You Get What You Give: Theory and Evidence of Reciprocity in the Sharing Economy

Davide Proserpio
Marshall School of Business
University of Southern California

Wendy Xu
Department of Economics
Boston University

Georgios Zervas
Questrom School of Business
Boston University

June 26, 2018

Abstract

We develop an analytical framework of peer interaction in the sharing economy that incorporates reciprocity, the tendency to increase (decrease) effort in response to others' increased (decreased) effort. In our model, buyers (sellers) can induce sellers (buyers) to exert more effort by behaving well themselves. We demonstrate that this joint increased effort can improve the utility of both parties and influence the market equilibrium. We also show that bilateral reputation systems, which allow both buyers and sellers to review each other, are more responsive to reciprocity than unilateral reputation systems. By rewarding reciprocal behavior, bilateral reputation systems generate trust among strangers and informally regulate their behavior. We test the predictions of our model using data from Airbnb, a popular peer-to-peer accommodation platform. We show that Airbnb hosts that are more reciprocal receive higher ratings, and that higher rated hosts can increase their prices. Therefore, reciprocity affects equilibrium prices on Airbnb through its impact on ratings, as predicted by our analytical framework.

1 Introduction

Over the past few years, we have witnessed the rapid growth of peer-to-peer marketplaces like Airbnb (accommodation), Uber (transportation), and TaskRabbit (chores and errands). Compared to traditional marketplaces for similar services, peer-to-peer markets have two distinguishing features. First, service quality is heterogeneous and can vary significantly both between suppliers and from one occasion to the next. Second, transactions in peer-to-peer markets entail the close interaction of individual buyers and sellers.

In this paper, we study these peer interaction patterns, focusing in particular on reciprocity: a social norm under which people respond to others' actions with an equivalent action. We provide theoretical and empirical evidence that reciprocity regulates the behavior of peer-to-peer market participants thus affecting both service quality and market outcomes. In doing so, we provide an interesting connection between peer-to-peer marketplaces and behavioral economics. Our model is also the first, to the best of our knowledge, to apply interdependent preferences to study peer interactions in the sharing economy, and to show that this interdependence is more salient in these new marketplaces when compared to traditional marketplaces for similar services.

Before proceeding, we note the term reciprocity has been used to refer to two different concepts in two main strands of literature that we build upon: the literature on online reviews and reputation, and the behavioral economics literature on trust, fairness, and cooperation. In the online reputation literature, reciprocity refers to strategic reviewing behavior where positive (negative) feedback from one party is likely to be "reciprocated" with positive (negative) feedback from the other party (Dellarocas and Wood, 2008; Resnick and Zeckhauser, 2002). In the behavioral economics literature, reciprocity refers to the tendency of market participants to respond to good (bad) behavior with good (bad) behavior (Sobel, 2005). We use the term reciprocity in the behavioral economics sense.

The sharing economy is a natural setting for reciprocity to arise. Take for instance Airbnb, an online marketplace for short-term accommodation rentals that has emerged as an alternative to hotels (Zervas et al., 2017; Farronato and Fradkin, 2018). According to Airbnb, most hosts rent the place they live in, and are likely to closely interact with their guests.[1] Some Airbnb hosts provide home cooked meals, take their guests on neighborhood tours, or pick their guests up from the airport. Airbnb guests can reward exceptional hospitality by keeping the place clean, being respectful of neighbors, or leaving small gifts for their hosts. Airbnb user testimonials suggest that this type of behavior is both common and greatly

[1] See: http://blog.airbnb.com/economic-impact-airbnb/.

valued by guests and hosts alike.[2]

We formalize this intuition by developing an analytical framework of peer interaction in the sharing economy. Taking Airbnb as our use case, we model the quality of each Airbnb stay as a function of the joint effort of hosts and guests. In deciding how much effort to exert during each stay, hosts and guests take into account each others' reciprocity preferences. These preferences can vary. For example, some hosts will respond to a small gift by taking better care of their guests, while others won't. Airbnb guests can induce reciprocal hosts to exert more effort by behaving better themselves. Reciprocal hosts respond to increased guest effort by increasing their own effort levels and, as a result, earn higher ratings from their guests. Therefore, in our model hosts who are more reciprocal have higher ratings. In turn, higher ratings lead to increases in demand and, consequently, more reciprocal hosts can charge higher prices. Thus, our model predicts that reciprocity influences prices through its effect on ratings.

