW indicates an inverse Wishart distribution.

MCMC Estimation Steps

The MCMC estimation steps follows a strategy of Gibbs sampling combined Metropolis steps (Jiang et al. 2009, Rossi et al. 2005). Briefly speaking, in each iteration of the MCMC, we first use Gibbs Sampler to draw the conditionals of $\bar{\theta},\delta,\Omega|r,\{ms_t,P_t,W_t,Z_t\}_{t=1}^{T},\bar{\theta}_0,V_{\bar{\theta}},\bar{\delta},V_{\delta},\upsilon_0,V_{\Omega}$ in a sequence. Then, conditional on updated $\{\bar{\theta},\delta,\Omega\}$, data $\{ms_t,P_t,W_t,Z_t\}_{t=1}^{T}$, and priors $(\sigma_{r_{ll}}^2,\sigma_{r_{off}}^2)$,we update the variance-covariance matrix, Σ, by making draws of r through a Random-Walk (RW) Metropolis chain. Specifically, we draw a proposal of r, given the accepted r in the previous iteration: $r_{new}=r_{old}+MVN(\mathbf{0},\sigma^2D_r)$, where σ^2 is one of $(\sigma_{r_{ll}}^2,\sigma_{r_{off}}^2)$ depending on whether we're drawing a diagonal or off-diagonal element of Σ. D_r is a candidate covariance matrix. r_{new} is either accepted or rejected, based on ratio computed using

[33] For the setup of hyper-parameters, we used diffuse priors. In appendix, we describe details on the choices of priors.

Equation (16). The intuition is that, if conditional on data, priors, and other parameters updated in the Gibbs sampling step, $\Sigma_{new}(r_{new})$, relative to $\Sigma_{old}(r_{old})$, is closer the true posterior of Σ, then we should have $(\bar{\theta}, \Sigma_{new}, \Omega, \delta \big| \{P_t, ms_t, X_t\}_{t=1}^{T}) > (\bar{\theta}, \Sigma_{old}, \Omega, \delta \big| \{P_t, ms_t, X_t\}_{t=1}^{T})$. In appendix we provide detailed technical notes of our estimation steps.

Estimating Supply Model

Conditional on one's current state $s = (s_{jt}, s_{-jt})$, her mage and effort decisions can be described as sequentially solving a DP problem:

$$\{a_{jkt}\}_{t=0}^{\infty} = \underset{\{a_{jkt}\}_{t=0}^{\infty}}{argmax}\, E_{\varepsilon_{jkt}}\{\sum_{t=0}^{\infty} \cdot (\widetilde{U}_{jkt}(a_{jkt} \big| s_{jt}, s_{-jt}) + \varepsilon_{jkt})\} \tag{20}$$

where $\widetilde{U}_{jkt}$ is property j's expected utility from choosing action k in period t and ε_{jkt} is the random shock associated to action *k* that is received before *j* makes a decision.

As discussed in section 3.8., in such dynamic game with many players, computing an MPE is computationally infeasible, hence we use OE to approximate MPE. In a OE, the individual's conditional choice probability is a function of her own state s_{jt} only. The set of states of her peers, s_{-jt}, is captured by tracking an average industry sate $\bar{s}_t$, which reflects the distribution of the number of the reviews across the properties. It can be seen one's action and utility is influenced by her peers' state—s_{-jt}, as it is the action of *j* and her peers and the subsequent state transitions that determine the average state in the next period. Then solving for an OE provides substantial computational advantage, as it converts a many-agent game problem into a problem similar to single-agent optimization, treating $\bar{s}_t$ as a single state variable that is common across all individuals at time *t*. Thus, one can use any existing estimation method that can be applied to a single-agent discrete-choice dynamic programming (DDP) model to solve for an OE. Widely-used estimation strategy includes the nested fixed-point (NFXP) algorithm (Rust 1987) and conditional choice probability (CCP) based estimation (Hotz and Miller 1993, Aguirregabiria and Mira 2007).

In this paper, we use a Bayesian estimation strategy as this way we can flexibly incorporate individual heterogeneity—a key element in our model—in a hierarchical Bayesian framework (developed by Imai, Jain and Ching (2009), hereafter IJC). IJC algorithm allows estimating a heterogeneous model with a relatively low computational burden. In addition, it overcomes the problem of "curse of dimensionality[34]

[34] The state space grows exponentially with the dimensionality of sate variables, causing evaluating Bellman operator at every point in the state space infeasible.

when approximating the DP solution and avoids the complexity of searching for a global optimum in the space of the data likelihood function (IJC provides DP approximation that is comparable to state-of-the-art likelihood-based approaches, e.g., Keane and Wolpin (1994), Ackerberg (2009). See Ching et al. (2012) for detailed discussions). The advantage of avoiding of searching in the parameter space, which usually requires the use of an optimization tool, is another reason we choose IJC algorithm. As we will discuss in section 4.4., Bayesian estimation approach can be easily combined with parallel computing and GPU computing techniques, without which it would be computationally infeasible given the large number of individuals and state space in our study.

