

How Much Is An Image Worth? An Empirical Analysis of Property's Image Aesthetic Quality on Demand at AirBNB

Completed Research Paper

Shunyuan Zhang

David A. Tepper School of Business
Carnegie Mellon University,
Pittsburgh, Pa
shunyuaz@andrew.cmu.edu

Dokyun Lee

David A. Tepper School of Business
Carnegie Mellon University,
Pittsburgh, Pa
dokyun@cmu.edu

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

Abstract

Consumers using sharing economy platforms such as Airbnb are challenged with high product uncertainty and search cost. To ameliorate these issues, Airbnb has implemented many strategies such as professionally taking high quality photos for hosts and calling them verified. In this paper we study the impact of having unit list's photos verified. To assess the aesthetic quality of images, we use machine learning techniques. Employing Difference-in-Difference analysis, we find that on average, rooms with verified photos are 9% more frequently booked. We further separate the effect of photo verification from photo quality and room reviews and find an extra \$2,455 in yearly earnings brought by high photo quality. Lastly, we look at the properties in the same neighborhood and find asymmetric spillover effects. On the neighborhood level, the results suggest higher overall demand if more rooms have verified photos.

Keywords: sharing economy, Airbnb, image quality, aesthetic quality classification, treatment effect

Introduction

Sharing economy, also known as collaborative consumption, is a market model that provides peer-to-peer sharing of access to goods and services (Sundararajan 2016, Cohen and Sundararajan 2015, Avital et al. 2014, Hamari et al. 2015). The global sharing economy market will generated around \$15 billion in 2013 and is projected to increase the global revenues to roughly \$335 billion by 2025 (PwC report 2015). The rapid growth has led to sharing economy being viewed as disruptive to traditional markets. Some firms under sharing economy model are now valued higher than their counterparts in traditional market. For

example, Uber was valued over \$18 billion in 2014, roughly the sum of valuations of Avis (\$6.3 billion) and Hertz (\$12.4 billion) (Rusli and Macmilian 2014). Another example---Airbnb, the research context of this paper, is a sharing economy platform for people to list, find, and rent lodging. Airbnb was recently reported to be valued at \$24 billion, while the valuation of Marriott is around \$21 billion (Winkler and Macmilan 2015). In 2014, Airbnb was reported to average over 4 million guests per night, 22% more than the number for Hilton Worldwide (PwC report 2015). Airbnb has therefore become one of the most important markets where hosts expect to generate income by renting their properties (rooms, apartments, houses etc.).

Despite the explosive growth, sharing economy platforms such as Airbnb are facing big concerns from consumers (Ufford 2015). In a recent consumer report, more than half of the consumers (who have tried sharing economy) said they have concerns. Even among those who are familiar with sharing economy, more than two thirds agree that they will try only if someone they trust recommend (PwC report 2015). Survey shows that main concerns are about security and uncertain or inconsistent quality, which lead many consumers to choose trusted hotel brands over Airbnb (PwC report 2015). Complaints are also heard from host' side when it comes to compare Airbnb room photos to hotel room photos. Hosts said the room photos by themselves "actually makes the room look smaller than it is", while photos taken by professional photographers (such as hotel room photos) makes "the room looks more". To address these concerns, Airbnb has implemented a couple of strategies such as building trust system (ID verification, guest reviews etc.) and launching so called "photography program" (launched in 2011, which provides access for Airbnb hosts to free photography service). These actions were intended to add transparency and trust, and to reduce friction and concerns when consumers/hosts are doing business with strangers (Ufford 2015). For example, the photography program was launched to solve issues from two aspects: 1. professional photographers shoot high resolution hence high quality photos for a property; 2. verified photos give the guests added trust since an Airbnb's professional photographer has visited the property unseen by the guests.

Although the photography program has been running for 5 years, studies in evaluating its effect is lacking in some aspects, thus in this paper we aim to study the impact of verified photos, which are taken by Airbnb professional photographer. Specifically, we ask the following three main research questions:

- 1) What is the impact of room photo verification on room demand, i.e., room booking frequency?
- 2) How much, if any, of the effect of verified photos comes through image quality and added trust? That is, can we decouple the high photo quality and photo verification, and quantify them?
- 3) Are there any spillover effect, i.e., does rooms with verified photos affect the demand of rooms in the same neighborhood? Moreover, what's the impact of percentage of rooms with verified photos in a neighborhood on the aggregated demand of that neighborhood?

We collected a unique dataset of over 17,000 Airbnb listings (rooms) and over 15,000 hosts. The dataset contains rich information about room's daily availability, room characteristics, room photos, and information about hosts. Employing Difference-in-Difference (DD) analysis, we find that the effect of verified photos is positive and significant on room demand. On average, a room will be 9% more frequently being booked by having verified photos, which will roughly bring the host an increased income of 3,285 USD if room's daily price is set at 100 USD¹. We further separate the effect of photo verification from photo quality and room reviews. The estimation results suggest an increase of \$2,455 in earnings in a calendar year to the host, if the host price his/her room at \$100 and update his/her 15 low quality photos to high quality (but unverified) photos². We then investigate possible spillover effect across properties in the same neighborhood. The results suggest that a higher portion of properties having verified photos in the neighborhood has a positive impact on the demand of rooms with verified photos and a negative impact on the demand of rooms without verified photos. On average, a representative property receive a negative spillover effect, possibly due to the fact that relatively small number of

¹ In our data, the mean and the median of room's daily price are \$162.2 and \$109, respectively.

² 15 is the average number of photos of rooms in the sample.

properties have verified photos. When we look at aggregated demand on neighborhood level, the results suggest that a neighborhood will have a higher overall demand if it has more (in percentage) rooms with verified photos. One explanation is that, more rooms with verified photos bring a better “neighborhood image” for that neighborhood, and that as a result it’s able to “steal” demand from neighborhoods near it. One challenging but essential task in this study is to quantify image quality. To assess the quality of room photos, we adopt method of image aesthetic quality assessment, which is widely used in computer vision. We combine techniques in machine learning and computer vision, and train a classifier to predict image aesthetic quality. Our classifier achieves a relatively high accuracy in predicting image quality (error rate < 0.15 on test set).

Literature Review

Our work is related to the relatively new yet emerging literature on sharing economy. The rise of sharing economy has drawn massive attention (Avital et al. 2014, Sundararajan 2016, Cohen and Sundararajan 2015, Zervas et al. 2016, Martin et al. 2010, Cervero et al. 2007, Cusumano 2015, Bardhi and Eckhardt 2012, Hall and Krueger 2015). A recent study on Airbnb (Zervas et al. 2016) investigates the impact of Airbnb model on hotel’s revenue and finds that lower-priced hotels are most affected by Airbnb. Our paper differs from the majority of the literature in that, instead of looking at the impact of sharing economy model on the traditional market (Cusumano 2015, Zervas et al. 2016), we are interested in the outcome of individual participants in a sharing economy. Specifically, we aim to identify determinates of the demand of a property listed on Airbnb.com. We are particularly interested in identifying and quantifying the (monetarized) effects of photo verification as well as of photo quality.