After presenting our theoretical framework, we use data from Airbnb to test the predictions of our model. First, we show that reciprocity and ratings are positively correlated. To do so, we begin by identifying Airbnb hosts who are likely to be reciprocal, *i.e.*, hosts who value their guests' behavior. Because we cannot directly measure reciprocity, we rely on proxies. Our preferred proxy is the length of the reviews hosts leave for their guests. Unlike listing reviews, which usually describe both the host and her property, the reviews hosts leave for their guests focus on the person being accommodated. We hypothesize that hosts who take the time to write detailed reviews about their guests, are more likely to care about their guests' conduct, and, therefore, to value reciprocity more. As hypothesized, we find that hosts who write longer guest reviews have higher ratings. This result is robust to alternative proxies for reciprocity, which we discuss in detail in Section 5.

Our model is also able to explain a novel pattern we identify in the Airbnb data that is not easily explained under alternative plausible models of peer behavior. We show that Airbnb hosts who list their properties for rent on Airbnb frequently (whom, for simplicity, we will refer to as *professional* hosts) have lower ratings than those with less frequent market participation (whom we will refer to as *casual* hosts). Moreover, this relationship is independent of the degree of social interaction between guests and hosts, and thus unlikely to be the effect of social bias arising from face-to-face interaction between hosts and guests (Fradkin et al., 2017). Taken together, these two observations suggest that casual hosts offer higher levels of quality than professional hosts. To explain this finding, we show that there is a negative association between reciprocity and market participation. Thus, we argue that the higher ratings of casual hosts can, at least in part, be explained by their increased preference

[2]See, for instance, `https://www.airbnb.com/info/why_host`.

for reciprocity.

We conclude our analysis by showing that reciprocity has important implications for market equilibrium. Our model predicts that reciprocity affects ratings which in turn affect prices. We test the prediction that Airbnb listing prices respond positively to ratings by exploiting a unique feature of the Airbnb reputation system: Airbnb discloses a host's average rating only after the host collects three reviews. We argue that the timing of these rating disclosures is exogenous and show that listing prices respond positively to the disclosure of high ratings, even after controlling for changes in listing quality. Taken together, our results suggest that reciprocity can explain the higher ratings, and thus higher prices, of casual hosts.

Besides demonstrating the impact of reciprocity on market equilibrium and prices, our work contributes to a more comprehensive understanding on how trust is generated in peer-to-peer markets (Tadelis, 2016; Einav et al., 2016). Reciprocity in the sharing economy can act as a social norm (Gouldner, 1960), generating trust between market participants (Malmendier et al., 2014). When two agents trade, one agent's poor conduct can trigger negative reciprocity and the agent is punished by the other party. Similarly, an agent's good conduct triggers positive reciprocity and the agent is rewarded by the other party. Thus, agents may choose to behave well both to motivate good conduct from the other party, and to avoid the threat of negative reciprocity. Through this mechanism, reciprocity can regulate behavior and generate trust among users.

2 Related literature

Next, we discuss two relevant strands of literature, on reciprocity and reputation, and highlight connections to our work.

Our work studies the role of reciprocity in promoting improved behavior and cooperation in peer-to-peer platforms. The definition of reciprocity that we adopt in this paper is similar to that of Sobel (2005), who defines reciprocity as "a tendency to respond to perceived kindness with kindness and perceived meanness with meanness and to expect this behavior from others". A large body of research finds that reciprocity is a important determinant of behavior (see, for example, Kahneman et al. (1986); Fehr and Gächter (2000)). Fehr and Gächter (2000) survey applications of reciprocity in several areas of economics, emphasizing the role of reciprocity in encouraging collective action. In our setting, we investigate how hosts and guest act collectively to improve their joint experience.