IJC Algorithm

We briefly introduce the logics and estimation procedure in IJC algorithm. In appendix we provide technical notes and details of implementing IJC.

IJC algorithm combines MCMC with DDP approximation, solving for the DP problem and making draws of structural parameters from the posterior distribution simultaneously. At each iteration m in the MCMC, IJC saves the simulated parameter vector θ_{IJC}^{*m} and computes a corresponding pseudo-value function $\widetilde{W}^{m}\left(\theta_{IJC}^{*m}\right)$[35]. A total of the most recent N iterations of $\{\theta_{IJC}^{*m}, \widetilde{W}^{m}\left(\theta_{IJC}^{*m}\right)\}$ are saved. When at new iteration m', the simulated vector $\theta_{IJC}^{*m'}$ is rejected or accepted by comparing the pseudo- posterior likelihood evaluated at the accepted parameters from the previous iteration, $\theta_{IJC}^{m'-1}$, and at the proposed parameters at current iteration, $\theta_{IJC}^{*m'}$. When computing the pseudo-likelihood function, one needs to calculate the choice probability for each choice alternative. Recall that one solves for the DP problem by taking into account the value function (see Equation 17)), hence the likelihood function is also 'pseudo-' because the conditional choice probabilities are computed based on pseudo-value functions $\{\widetilde{W}^{n}\left(\theta_{IJC}^{*n}\right)\}_{n=m'-1}^{n=m'-N}\}$ saved in the past N iterations. Specifically, $\widetilde{W}^{m'}\left(\theta_{IJC}^{*m'}\right)$ is approximated by computing a (kernel-based) weighted average of the past N history draws of $\{\widetilde{W}^{n}\left(\theta_{IJC}^{*n}\right), \theta_{IJC}^{*n}\}_{n=m'-1}^{n=m'-N}\}$, with the more weights attributed to history that have θ_{IJC}^{*n} closer to current draw $\theta_{IJC}^{*m'}$. Hence, IJC algorithm is efficient in providing a full solution to the DP problem as it keeps the simulated parameter draws and computed pseudo-value functions to approximate the current pseudo-value function, with Bellman operator evaluated exactly once at each

[35] The pseudo-value function is obtained by applying the Bellman operator (i.e., solving for the value function) at the trial parameter vector. It is called 'pseudo' as the functions are evaluated at the simulated parameter vector not at the true parameter vector. Here the * denotes that this is proposed parameter (regardless of whether it was accepted or rejected) at that iteration.

interaction. As Imai et al. (2009) proved, such an interactive steps of simulating parameter vector through the pseud0-Markov chain can effectively approach a steady state (after burn-in), where most of the structural parameters will be drawn from a distribution close to the true posterior distribution of the parameter vector.

In summary, we a Gibbs Sample to sequentially simulate parameters of $\left(\left\{\lambda_j^{effort*n}, \lambda_j^{operate*n}\right\}_{j=1}^{J}, \lambda, \Sigma_\lambda, \lambda^{MedImg}, \lambda^{HighImg}, \sigma_\varepsilon\right)$, with $\left(\lambda_j^{effort}, \lambda_j^{operate}\right) \sim MVN(\lambda, \Sigma_\lambda)$. At each iteration m, we have the history of the drawn parameters and the associated pseudo-value functions:

$$\left\{ \begin{array}{l} \widetilde{W}^n\left(\cdot; \left\{\lambda_j^{effort*m}, \lambda_j^{operate*m}\right\}_{j=1}^{J}, \lambda^{MedImg*m}, \lambda^{HighImg*m}, \sigma_\varepsilon^{*m}\right)\right), \\ \left\{\lambda_j^{effort*m}, \lambda_j^{operate*m}\right\}_{j=1}^{J}, \lambda^{MedImg*n}, \lambda^{HighImg*n}, \sigma_\varepsilon^{*n}, \bar{s}^{*n}\} \end{array} \right\}_{n=m-N}^{n=m-1}$$

where $\widetilde{W}^n(\ \cdot;)$ indicates the pseudo-value functions at all possible state and $\bar{s}^{*n}$ is proposed industry average state. Then each iteration *m* consists the following steps (we pre-specify of M as the total number of iterations):

1) given $\left\{\lambda_j^{effort^{m-1}}, \lambda_j^{operate^{m-1}}\right\}_{j=1}^{J}$, generate the conditional posterior mean $\bar{\lambda}^m$ from a multivariate normal distribution and the posterior variance-covariance matrix Σ_λ^m from an inverse-gamma distribution
2) let $\rho_j = (\lambda_j^{effort}, \lambda_j^{operate})$ indicate the individual-specific parameters, we want to make a draw from its posterior distribution

$$f\left(\rho_j \middle| \rho, \Sigma_\rho, \lambda^{MedImg}, \lambda^{HighImg}, \sigma_\varepsilon, b_j, data_j\right) \sim g\left(\rho_j \middle| \rho, \Sigma_\rho\right) \cdot L_j(b_j | \rho_j, \lambda^{MedImg}, \lambda^{HighImg}, \sigma_\varepsilon; data_j)$$