This paper is also related to the literature studying aesthetic classification on images. In computer vision, image aesthetic quality assessment is usually used to design algorithm that automatically predicts or orders images in terms of the quality (Wang et al. 2003). An application for example, is to automatically select the photo with highest quality from one’s album. A typical task in aesthetic quality assessment is classification, e.g., classify an input image to two categories—‘high quality’ versus ‘low quality’. Early work in aesthetic classification task relies on forming a list of features that are in line some “rules of thumb” in photography, for example, “The Rule of Thirds” (Datta et al. 2006, Ke et al. 2006, Desnoyer and Wettergreen 2010, Luo and Tang 2008, Nishiyama et al. 2011, Aydin et al. 2015). Explicit features are then extracted from images to be fed into a supervised learning model. These features usually capture specific aesthetic principles that a professional photographer may consider to shoot a good photo, for example, contrast, shape convexity, region composition, and depth etc.. Method that assesses the aesthetic quality of images by exploring the relation between image complexity and aesthetic quality and by applying image compression technique is also proposed (Romero et al. 2011). However it is often difficult to search and form an exhaustive set of image descriptors that are able to capture aesthetic quality. Recently, method of extracting generic local descriptors and implicitly characterizing images have been proposed and shown to achieve state-of-the-art results (Marchesotti et al. 2011, Murray et al. 2012, Marchesotti et al. 2015). In this paper, we perform the method of using generic image descriptors, specifically, upon extracting Scale Invariant Feature Transform descriptors (Lowe 1999), Fisher Vector (Jaakkola and Haussler 1999) is calculated to represent an image and to produce an aesthetic quality label with a trained SVM classifier (Vapnik 1982).

Empirical Framework

In this section we discuss our research context, data (data collection, definition and measurement of key variables) and our method.

Research Context


Our research context is Airbnb.com, a sharing economy platform for people to list, find, and rent lodging. As Airbnb describes itself, “Airbnb is a trusted community marketplace for people to list, discover, and book unique accommodations around the world”. By 2016, Airbnb has over 2,000,000 listings in over 34,000 cities and 190 countries. More than 60,000,000 guests have experiences with Airbnb. Airbnb makes a profit from both sides of the market—guests are charged a 6%-12% service fee and hosts are charged a 3% fee for each completed booking. Airbnb has dedicated numerous effort to build a

mechanism that makes the community trusted, including building online review/rating system, verifying identities of hosts/guests, and providing the options to hosts/guests to link their Airbnb accounts to their other social media accounts (e.g., Facebook, LinkedIn). One practice particularly of our interests is so called “photography program” launched by Airbnb in 2011. This program give hosts access to local profession photographers who will take photos of the hosts’ properties, and upload the shoots that meet Airbnb’s standards. A photo that is shot and uploaded by Airbnb’s professional photographer comes up with a “verified” watermark that appears in a textbox below the image. Airbnb contended, when launched this program, that a property with verified photos will get bookings 2.5 times more frequently than a property without verified photo³.


Photo shoots by Airbnb professional photographers are likely to benefit hosts in two ways: 1, the verification of phots may help hosts to build trust, because verified photos means this property was reviewed by professional photographers and the photos are real; 2, professional photographers in general are able to make photos of high aesthetic quality, hence a guest, when reviewing the property’s webpage, is likely to have a higher perceived quality of the property. Moreover, professional photos may convey quality better due to less friction in information transfer⁴. To illustrate the differences between verified and unverified photos, in Figure 1 we compare an unverified photo (left) to a verified photo (right).

Figure 1 Compare Unverified To Verified Photos

Unverified Photo



Verified Photo



Data

The source of our data set is Airbnb.com. The dataset spans four months and includes near 18,000 listings (properties) in 7 cities in the U.S., as well as over 15,000 hosts of the listings. Our data contains the following three parts:

Information of hosts

The data of a host includes all the listings of this host, whether this host has a verified Airbnb account, when did the host become a member on Airbnb.com, etc.

Static information of listings

The static data of a listing contains the listing’s characteristics that are unlikely to change, such as location, size, types (e.g., house vs apartment), and amenities.

³ <http://thenextweb.com/apps/2011/10/06/airbnb-launches-its-photography-program-with-13000-verified-properties/#gref>

⁴ An Airbnb host said that verified phots more accurately reflect the look of the room. Quote: “but I can tell you this shot MUCH more accurately reflects the actual color on the walls and the amount of light at this time of day/year in the room. It also more accurately show the size of the desk, and of course that there’s a chair and reading lamp in the room”,... “and in fact in MOST ways I think it’s MORE accurate than the photo I took myself.”

Dynamic information of listings

The dynamic data of listing contains the listing’s information that are likely to change over time, we consider three sources of the dynamic information:

Room availability For each listing, we collect the room’s availability on a daily base. From each day’s collected availability information, we know whether a room is occupied or is available to be booked from the current date to three months in advance.

Room reviews For each listing, we collect all reviews posted on the listing’s webpage every two weeks. Review data include comments posted by the listing’s guests, replies (if any) from the host responding to comments, and overall ratings by guests.

Room photos For each listing, we collect all photos on the listing’s webpage every two weeks. Hence we know at the date of each collecting, whether the room has verified photos.

We look at two periods in our observed window. Period 1 (pre-treatment period) spans from 2016/01/16 to 2016/01/31 and period 2 (post-treatment period) spans from 2016/04/02 to 2016/04/16. We leave out a treatment window for two months and has two inconsecutive periods mainly because the rooms receive treatment (getting verified photos) on different dates (similar construction of pre-treatment and post-treatment periods in Yang 2008, Tella and Schargrotsky 2004, Ashenfelter and Card 1985). We collected data from 17,826 rooms and 15,535 hosts. 4115 out of the 17,826 already had verified photos by pre-treatment period, hence we exclude them from data for conceptual clarity⁵. In remaining 12,776 rooms, we observed that 410 rooms had verified photos in post-treatment period, and the rest remained the status of having no verified photos. Further removing rooms with missing data left us a sample of 12634 rooms, with 410 rooms, those had verified photos in post-treatment period, in treated group and the rest 12224, those had zero verified photo in post-treatment period, in control group. We report summary statistics for variables in Table 1.

Table 1 Summary Statistics

Variable		Treated Group (410 observations) ⁶		Control Group (12,224 observations)	
		Mean	Std. Dev.	Mean	Std. Dev.
Pre-treatment	Num. Room Photos	21.85	13.12	13.47	10.87
	Num. Room Reviews	18.5	41.07	9.83	28.57
	Overall Room Rating Score	32.6	44.73	24.64	41.26
	Demand	0.34	0.39	0.38	0.43
Post-treatment	Num. Room Photos	23.87	13.58	13.58	10.9
	Num. Room Reviews	33.14	41.71	16.92	31.88
	Overall Room Rating Score	90.02	19.59	58.55	45.42
	Demand	0.28	0.34	0.24	0.36
	Room Daily Price	145.7	151.3	163.2	248.2
	Num. Max. Accommodates	3.31	2.18	3.22	2.19

⁵ In our data many units didn’t have verified photos, we think possible reasons are: 1. hosts are on a wait-list waiting for available photographer due to limited local Airbnb photographer; 2. hosts share a property with others hence do not fully control whether they can have a photographer visit and take photos; 3. hosts have access to people who can take nice-looking photos hence are less motivated to join the photography program; 3. personal characteristics.

⁶ The number of treated units is only 410 because for conceptual clarity we’ve removed 4115 units that had got verified photos before pre-treatment period from dataset set.