Rabin (1993) was the first to propose intention-based reciprocity, the tendency of players to reward good intentions and punish bad intentions. Falk and Fischbacher (2006) investigate the role of intentions in driving reciprocity whereby actions with identical consequences can result in different reciprocal responses depending on how "fair" these action are perceived to be. Malmendier et al. (2014) discuss recent theoretical and experimental developments in reciprocity research, arguing that "reciprocity is internal in that it arises from an individual's preference to act in a way that rewards good behavior by others". Our concept of shared experience utility builds on these insights, by enabling marketplace participants to rewards or punish each others' behavior.

Finally, Rotemberg (2006) survey economic models of reciprocity in the context of organizational economics, distinguishing reciprocity from other types of social preference, such as altruism. We also consider altruism as a possible way to explain our findings, but argue that reciprocity is better suited to explaining the patterns in our data.

Our work also relates to a large marketing literature on reputation, feedback systems, and rating biases in online markets. A number of papers have studied the effects of seller reputation across different settings. For example, Luca (2016); Anderson and Magruder (2012) show that Yelp ratings affect restaurant revenue and the likelihood of being sold out, and Yoganarasimhan (2013) shows that better rated freelancers in an online labor market are more likely to be chosen by buyers and can charge higher prices. Our theoretical model makes similar predictions to Yoganarasimhan (2013), and we find similar patterns in the Airbnb data: hosts with higher ratings tend to charge higher prices.

Closely related to our setting, a number of papers have studied bilateral reputation systems. Early studies focused on eBay, one of the first platforms introducing a bilateral reputation mechanism, in which buyers and sellers review each other after a transaction. Among others, Dellarocas and Wood (2008) and Resnick and Zeckhauser (2002) show that ratings on eBay are extremely positive, and that there is substantial amount of feedback reciprocation – a practice in which the receipt of a positive (negative) feedback from a transaction party increases the likelihood of the other party to also report a positive (negative) feedback. More recently Fradkin et al. (2017) studied the bilateral reputation system of Airbnb, and reported similar reporting bias. Fradkin et al. (2017) argue that the bias in this case is generated by socially induced reciprocity whereby users tend to under-report negative feedback following an in-person interaction with the other party. Horton and Golden (2015) also study reputation inflation in online markets, and show that negative ratings are more likely to be under-reported when they are public.

For an extensive review of the literature on reputation and feedback systems on online platforms, we refer the reader to Tadelis (2016), Edelman (2017), and Seiler et al. (2018).

3 Theoretical framework

In this section, we introduce our analytical framework of peer interaction in sharing economy. Below, we describe the model in the setting of Airbnb, but the model can be extended to other markets.

3.1 Setup

Since our main focus is the interaction between hosts and guests, we assume a monopolistic host and a continuum of guests in our model. Prices are posted by the host, which matches the price-posting feature on Airbnb.

The extensive-form game of our model is as follows:

1. In period 1, the host chooses price P_1 for her listing, and every customer decides whether to enter the market, *i.e.*, pay the host to book the accommodation.

2. In period 2 (the accommodation period), the guest stays at host's listing. Both the host and the guest i determine the effort level to exert during the stay. Then, each guest i who stays at the host's property publishes a rating $r_{h,i}$, and the host publishes a rating r_i for the guest i.

3. In period 3, another unit mass of guests enters the market. The host and period 3 guests observe the first period demand and the average rating disclosed in period 2. The host chooses price P_3 for her listing, and the customers decide whether to enter the market according to the disclosed rating and the price observed.

We solve the equilibrium by backward induction from period 2 to period 1. Since the subgame in period 2 is the main focus of our analysis, we introduce the setup of period 2 first.

Period 2: Accommodation In period 2, guest i and the host h choose effort levels by maximizing their respective ex-post utilities:

$$U_i(e_i|e_h, r_i) = v_h + \alpha_i u(e_i, e_h) - C_i(e_i) + \beta_i r_i \tag{1}$$

$$U_h(e_h|e_i, r_{h,i}) = v_i + \alpha_h u(e_i, e_h) - C_h(e_h) + \beta_h r_{h,i} \tag{2}$$

The two utility functions are symmetric in the identities of the agents, i and h. To simplify notation, we will use the subscript j – *e.g.* U_j – to refer to either utility function.