where $g\left(\rho_j \middle| \rho, \Sigma_\rho\right)$ indicates is the probability density function of ρ_j given the population mean ρ and variance-covariance matrix Σ_ρ. $L_j(b_i | \cdot; data_j)$ is the likelihood for individual *j* with the tracked action and data (covariates) across time, i.e., $b_j = \left\{a_{jkt}\right\}_{t=1}^{T}$, $data_j = \left\{X_{jt}, s_{jt}, s_{-jt}\right\}_{t=1}^{T}$, evaluated at the trial parameter vector. Since this posterior distribution $f\left(\rho_j \middle| \cdot\right)$ does not a closed-form from which we can easily make a draw, we use Metropolis-Hasting to draw a proposal, $\rho_j^{*m} = N(\rho_j^{m-1}, \Upsilon)$, which will be evaluated to be either accepted or rejected, determined by a computed acceptance ratio:

$$accept_j^m = min\{1, \frac{g(\rho_j^{*m}|\rho^m, \Sigma_\rho^m) \cdot \tilde{L}_j^m \left(b_j \middle| \rho_j^{*m}, \lambda^{MedImg^{m-1}}, \lambda^{HighImg^{m-1}}, \sigma_\varepsilon^{m-1}; data_j\right) \cdot N(\rho_j^{m-1}|\rho_j^{*r}, \Upsilon)}{g(\rho_j^{m-1}|\rho^m, \Sigma_\rho^m) \cdot \tilde{L}_j^m \left(b_j \middle| \rho_j^{m-1}, \lambda^{MedImg^{m-1}}, \lambda^{HighImg^{m-1}}, \sigma_\varepsilon^{m-1}; data_j\right) \cdot N(\rho_j^{*m}|\rho_j^{m-1}, \Upsilon)}\}$$

Let $\gamma_c = (\lambda^{MedImg}, \lambda^{HighImg}, \sigma_\varepsilon)$ indicate the common cost vector, we evaluate the pseudo-likelihood $\tilde{L}_j^m$ at the corresponding parameters and approximated pseud-value functions

$$\tilde{E}_j^m[W(s'; \rho_j^{*m}, \theta_c^{m-1}))|s, a] = \sum_{n=m-N}^{m-1} \widetilde{W}^n(s'; \rho_j^{*n}, \theta_c^{*n}) * \frac{K_h(\rho_j^{*n}, \rho_j^{*m}) K_h(\gamma_c^{*n}, \gamma_c^{m-1}) K_h(\bar{s}^{*n}, \bar{s}')}{\sum_{l=m-N}^{m-1} K_h(\rho_j^{*l}, \rho_j^{*m}) K_h(\gamma_c^{*l}, \gamma_c^{m-1}) K_h(\bar{s}^{*l}, \bar{s}')} \tag{21}$$

where $K_h(\cdot)$ represents a multivariate Gaussian Kernel with bandwidth h. s' and $\bar{s}'$ indicates the individual's state and the industry average state in the next period, conditional on action a. Note that the likelihood function obtained by computing the choice probability. In a DDP with Type-I EV distribution of idiosyncratic shocks, the probability of choosing alternative *k* is:

$$prob_{jk} = \frac{\exp\left(\widetilde{U}_J(s, k) + \beta \widetilde{W}^n(s'; \rho_j^{*n}, \gamma_c^{*n}|s, k)\right)}{\sum_{k'=0}^{J} \exp\left(\widetilde{U}_J(s, k') + \beta \widetilde{W}^n(s'; \rho_j^{*n}, \gamma_c^{*n}|s, k')\right)} \tag{22}$$

3) conditional on the accepted ρ_j, we draw the common cost parameters $\gamma_c = (\lambda^{MedImg}, \lambda^{HighImg}, \sigma_\varepsilon)$: $\gamma_c^{*m} \sim N(\gamma_c^{m-1}, \Upsilon)$. Here similar to step 2), we use Metropolis-Hasting method to evaluate the drawn parameters by approximating the averaged pseudo-value functions, computing likelihood functions evaluated at corresponding parameters, and then computing the acceptance ratio.
4) with the simulated parameters at *m*th iteration, we now update and store the pseudo-value function:

$$\widetilde{W}^m\left(. ; \left\{\lambda_j^{effort^{*m}}, \lambda_j^{operate^{*m}}\right\}_{j=1}^{J}, \lambda^{MedImg^{*m}}, \lambda^{HighImg^{*m}}, \sigma_\varepsilon^{*m})\right) = \{ln\left[\sum_{k'=0}^{J} exp\left(\widetilde{U}_J(s, k') + \beta \widetilde{W}^n\left(s'; \lambda_j^{effort^{*m}}, \lambda_j^{operate^{*m}}, \lambda^{MedImg^{*m}}, \lambda^{HighImg^{*m}}, \sigma_\varepsilon^{*m}|s, k'\right)\right)\right]\Bigg\}_{j=1}^{J}$$

5) m=m+1 (iteration continues with step 1)-4) repeated, till m>M).