Time-Invariant Variables	Num. Bathrooms	2.25	0.64	2.41	1.13
	Num. Bedrooms	2.64	1.40	2.58	1.34
	Years of Host Membership	3.34	1.46	3.09	1.50
	Num. Listings of Host	2.26	2.65	5.93	32.76

Extraction of Aesthetic Quality of Images

Before moving empirical framework, let’s discuss how we quantify quality of photos, a challenging but essential task in this study. Since the total number images to be classified in our sample is over 380,000, we need an algorithm that is able to automatically assess the quality of room photos. Largely due to the scale of the data set⁷, we combine techniques in machine learning and computer vision, and train a classifier to predict image aesthetic quality. We use supervised learning to train a binary classifier on a human-labeled training set, which consists of 322 high quality images and 322 low quality images. This binary classifier assigns “1” to high quality image and “0” to low quality image.

Aesthetic Quality Classification We adopt method of image aesthetic quality assessment, which is widely used in computer vision. In semantic task such as image classification and object retrieval, initially a standard approach is Bag-Of-Visual words (BOV), where an image can be represented as a histogram of local features such as Scale Invariant Feature Transform (SIFT, see Lowe 2004) descriptors (Csurka et al. 2004, Sivic and Zisserman 2003). More recent studies in semantic task show state-of-the-art results by employing the extend version of BOV—Fisher Vector (FV) (Perronnin and Dance 2007, Peronnin et al. 2010 b, Marchesotti et al. 2011). In this paper, we employ FV-based aesthetic classification strategy and describe how we train our classifier. We start with brief introduction of FV.

Fisher Vector (FV)

We discuss the difference between BOV and FV, then describe how FV is calculated and then used to represent an image.

Compare BOV to FV In the method of BOV representation, a visual vocabulary is first trained with clustering algorithm (such as K-means clustering) applied on a large number of local features extracted from training images. Then an image is described with a histogram of local features by counting number of local descriptors falling into each visual word (i.e., each cluster). The FV representation extends BOV by encoding higher order statistics (mean, variance) rather than zero-order statistics (simply counting numbers in BOV). A continuous distribution (Gaussian Mixture Model, GMM) instead of discrete distribution (clutters) is trained in FV representation. Then an image is treated as continuous distributions and is represented by FV, which describe its deviation from the trained GMM. This extended version has been shown to outperform BOV method in semantic task (Perronnin and Dance 2007, Peronnin et al. 2010 a, Peronnin et al. 2010 b).

Extract Generic Features A set of generic local features, for example, SIFT descriptors are first extracted from an input image (Lowe 2004). SIFT divides a local patch, e.g., an image into multiple 4*4 grids, computes a Histogram of Oriented Gradients (HOG) of 128 bins in each of the grid, then produces 128-dimensional SIFT descriptors. Each images produces varying number of SIFT descriptors, depending on the size of the image. A large number of SIFT descriptors extracted from images are used to train GMM.

Train GMM The set of SIFT descriptors is used to train a GMM using Expectation-Maximization (EM) algorithm. The GMM parameters to be learned are means and standard deviations of each Gaussian. Number of Gaussians is a hyper-parameter of GMM that needs to be set before training the model. In this paper we used 256 Gaussians since 256 Gaussians in GMM is widely used and have been shown to achieve good performance in image classification (Sánchez et al. 2013, Marchesotti et al. 2011). The trained GMM

⁷ The over 380,000 images contain all images collected during data collection periods, including images of the rooms removed from sample for estimation, and those images updated during treatment period.

can be seen as a probabilistic visual vocabulary that is trained on a set of training images. Moreover, it models the generative process of local features by assigning each local feature with uncertainty (probabilistic) to all visual words in the trained visual vocabulary.

Represent an Image with FV For an input image, we first extract a set of SIFT descriptors $\aleph = \{x_i, i = 1, 2, \dots, I\}$ with each x_i is a SIFT descriptor. Given a trained GMM distribution

$$\emptyset = (\mu_k, \Sigma_k, \pi_k, k = 1, 2, \dots, K)$$

where $k=1,2,\dots,K$ denotes each of the Gaussians, μ_k, Σ_k, π_k denote the mean, covariance (we used diagonal covariance matrix), and weight of the k^{th} Gaussian in the mixture model. $\{\mu_k, \Sigma_k, \pi_k\}_{k=1}^K$ are the parameters learned from previous step.

Sample is characterized by its deviation from the trained GMM distribution with respect to the mean and covariance:

$$u_{jk} = \frac{1}{N\sqrt{2\pi_k}} \sum_{i=1}^I q_{ik} \frac{x_{ji} - \mu_{jk}}{\sigma_{jk}}$$

$$v_{jk} = \frac{1}{N\sqrt{2\pi_k}} \sum_{i=1}^I [q_{ik} \left(\frac{x_{ji} - \mu_{jk}}{\sigma_{jk}} \right)^2 - 1]$$

where $j=1,2,\dots,J$ denotes the vector dimensions (in this paper with SIFT descriptors, $J=128$). q_{ik} is the posterior probability of descriptor x_i being assigned to k^{th} mode in the mixture model:

$$q_{ik} = \frac{\exp\{-\frac{1}{2}(\mathbf{x}_i - \mu_k)^T \Sigma_k^{-1}(\mathbf{x}_i - \mu_k)\}}{\sum_{k'=1}^K \exp\{-\frac{1}{2}(\mathbf{x}_i - \mu_{k'})^T \Sigma_{k'}^{-1}(\mathbf{x}_i - \mu_{k'})\}}$$

Hence for an input image IMG, its FV is a vector representation by stacking vector of deviations with respect to the mean and covariance of each of the K modes in the GMM, that is, we represent the image

$$\Phi(\text{IMG}) = \begin{pmatrix} \vdots \\ \mathbf{u}_k \\ \vdots \\ \mathbf{v}_k \\ \vdots \end{pmatrix}$$

At the end, vectors $\Phi(\text{IMG})$ are further normalized by L^2 norm since L^2 -normalization has been shown to improve the classification accuracy (Perronnin et al. 2010). The normalized FV will be fed into a classifier such as Support Vector Machines (SVMs) to train the classifier.

Train a SVM model

The SVM algorithm was developed from statistical learning theory in 70s (Vapnik 1979) and was introduced in 90s (Boser et al. 1992). SVM model is among the best classification methods for supervised learning and is widely used due to the high accuracy (Cortes and Vapnik 1995, Blanz et al. 1996, Schmidt 1996). We train a non-linear SVM classifier by (implicitly) mapping features into higher-dimension feature space with Radial Basis Function (RBF) Kernel and L^2 regularization is used.

Perform Prediction

With trained GMM and SVM classifier, we can predict the label (“high quality” versus “low quality”) of an input image. SIFT descriptors are extracted from the image to calculate FV, which characterizes the input image. Then FV is fed into the trained SVM classifier to produce the predicted label. In our classification, we assign a binary variable to each input image. The image has an aesthetic quality of 1 if it’s classified as “high quality image” and an aesthetic quality of 0 if otherwise.

Performance of Trained Classifier

The performance of the classifier was assessed on a test set (hold-out sample) for 162 images. We present the performance of our image aesthetic classifier on the test in Table 2. Note that the hold-out sample was not used in training step.