Utility is composed of three parts. The first part, $v_j + \alpha_j u(e_i, e_h)$, denotes utility obtained during accommodation. The terms v_h and v_i are exogenous factors affecting accommodation quality. For example, v_h could be the location of the listing. Similarly, v_i are the characteristics of guest i that affect the host's utility and that are not endogenously determined by the guest's effort, for instance whether guest i has a pet. We assume v_h and v_i to be fixed and exogenously given, and to follow a uniform distribution over $[0, \bar{v_h}]$ and $[0, \bar{v_i}]$ respectively.

The term $u(e_i, e_h)$ captures the utility the host h and the guest i derive from interacting with each other, which we refer to as the *shared experience utility*. The shared experience utility is a function of host effort e_h and guest effort e_i. An example of e_h may be a smooth check-in experience, or providing information about local attractions. An example of e_i may be a guest who follows the house rules. Introducing the shared experience term in the utility function allows us to formally model the social interaction between hosts and guests. Related work in marketing and economics, has also considered the idea that interpersonal interaction can directly affect agents' utility (see, for example, Charness (2004) and Blanchard et al. (2016)).

We allow for heterogeneity on the weight agents place on the shared experience utility by introducing the terms α_j, which we will refer to as *reciprocity weights*. This allows for some hosts to care more than other about the behavior of their guests.

We assume that both host and guest effort positively impact the shared experience utility, *i.e.*:

$$\frac{\partial u(e_i, e_h)}{\partial e_i} > 0, \; \frac{\partial u(e_i, e_h)}{\partial e_h} > 0. \tag{3}$$

Further, we model reciprocity by assuming:

$$\frac{\partial^2 u(e_i, e_h)}{\partial e_i \partial e_h} > 0. \tag{4}$$

This condition states that increased host effort improves the guest's marginal shared experience utility, *i.e.*,

$$\text{if } e_h > e_h^{'}, \text{ then } \tfrac{\partial u(e_i, e_h)}{\partial e_i} > \tfrac{\partial u(e_i, e_h^{'})}{\partial e_i}, \text{ given the same level of } e_i.$$

Similarly, increased guest effort improves the host's marginal shared experience utility. Our reciprocity assumption is in line with similar assumptions invoked in intrinsic reciprocity models, such as the intention-based reciprocity model (Rabin, 1993) and type-based model (Levine, 1998). As pointed out by Malmendier et al. (2014), "under these theories, people reciprocate because another person's kind act or benevolent nature increases the intrinsic utility of acting kindly toward this person. Thus, such preferences are internal in that they arise from an individual's preference to act in a way that rewards good behavior by others". Our assumption $\frac{\partial^2 u(e_i,e_h)}{\partial e_i \partial e_h} > 0$ incorporates the above intuition in the setting of the sharing economy in a tractable way.

The second part of the utility function, $C_j(e_j)$, represents the cost of exerting effort. In the analysis below, we adopt the widely-used quadratic functional form for the effort cost function, *i.e.*, $C_j(e_j) = \frac{1}{2}c_j e_j^2$.

Finally, the third part of the utility function, r_i, denotes the rating guest i receives from the host, and $r_{h,i}$ denotes the rating host h receives from the guest. The term β_j, which we will refer to as the *reputation weight*, denotes how much agent j cares about this rating. On Airbnb, the rating a guest receives affects her chance of being accepted by future hosts when they apply for accommodation. Hosts care about ratings because ratings signal quality and hence affect the expected demand of the listing. Here, we introduce the term $\beta_j r_j$ in the utility function as a simple and tractable device for comparing unilateral and bilateral rating systems (see Section 3.2). However, even if we do not assume that $\beta_h r_h$ enters host utility, hosts still care about the ratings since guests infer that higher ratings imply higher quality, and thus affect expected demand.