Table 2 Performance of Aesthetic Classifier on Test Set

Property	Performance
Correct Rate	0.8519
Error Rate	0.1481
Sensitivity	0.9136
Specificity	0.7901
Positive Predictive Value	0.8132
Negative Predictive Value	0.9014
Positive Likelihood ⁸	4.3529
Negative Likelihood ⁹	0.1094
Prevalence	0.5000

Table 2 shows that our trained aesthetic classifier performs pretty well on predicting/classifying images based on their aesthetic qualities, with averaged error rate < 15%, that is, in the test set, less than 15% of the samples are incorrectly classified/labeled. The measures of sensitivity and specificity imply that the classifier performs better on recognizing a high quality image than recognizing a low quality image. The measures of positive predictive value and negative predictive value suggest that the classifier is more like to incorrectly label a low quality image as high quality than the other way around. In Figure 2, we compare images with predicted high versus low quality by presenting the output of our classifier given input of four images. Note that the two images at each row are taken for the same room, with one being classified as “high quality” and the other “low quality”. The comparison of two classes of image displayed in Figure 2 gives some intuition about the underlying aesthetic differences between high and low quality images.

Figure 2 Examples of Photos Classified High Quality and Low Quality

Labeled High Quality Image (label Y=1)

Labeled Low Quality Image (label Y=0)

⁸ Positive Likelihood= Sensitivity / (1 – Specificity). This explains why it is greater than 1.

⁹ Negative Likelihood= (1 – Sensitivity) / Specificity



Methods

We're interested in identifying treatment effect, i.e., the impact of having verified photos of a room on its demand. Next we describe the measures of key variables and the treatment effect of estimator.

Key Variables Definitions and Measurement

Room Demand For each room, we check whether a room is occupied on each day within our observed window. For each room, each day, we check its availability from an online dynamic room availability calendar. Then in each period, the room demand is defined as portion of days being occupied during the period. Since we can only observe the room availability not actual room occupancy, the dependent variable (demand) is measured with error¹⁰.

Room Reviews The room ratings takes on two dimensions--the total number of comments and the overall rating of the room. Comments are posted by guests who have stayed in the property and rating is an average score evaluated by guests based on multiple criteria such as cleanliness and location.

¹⁰ For example, the host marks on the room availability calendar that the room is unavailable on a day because he has friends visiting and staying in his property on that day, however, from data this "error observation" is not reflected and we can only infer that the room is booked on that day.

Photo Quality and Quantity For each room, we first predict the aesthetic quality of each of its photos (prediction is implemented with our aesthetic quality classifier described in previous section), then calculate an average quality score over all the photos. Since the quality of each photo is a binary response (0 or 1), the average quality score is between 0 and 1. Photo quantity is simply the number of photos of the property.

Neighborhood Our data contains geographic information about each property including zip code and neighborhood name. In this paper, we present results associated with neighborhood labeled by neighborhood name (for example, all properties with neighborhood name of “Browns Hill” are in the same neighborhood). As a robustness check, we implement the analysis when zip code is used to label neighborhood (that is, properties with the same zip code are in the same neighborhood). The two sets of results are comparable and consistent.

Neighborhood Treatment Level To investigate possible spillover effect, we study the impact of the neighborhood treatment level on a property in the neighborhood. The variable neighborhood treatment level measures the percentage of properties that have verified photos in a given neighborhood and given period. For example, a neighborhood with half of the properties having verified photos has a neighborhood treatment level of 5%.

Treatment Effect Estimator

We adopt the method of combining propensity score and Difference-in-Differences (DD). DD estimator is a popular tool of evaluating the effect of a treatment on an outcome variable that is of researchers’ interests (Card and Krueger 1994, zervas 1990, Garvey and Hanka 1999, Heckman and Payner 1989, Jin and Leslie 2003). DD method estimates treatment effect by comparing two differences: the difference before and after treatment in treated group, and the difference before and after treatment in control group (i.e., untreated group). A simple but essential assumption for DD to provide unbiased estimator is so called “parallel trend”. That is, without treatment intervention, the trend for treated and control groups over time would have been the same. However, the assumption will not hold if, for example, the pre-treatment covariates between treated and control groups are unbalanced, so that the two groups differ in ways that are associated with the dynamics of outcome variable over time (Besley and Case 2000, Angrist and Krueger 1999, Athey and Imbens 2006). Methods have been proposed to account for the possibility that unbalanced observed covariates in treated and control groups create “non-parallel pre-treatment trend”. Two-step framework, where propensity score is applied to adjust for the unbalanced observed covariates between two groups, is one of the strategy to estimate average treatment effect while flexibly adopt observed covariates (Abadie 2005, Stuart et al. 2014). This strategy generally takes two steps: 1) estimate propensity score and 2) based on step 1), estimate Average Treatment Effect (ATE) after adjusting for the unbalanced covariates between treated and control groups.

Propensity Score (PS)

Propensity score is defined as the probability of each individual receiving a treatment, conditional on a set of observed covariates (Rosenbaum and Rubin 1983). Since in practice true propensity scores are often unknown, we need to estimate propensity score with logit or probit regression model that maximizes the data likelihood of observing treatment received in our sample. We use logistic regression

$$\theta = \max_{\theta} \prod_{n=1}^N \left\{ \left(\frac{\exp(\theta X_n)}{1 + \exp(\theta X_n)} \right)^{I(n \text{ in treat group})} \cdot \left(1 - \frac{\exp(\theta X_n)}{1 + \exp(\theta X_n)} \right)^{I(n \text{ in control group})} \right\} \quad \text{Equation 1}$$

where N is the total number of rooms in sample, X_n is a 1*k vector that describes room n’s specific covariates, including either fixed or pre-treatment variables. θ is a k*1 vector parameter to be estimated. $I\{\cdot\}$ is indicator function. The selection of covariates is based on performing covariates balance check. That is, the differences between means in covariates of control rooms and means of treated rooms should be minimized.

Difference-in-Differences Estimator

We implement a linear regression model adjusted for propensity score:

$$\text{demand}_{nt} = \text{intercept} + ATE \cdot \text{treatind}_n + \alpha \cdot \text{treatgroup}_n + \beta \cdot \text{posttreat}_t + \gamma \cdot \text{pscore}_n + \rho \cdot \tilde{X}_{nt} + \varepsilon_{nt}$$

Equation 2

where demand_{nt} is demand measure for room n in period t . treatind_{nt} is 1 if room n has received treatment (having verified photos) in period t and is 0 if otherwise. treatgroup_n is 1 if room n belongs to treated group and is 0 if otherwise. posttreat_t is 1 if period t is post-treatment period and is 0 if otherwise. pscore_n is n 's estimated propensity score calculated in PS step. \tilde{X}_{nt} is a $M \times 1$ covariates vector of room n in period t . The DD estimate of interest is the ATE estimator (\overline{ATE}), is be calculated with OLS (ordinary least square).

Results

We start with the results of propensity score estimation, then we will present the estimated parameters under linear specification model and extended models. We start with the results of propensity score estimation, then we will present the estimated parameters under linear specification model and extended models.


Propensity Score Estimation

In Table 3 we present the estimated coefficients from propensity score step, as in Equation 1. The model suggests that pre-treatment photo quality and quantity, number of bathrooms, number of host's listings, host verification as well as room daily price are strong indicators of predicting the likelihood one property receiving treatment. In Figure 3 we plot the histograms of estimated propensity score for treated and control groups. Plots in Figure 3 suggest that two groups have a large overlap on estimated propensity scores.

Table 3 Estimated Coefficients of Propensity Score Model

	Estimate	Std. Error
Intercept	-5.0578***	0.274874
Num. of Bathroom	-0.3362***	0.101225
Property Type	-0.1072	0.109105
Num. Accommodations	0.0049	0.047377
Num. of Beds	-0.1163	0.079519
Num. of Bedrooms	0.0547	0.054738
Num. of Host's Listings	-0.0872***	0.020073
Years of Member of Host	-0.0066	0.039312
Host ID Verified	0.3600**	0.139922
Quality of Photos	3.6388***	0.194107
Num. of Photos	0.0424***	0.00346
Num. of Comments	0.0020	0.001636
Overall Rating	0.0021	0.001568
Room Daily Price	-0.0007(·)	0.000392
*** p<0.001, ** p<0.01, * p<0.05, (·) p<0.1		

Figure 3 Histogram Plots of Estimated Propensity Scores For Two Groups



DID Estimators

In this section, we present the ATE estimate (\widehat{ATE}) by analyzing an OLS specification as in **Error! Reference source not found.**. We first present results from our linear model specified, without adding covariates \tilde{X}_{nt} . Then we extend this base model by adding covariates of image and review variables. We will then explore the impact of increased portion of properties with verified photos on the demand of individual property in the same neighborhood, and on the aggregated demand of the neighborhood¹¹.

The Effect of Verified Photos

We start with presenting the estimation results of the base model in Table 4. The key variable of interest is $treatind=posttreat*treatgroup$, which equals 1 for a room if and only if it is in treated group and the period is post-treatment period. This key variable captures the effect of having verified photos for a treated property in the post-treatment period.

Table 4 The Effect of Verified Photos

Coefficients	Estimate	Std. Error	t value
<i>Intercept</i>	40.09406***	0.390257	102.7375
<i>treatind (ATE)</i>	8.542135**	2.804269	3.046118
<i>posttreat</i>	-14.9032***	0.505174	-29.5012
<i>treatgroup</i>	-0.3258	2.017353	-0.1615
<i>pscore</i>	-0.53587***	0.052141	-10.2773
*** p<0.001, ** p<0.01, * p<0.05, () p<0.1			

Estimated ATE in Table 4 is positive and significant, suggesting a positive impact of verified photos on room’s demand. This result is in line with Airbnb’s expectation that a room with verified photos will have bookings more frequently than a room without verified photos. The estimation results in Table 4 suggest

¹¹ As a robustness check, we employ two ways of defining a “neighborhood”—by zip code and by neighborhood name. We present results with neighborhood name used. The analysis with zip code show comparable and consistent results.

that, given everything else equal, having verified photos can increase near 9 points of demand than without verified photos. That is, on average, a room will be 9% more frequently being booked by having verified photos, which will roughly bring the host an increased income of $9\% \times 365 \times 100 = 3,285$ USD if room's daily price is set at 100 USD. The coefficient of variable *posttreat* is negative and significant, suggesting time-varying shocks to both groups, for example, higher overall demand in Holidays.

Separating the Effect of Photo and Reviews from Photo Verification

We extend the analysis by adding covariates associated to room photos and room reviews as controls. In Table 5 we present the estimated coefficients in two extended models: extended model A where photos (quality and quantity) are added in covariates as control, and extended model B where both photos (quality and quantity) and reviews (number of comments and overall ratings) are added in covariates as control. This extended analysis allows us to tease out the effect of verification of photos from possible positive effect brought by photo quality and room reviews. In extended model A, we are interested in key variables *treatind*, which captures the effect of verification of photos, and *Photo_Quality*, which captures the effect of having a higher portion of high quality photos. As reported in column A, the treatment effect (\widehat{ATE}), compared to the result in Table 4, is estimated smaller in size as well as with a lower precision. Furthermore, the effect of key variables *Photo_Quality* and *Num.of_Photos* are both positive and significant. Given the estimated effect of *Photo_Quality* variable, we calculate the monetarized impact of having high quality room images. Suppose a room (priced at 100USD) has 15 photos (average number of room photos in our sample), all of which are low quality. If the host were to replace all the 15 low quality photos with 15 high quality though unverified photos, his/her room would get additional $4.22 + 0.167 \times 15 = 6.725$ points of bookings. Consequently, he/she would earn extra $6.725\% \times 365 \times 100 = 2455$ USD from renting his/her property in a calendar year.

Similar findings are presented in Column B, where room reviews are added a second control. In column B, the effects of the four key variables *Photo_Quality*, *Num.of_Photos*, *Num.of_Comments* and *Property_Ratings*, are all positive significant. The estimated effect of key variable *treatind* is estimated in a smaller size, and not significant at 95% confidence level (however significant at 90% level). The results suggest that some of positive impact of photo verification on room demand come through photo quality and room reviews. As we discussed in research context section, the verification of room photos may bring two effects: pleasing photos and trust, which are likely to be captured by aesthetic quality of photos, and by room reviews, respectively. The smaller and less significant impact of verified photos on room demand (when we control photo quality and room reviews) implies that the effect of verified photos can come from: 1) efficient transfer of information due to high quality photo; 2) pleasing effect changing the perceived quality and preference; 3) higher trust in property (host) due to the verification tag.


To compare the effects that come through photo verification, image quality and room reviews, respectively, we contrast their impacts on room demand. We consider the impact of a room getting verified photos on its demand, then calculate the needed amount of increase in room photo quality, and in number of comments, respectively, to create equivalent impact on room demand. In we plot the needed increases, given 3 different cases: for photo quality, we consider cases of a room with number of photos of 10, 15, and 20; for room comments, we consider cases of a room with room ratings of 30, 50, and 70. Results are calculated based on estimates in Column B of Table 5.

Table 5 The Effect of Verified Photos (Photos and Reviews as Control)

	Column A		Column B	
	Extended Model A		Extended Model B	
	Estimate	Std. Error	Estimate	Std. Error
<i>Intercept</i>	36.933***	0.6501	35.158***	0.6489
<i>Treatind</i> (\widehat{ATE})	7.594**	2.8007	4.840 (·)	2.7771
<i>posttreat</i>	-14.927***	0.5041	-19.256***	0.5376
<i>treatgroup</i>	-0.848	2.0149	-1.510	1.9961

<i>pscore</i>	-1.060***	0.0785	-0.997***	0.0789
Photo_Quality	4.225**	1.4333	2.994*	1.4212
<i>Num.of_Photos</i>	0.183***	0.0421	0.080 (.)	0.0421
<i>Photo_Quality * Num. of_Photos</i>	0.167(.)	0.0897	0.247**	0.0891
Num. of_Comments	---	---	0.805***	0.1515
Property_Ratings	---	---	0.136***	0.0065
<i>Num.of_Comments * Property Ratings</i>	---	---	-0.009***	0.0016
*** p<0.001, ** p<0.01, * p<0.05, (.) p<0.1				

Figure 4 Compare Needed Change in Photo Quality and in Room Comments to Create Equivalent Increase in Room Demand by Verified Photos



The Spillover Effect in Neighborhood

To investigate the potential spillover in room demand across properties in the same neighborhood, we add the neighborhood treatment level as an extra control. We are interested in finding the impact of neighborhood treatment level on the demand of any property in the neighborhood. In particular, we ask the following questions, if a neighborhood (assuming here the size of the neighborhood is fixed) has more Airbnb properties with verified photos:

- 1) Will every property in the neighborhood benefit from a likely increased demand brought by a “better neighborhood image”? Or will they suffer from a higher competition with increased number of properties in the same neighborhood, who are actually “stealing” demand from them?
- 2) Are the (positive or negative) spillover effects symmetric to rooms with and without verified photos?
- 3) What is the impact of neighborhood treatment level on the neighborhood-level aggregated demand?