In our main analysis, we assume that agents truthfully report their utility of the accommodation experience in ratings, *i.e.*,

$$
\begin{aligned}
r_i &= v_i + \alpha_h u(e_i, e_h) && (5) \\
r_{h,i} &= v_h + \alpha_i u(e_i, e_h). && (6)
\end{aligned}
$$

However, in Section 3.3, we relax this truth-telling assumption and show that our model is robust to this assumption.

Lastly, we assume the vector of model parameters $\Gamma_j \equiv [\alpha_j, \beta_j, c_j, v_j]$ to be private information of agent j prior to accommodation. However, the distribution of Γ_j, which we denote by $F(\Gamma_j)$, is assumed to be common knowledge.

Period 1: Pre-accommodation In period 1, the host posts the price P_1, and the guests decide whether to request accommodation. The transaction volume Q_1 is determined at this stage. The utility of the host is composed of two parts: the monetary revenue, $P_1 Q_1$, and the expected utility in period 2. The utility of the host in period 1 is given by:

$$V_h(P_1) = P_1 Q_1 + \int U_h(e_h^*, e_i, v_i, r_{h,i}) dF(\Gamma_i), \tag{7}$$

where $U_h(e_h^*, e_i, v_i, r_{h,i})$ is the ex-post utility the host obtains during the period 2. e_h^* is the optimal effort the host exerts during guest i's accommodation; e_i is the equilibrium effort exerted by guest i, v_i is the objective "value" of guest i affecting the welfare of the host, and $r_{h,i}$ is the rating the host receives from guest i.

In period 1, the early guests choose whether to make a booking request based on their ex-ante expected utility. The ex-ante utility of guest i who chooses to enter the market is given by:

$$V_i(P_1) = \int U_i(e_i^*, e_h, v_h, r_i) dF(\Gamma_h) - P_1, \tag{8}$$

where $U_i(e_i^*, e_h, v_h, r_i)$ is the ex-post utility of period 2. As before, the ex-post utility depends on e_i^*, the optimal effort exerted by guest i; e_h, the equilibrium effort level exerted by the host in the transaction with guest i, and r_i, the rating guest i receives and v_h.

3.2 Propositions

To simplify proofs, we assume that the shared experience utility takes the Cobb-Douglas functional form, *i.e.*,

$$u(e_i, e_h) = e_i^k e_h^{1-k}, \ k \in (0,1). \tag{9}$$

However, our results hold under any non-separable form of $u(e_i, e_h)$ satisfying the following conditions:

$$\frac{\partial^2 u(e_i, e_h)}{\partial e_i \partial e_h} > 0, \ \frac{\partial u(e_i, e_h)}{\partial e_i} > 0, \ \frac{\partial u(e_i, e_h)}{\partial e_h} > 0. \tag{10}$$

We solve the equilibrium by backward induction. In period 2, the host h and the guest i choose their effort levels. At this stage, the uncertainty on the parameters is resolved, and thus the optimization problem of the guest i and the host are given by:

$$\max_{e_i} U_i(e_i|e_h, r_i) = \max_{e_i} \{v_h + \alpha_i u(e_i, e_h) - C_i(e_i) + \beta_i r_i\} \tag{11}$$

$$\max_{e_h} U_h(e_h|e_i, r_{h,i}) = \max_{e_h} \{v_i + \alpha_h u(e_i, e_h) - C_h(e_h) + \beta_h r_{h,i}\}, \tag{12}$$

where v_h and v_i are just exogenous variable and do not affect the optimization problem.

After invoking the truth-telling assumption, the first order conditions for the two optimality problems are:

$$k(\alpha_i + \beta_i \beta_h)(\frac{e_i^*}{e_h^*})^{k-1} - c_i e_i^* = 0 \tag{13}$$

$$(1-k)(\alpha_h + \beta_h \beta_i)(\frac{e_i^*}{e_{h,i}^*})^k - c_h e_h^* = 0. \tag{14}$$