As described previously, key variable, *Neighborhood_Treatment_Level*, is defined as the percentage of properties having verified photos in the same neighborhood. We report the results in Table 6.

In column A, the key variable of interest *Neighborhood_Treatment_Level* is negative and significant, suggesting that an increased percentage of treated properties in neighborhood has, on average, a negative impact on the demand of properties located in the same neighborhood. The result suggests a 0.4 point decrease in demand as every 1 percentage increase in neighborhood treatment level. For example, given everything else equal, the demand for an average property in a neighborhood will decrease 4 points if the percentage of properties having verified photos increase from 10% to 20%. Approximately, the average host earn $4\% \times 365 \times 100 = 1460$ USD less in a calendar year if his/her room is priced at 100USD.

We further look at whether this negative spillover effect is universal across properties, or is it possible that properties with verified photos are able to moderate the spillover effect. In column B, we report the

estimation results with an interaction term *Neighborhood_Treatment_Level*treatind* added as a control. The positive and significant of this interaction variable suggests an asymmetric spillover effect across two groups of properties. The results suggest that properties with verified photos receive a positive spillover effect while other properties receive a negative spillover effect. On average, however, the effect on a representative property is negative, probably due to the relative small portion of properties with verified photos in the sample (our data contains 410 treated properties and 12224 untreated properties).

Table 6 The Spillover Effect of Verified Photos

	Column A		Column B	
	Spillover effect on individual property		Spillover effects on two groups	
	Estimate	Std. Error	Estimate	Std. Error
<i>Intercept</i>	35.156***	0.6487	35.147***	0.6487
<i>Treatind (ATE)</i>	6.074*	2.7973	0.449	3.7297
<i>posttreat</i>	-17.994***	0.6402	-	17.764***
<i>treatgroup</i>	-1.510	1.9956	-1.520	1.9955
<i>pscore</i>	-0.995***	0.0789	-0.994***	0.0789
<i>Neighborhood_Treatment_Level</i>	-0.417***	0.1150	-	0.493***
<i>Neighborhood_Treatment_Level*treatind</i>	---	---	0.975*	0.4278
<i>Photo_Quality</i>	2.996*	1.4209	3.021*	1.4208
<i>Num. of_Photos</i>	0.080 (.)	0.0421	0.080 (.)	0.0421
<i>Num. of_Comments</i>	0.799***	0.1515	0.805***	0.1515
<i>Property_Ratings</i>	0.136***	0.0065	0.136***	0.0065
<i>Photo_Quality * Num. of_Photos</i>	0.244**	0.0891	0.243**	0.0891
<i>Num. of_Comments * Property_Ratings</i>	-0.009***	0.0016	-	0.009***
*** p<0.001, ** p<0.01, * p<0.05, (.) p<0.1				

Next, we investigate whether neighborhood treatment level affect the aggregated demand at neighborhood level. To show the impact at neighborhood level, we treat each neighborhood as an individual, hence one neighborhood has two observations, each in one period. We first calculate the change in aggregated demand (at neighborhood level) across two periods, then aggregate demand is defined as the average demand over all properties in the same neighborhood in a given period. We then look at the change in neighborhood treatment level, and divide neighborhoods into two groups based on their increased percentage of verified rooms across two periods. For example, neighborhoods with none or few new rooms having verified photos are “control neighborhoods”, while those with some new rooms having verified photos in post-treatment period are “treated neighborhoods”¹². We compare the means of change in aggregated demand in two groups of neighborhood. Note that the negative change in aggregated demands are due to seasonal factors. We then test the following null hypothesis:

¹² We did this analysis with two “cut-off” points of change in neighborhood treatment level: 0%, and 1%. Both give consistent results. Results presented in Table 7 are obtained under 1% criteria.

H_0 : Patterns of change in demand over two neighborhood groups are **indifferent**. ($\Delta\mu_{control} = \Delta\mu_{treat}$)

H_1 : Patterns of change in demand over two neighborhood groups are **different**. ($\Delta\mu_{control} \neq \Delta\mu_{treat}$)

The purpose of testing the null hypothesis is that, if the change in aggregated demands are significantly different across the two groups of neighborhoods, then we may interpret this result as the difference caused by the impact of increasing rooms with verified photos in a neighborhood. In particular, if $\Delta\mu_{treat} > \Delta\mu_{control}$, this may suggest that increased neighborhood treatment level has, on average, a positive impact of the overall (aggregated) demand at neighborhood level.

In Table 7, we report the means of change in aggregated demand in two groups of neighborhoods, difference in means, t statistics (with p-value in parentheses), and confidence level to reject the null hypothesis. As a robustness check, we construct the two groups of neighborhoods by restricting the analysis to neighborhood with not too few rooms. For example, in the first column, if neighborhood size ≥ 20 , then we only look at those neighborhoods with at least 20 rooms. We present results with the different thresholds of selecting neighborhoods. As results shown in Table 7, compared to the control neighborhood group, the treated neighborhood group has a relative increase in aggregated demand, suggesting a positive effect of neighborhood treatment level. The differences are statistically significant, hence we conclude that, if a neighborhood gets more (in percentage) rooms with verified photos, not only this has a (negative) impact on the demand of rooms without verified photos and a (positive) impact on the demand of rooms with verified photos in the same neighborhood, but it also has a positive impact of the neighborhood’s overall demand at an aggregated level. One explanation for this is that, more rooms with verified photos in a neighborhood leads to a better “neighborhood image” of the neighborhood. As a result, the neighborhood is able to “steal” demand from other neighborhoods near it.

Table 7 Compare Change in Aggregated Demand at Neighborhood Level

Neighborhood Size \geq	Mean of Change in Aggregate Demand		Dif. In Means	t-statistic (p-value)	Reject H_0 at the Confidence Level of
	Control Neighborhood Group ($\Delta\mu_{control}$)	Treated Neighborhood Group ($\Delta\mu_{treat}$)			
20	-0.1731	-0.1397	0.0334	1.9 (0.06)	10%
30	-0.1829	-0.1365	0.0464	2.1 (0.05)	5%
40	-0.2229	-0.1453	0.0776	3.6 (0.003)	1%
50	-0.2137	-0.1510	0.0627	3.3 (0.007)	1%

Sensitivity on Dependent Variable Calculation

In our main analyses, the Dependent Variable (D.V.) is calculated as the percentage of days a property being unavailable for booking, out of a calendar month. This is done by daily accessing the next day’s availability of a property. However the unavailability on a property’s booking calendar could be interpreted in two ways: 1. the property has been occupied, hence is unavailable for booking; 2. the property is marked by the host “unavailable” (for house cleaning etc.), and hence is unavailable for booking. This may induce a bias in estimated treatment effect, if our observations are a mix of the two scenarios, since the latter scenario doesn’t not reflect a property’s demand. Lack of transaction data, we’re unable to distinguish the two scenarios. However, we run a set of analyses, testing how sensitive our result is to possible “false occupancy”. We alter the way of calculating demand by checking property availability with a range of days in advance. That is, we check whether a property is “available” in m days, with m varies from 1 to 30. The logic behind is that, travelers often book a room in advance, and the bookings of room in next week also gives us information about demand. If we believe that a host on purpose mark his/her property as “unavailable” is rare, then varying the length of days in advance to check property availability should not alter our estimation results much. In Table 8 we present the matching estimator of photo verification when property availability is accessed 10, 20, 25, 30 days ahead (i.e., $m =$

10,20,25,30)¹³. The consistent results (in bold) suggest that the estimated treatment effect is not sensitive to the way that demand is calculated. Note that due to our data limitation, we don't distinguish between the case of "for rent" where the host lives in the property and the other case where the host doesn't live in the property. We can only infer from the property type, that if the property is a "shared" property, it's possibly that the host also lives in.

Table 8 Matching Estimators of Photo Verification with Varying Demand Calculations

Variable	Sample	Treated	Controls	Difference (%)	Std. Error	T-stat
<i>Demand (m=10)</i>	\widehat{ATE}	0.340179	0.242659	9.7520	0.0270	3.61
<i>Demand (m=20)</i>	\widehat{ATE}	0.375143	0.282885	9.2257	0.02793	3.30
<i>Demand (m=25)</i>	\widehat{ATE}	0.38929	0.308876	8.042	0.02885	2.79
<i>Demand (m=30)</i>	\widehat{ATE}	0.416978	0.342964	7.4014	0.029538	2.51

Conclusion

Although many sharing economy platforms are experiencing rapid growth, they face big challenges, one of which is the issue of trust. For sharing economy to continue to expand, the platforms need to understand the essential role of trust in the context of sharing economy. Airbnb, as an example of sharing economy companies has made a lot of effort to build trust system, including ID verification, transaction security system, and the launch of "Photography Program", i.e., photo verification. Research studies on evaluating these implementations, especially the photography program is lacking despite they have been running for a couple of years. In this paper, we investigate the effect of photography program, by looking at the impact of room photo verification on room demand. We further decouple the effect of high quality of photos and the effect of photo verification and quantify the two effects on room demand. Our study makes two contributions: on the one hand, with a rich and unique dataset consisting of over 17,000 Airbnb listings (rooms) and over 15,000 hosts, we employ Difference-in-Difference (DD) analysis and estimate the treatment (verified photos) effect; on the other hand, we bring methods in the literature on image aesthetic quality assessment to quantify the quality of images, where we use machine learning techniques to train a binary classifier that predicts quality of each photos in our sample (over 380,000 photos in total to predict) and classify each photo into two groups—high quality versus low quality. With the ability of assessing photo quality, the effects of image quality can be separated from effects of photo verification. Our estimation results suggest that, for a room priced at \$100/day, having verified photos will bring extra 9% of room booking frequency, leading an extra calendar year income of \$3,285 to the host. By decoupling the effects of photo quality and photo verification, we find that part, but not all, of the effect of verified photos come through high quality of photos taken by professional photographers. Results suggest an increase of \$2,455 in calendar year income to the host, if he/she replaced his/her 15 low quality (and unverified) photos with all high quality (and unverified) photos. We further investigate possible spillover effect across rooms in the same neighborhood and find asymmetric spillover effects over two groups of rooms. Specifically, a room with verified photos receives a positive spillover effect (higher demand) from an increase in the percentage of rooms with verified photos in the same neighborhood, while a room without verified photos receives a negative spillover effect. We also look at the impact of an increase in the percentage of rooms with verified photos in a neighborhood on the aggregate demand of the neighborhood. Our results suggest that neighborhood with more rooms having verified photos are likely to increase its overall demand at the neighborhood level. This is possible if, for example, guests perceive a better "neighborhood image" of a neighborhood which has more rooms with verified photos and hence are more likely to book rooms in that neighborhood. As a result, a neighborhood having fewer rooms with verified photos suffer from "stolen demand" by neighborhoods with "better neighborhood image" near it. With Airbnb as a test case, our finding may be generalized to a broader area in sharing economy/p2p platform.

¹³ Due to missing data when checking property availability 30 days in advance, in this analysis 383 treated units and 9249 untreated units are left for implementing PSM and calculating matching estimators.

References

- Abadie A. 2005. "Semiparametric Difference-in-Differences Estimators," *The Review of Economic Studies* (72:1). pp. 1-19.
- Angrist, J. D., Krueger, A. B. 1999. "Empirical Strategies in Labor Economics," in O. Ashenfelter and D. Card (eds.) *Handbook of Labor Economics*, Vol. 3A (Amsterdam: Elsevier Science), pp. 1277–1366.
- Ashenfelter, O., Card, D. 1985. "Using the Longitudinal Structure of Earnings to Estimate the Effect of Training Programs," *The Review of Economics and Statistics*, (67:4), pp.648-60.
- Athey, S., Imbens, G.W. 2006. "Identification and Inference in Nonlinear Difference-in-Differences Models," *Econometrica* (74:2), pp. 431–497.
- Avital, M., Andersson, M., Nickerson, J., Sundararajan, A., Alstyne, M.V., Verhoeven, D. 2014. "The Collaborative Economy: A Disruptive Innovation or Much Ado about Nothing?," Karahanna, Elena Srinivasan, Ananth Tan, Bernard (ed.). *Proceedings of the 35th International Conference on Information Systems*. Atlanta, GA.
- Aydin, T.O., Smolic, A., Gross, M. 2014. "Automated Aesthetic Analysis of Photographic Images", *IEEE Transactions on Visualization and Computer Graphics* (21 : 1), pp. 31 – 42.
- Bardhi, F., Eckhardt, G. 2012. "Access based consumption: The case of car sharing," *Journal of Consumer Research* (39), pp.881–898.
- Besley, T., Case. A. 2000. "Unnatural Experiments? Estimating the Incidence of Endogenous Policies," *Economic Journal* (110), pp.672–694.
- Blanz, V., Schölkopf, B., Ulthoff, H.B., Burges, C., Vapnik, C., Vetter, T. 1996. "Comparison of view-based object recognition algorithms using realistic 3d models," In C. von der Malsburg, W. von Seelen, J. C. Vorbrüggen, and B. Sendhoff, editors, *Artificial Neural Networks — ICANN'96*, pp.251 – 256, Berlin. *Springer Lecture Notes in Computer Science*, Vol. 1112.
- Boser, B. E. , Guyon, I. M., Vapnik, V. N. 1992. "A training algorithm for optimal margin classifiers,". In D. Haussler, editor, 5th Annual ACM Workshop on COLT, *ACM Press*, pp.144–152, Pittsburgh, PA.
- Card, D., Krueger, A. B. 1994. "Minimum Wages and Employment: A Case Study of the Fast-Food Industry in New Jersey and Pennsylvania," *American Economic Review* (84), pp. 772–793.
- Card, D. 1990. "The Impact of the Mariel Boatlift on the Miami Labour Market", *Industrial and Labor Relations Review* (44) pp. 245–257.
- Cervero, R., Golub, A., Nee, B. 2007. "City CarShare: Longer-Term Travel Demand and Car Ownership Impacts," *Transportation Research Record: Journal of the Transportation Research Board* (1992:1), pp.70–80.
- Cohen, M., Sundararajan, A. 2015. "Self-Regulation and Innovation in the Peer-to-Peer Sharing Economy," *The University of Chicago Law Review*. pp.116-133.
- Csurka, G., Dance, C., Fan, L., Willamowski, J., Bray, C. 2004. "Visual categorization with bags of keypoints," In: *ECCV SLCV Workshop*
- Cortes, C., Vapnik, V. 1995. "Support vector networks. Machine Learning," *Machine Learning*, 20:pp.273–297.
- Cusumano, M.A. 2015. "How traditional firms must compete in the sharing economy", *Communications of the ACM* (68:1), pp.32-34.
- Datta, R., Joshi, D., Li, J., Wang, Z. 2006. "Studying Aesthetics in Photographic Images Using a Computational Approach", *Computer Vision – ECCV* (953) pp. 288-301.
- Desnoyer, M., Wettergreen, D. 2010. "Aesthetic Image Classification for Autonomous Agents," *Pattern Recognition (ICPR)*, 2010 20th International Conference on, Istanbul, 2010, pp. 3452-3455.
- Elizabeth A. Stuart et al. 2014. "Using propensity scores in difference-in-differences models to estimate the effects of a policy change," *Health Serv Outcomes Res Methodol* (14:4), pp.166–182.
- Evelyn, M. R., Macmillan, D. 2014. "Uber Gets an Uber-Valuation", *The Wall Street Journal*/
- Garvey, G. T., Hanka, G. 1999. "Capital Structure and Corporate Control: The Effect of Antitakeover Statutes on Firm Leverage", *Journal of Finance* (54), pp. 519–546.
- Hamari, J., Sjöklint, M., Ukkonen, A. 2015. "The Sharing Economy: Why People Participate in Collaborative Consumption". *Journal of the Association for Information Science and Technology*. doi:10.1002/asi.23552.
- Hall, J.V., Krueger, A.B. 2015. "An Analysis of the Labor Market for Uber's Driver-Partners in the United States", <http://arks.princeton.edu/ark:/88435/dsp010z708z67d>.

- Heckman, J. J., Payner, B. S. 1989. "Determining the Impact of Federal Antidiscrimination Policy on the Economic Status of Blacks: A Study of South Carolina," *American Economic Review* (79). pp. 138–177.
- Jaakkola, T., Haussler, D. 1999. "Exploiting generative models in discriminative classifiers". In *NIPS*.
- Jin, G., Leslie, P. 2003. "The Effect of Information on Product Quality: Evidence from Restaurant Hygiene Grade Cards," *Quarterly Journal of Economics* (118), pp.409–451.
- Ke, Y., Tang, X., Jing, F. 2006. "The design of High-Level features for photo quality assessment," In *Proceedings of the 2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition* (1), pp.419–426. IEEE Computer Society Washington.
- Lowe, D. 2004. "Distinctive image features from scale-invariant keypoints,". *IJCV* (60:2), pp. 386-399.
- Luo Y., Tang, X. 2008. "Photo and video quality evaluation: Focusing on the subject,". In: *ECCV, 2008*
- Marchesotti, L., Perronnin, F., Larlus, D., Csurka, G. 2011. "Assessing the aesthetic quality of photographs using generic image descriptors", In *ICCV*, pp. 1784-1791
- Marchesotti, L., Murray, N., Perronnin, F. 2015. "Discovering Beautiful Attributes for Aesthetic Image Analysis", *International Journal of Computer Vision* (113:3). pp 246-266.
- Martin, E., Shaheen, S.A., Lidicker, J. 2010. "Impact of Carsharing on Household Vehicle Holdings,". *Transportation Research Record: Journal of the Transportation Research Board* (2143:1), pp.150–158.
- Murray, N., Marchesotti, L., Perronnin, F. 2012. "Learning to rank images using semantic and aesthetic labels,". In *BMVC*.
- Nishiyama, M., Okabe, T., Sato, I., Sato, Y. 2011. "Aesthetic quality classification of photographs based on color harmony," *Computer Vision and Pattern Recognition (CVPR)*, IEEE Conference. pp. 33-40.
- Perronnin, F., Dance, C. 2007. "Fisher kernels on visual vocabularies for image categorization,". In *CVPR*, pp. 1-8.
- Perronnin, F., Sánchez, J., Mensink, T. 2010. "Improving the fisher kernel for large-scale image classification,". In *ECCV* (6314), pp 143-156. (a)
- Perronnin, F., Liu, Y., Sánchez, J., Poirier, H. 2010. "Large-scale image retrieval with compressed fisher vectors," In *CVPR*. pp.3384 – 3391. (b)
- PwC report, "Consumer Intelligence Series: The sharing economy", pwc.com/CISsharing, 2015.
- Rolfe, W., Macmillan, D. 2015. "The Secret Math of Airbnb's \$24 Billion Valuation", *The Wall Street Journal*.
- Romero, J., Machado, P., Carballal, A., Osorio, O. 2011. "Aesthetic Classification and Sorting Based on Image Compression", *Applications of Evolutionary Computation*, Volume 6625 of the series Lecture Notes in Computer Science pp.394-403.
- Rosenbaum, P., and Rubin, D. 1983. "The Central Role of the Propensity Score in Observational Studies for Causal Effects." *Biometrik* (70), pp.41–55.
- Sivic, J., Zisserman, A. 2003. "Video Google: A text retrieval approach to object matching in videos," In *ICCV* (2), pp. 1470-1477.
- Sánchez, J., Perronnin, F., Mensink, T., Verbeek, J. 2013. "Image Classification with the Fisher Vector: Theory and Practice," *International Journal of Computer Vision* (105:3), pp. 222-245.
- Schmidt, M. 1996. Identifying speaker with support vector networks. In *Interface '96 Proceedings*, Sydney.
- Sundararajan, A. 2016. "The Sharing Economy: The End of Employment and the Rise of Crowd-Based Capitalism", The MIT Press.
- Tella, R.D., SCHARGRODSKY, E. 2004. "Do Police Reduce Crime? Estimates Using the Allocation of Police Forces After a Terrorist Attack", *The American Economic Review* (94:1), pp.115-133.
- Vapnik, V. 1979. "Estimation of Dependences Based on Empirical Data [in Russian]". Nauka, Moscow. (English translation: Springer Verlag, New York, 1982).
- Ufford, S. 2015. "The Future Of The Sharing Economy Depends On Trust", *Forbes*. <http://www.forbes.com/sites/theyec/2015/02/10/the-future-of-the-sharing-economy-depends-on-trust/#21a3370b58ff>
- Wang, W., Sheikh, H. R., Bovik, A. C. 2003. "The Handbook of video databases: design and applications," chapter 41. *CRC press*.
- Yang, D. 2008. "Can Enforcement Backfire? Crime Displacement in the Context of Customs Reform in the Philippines". *The Review of Economics and Statistics* (90:1), pp.1–14.

Zervas, G., Proserpio, D., Byers, J. 2016. "The Rise of the Sharing Economy: Estimating the Impact of Airbnb on the Hotel Industry," Boston U. School of Management Research Paper No. 2013-16.