By the optimality conditions, we obtain the following closed-form solution for e_i^* and e_h^*:

$$
\begin{aligned}
e_i^* &= A(\alpha_h + \beta_h \alpha_i)^{\gamma}(\alpha_i + \beta_i \alpha_h)^{1-\gamma} \quad (15)\\
e_h^* &= B(\alpha_h + \beta_h \alpha_i)^{\mu}(\alpha_i + \beta_i \alpha_h)^{1-\mu}, \quad (16)
\end{aligned}
$$

where $\gamma \equiv \frac{1-k}{2}$ and $\mu \equiv 1 - \frac{k}{2}$; $A \equiv (\frac{k}{c_i})^{1-\gamma}(\frac{1-k}{c_h})^{\gamma}$ and $B \equiv (\frac{k}{c_i})^{1-\mu}(\frac{1-k}{c_h})^{\mu}$.

We then present three propositions, which we explain intuitively leaving formal proofs for Appendix A.

Proposition 1. *A host's average rating on Airbnb positively depends on the host's reciprocity weight α_h,* i.e., *letting $R_{airbnb} \equiv \int r_{h,i} di$ denote the average rating of host h we have*

$$\frac{\partial R_{airbnb}}{\partial \alpha_h} > 0, \forall \alpha_h > 0. \quad (17)$$

Intuitively, Proposition 1 captures the following process: a host with higher reciprocity weight α_h is, on average, more willing to improve the shared experience by exerting effort. Because of reciprocity, her guests are willing to exert more effort themselves. Ultimately, the increased effort level of both agents increases the shared experience utility, which is reflected in higher ratings that guests leave for the host.

Formally, from Equations 15 and 16, we have:

$$
\begin{aligned}
\frac{\partial e_h^*}{\partial \alpha_h} &> 0 \quad (18)\\
\frac{\partial e_i^*}{\partial \alpha_h} &> 0. \quad (19)
\end{aligned}
$$

The above inequalities show that the equilibrium effort levels of the host and the guest increase with the host's weight on reciprocity. In turn, increased effort increases the shared experience utility $u(e_i, e_h)$. Finally, the higher level of $u(e_i, e_h)$ is reflected in the higher rating left by the guest i to for the host h, which leads to a higher average rating for the host.

In the next proposition, we show that Airbnb's bilateral reputation system reveals more information about hosts' reciprocity weights than a unilateral reputation system (where hosts do not rate their guests.) For instance, one consequence of this proposition is that bilateral ratings are more informative than unilateral ratings about host attributes like hospitality. Thus, to the extent that platforms like Airbnb want to encourage reciprocal behavior, the bilateral rating system is a better choice.

Proposition 2. *Let $R_{airbnb} \equiv \int_i r_{h,i}^{air} di$ denote the average rating of a host on Airbnb's* *bilateral reputation system, and* $R_{uni} \equiv \int_i r_{h,i}^{uni} di$ *denote the average rating of the same host on a unilateral reputation system. Then we have*

$$\frac{\partial R_{airbnb}}{\partial \alpha_h} > \frac{\partial R_{uni}}{\partial \alpha_h} > 0, \tag{20}$$

$$\frac{\partial R_{uni}}{\partial \beta_h} > \frac{\partial R_{airbnb}}{\partial \beta_h} > 0. \tag{21}$$

The proposition above states that ratings on both reputation systems respond positively to the reciprocity and reputation weights, α_h and β_h respectively. However, for the same host and guests, the reciprocity weight plays a bigger role in the bilateral rating system, while the reputation weight plays a bigger role on the unilateral system. Next, we illustrate the intuitive reasoning behind this proposition.

We begin with inequality 20, which quantifies the impact of the reciprocity weight, α_h. From Proposition 1, we know that $\frac{\partial R_{airbnb}}{\partial \alpha_h} > 0$. It is straightforward to show that $\frac{\partial R_{uni}}{\partial \alpha_h} > 0$ also holds, since inequalities 18 and 19 hold independently of the design of the reputation system. Thus, increased reciprocity weights leads to increased effort levels, which lead to increased ratings.

Next we show that ratings in the bilateral reputation system are more sensitive to reciprocity weight than ratings in a unilateral reputation system, *i.e.*, $\frac{\partial R_{airbnb}}{\partial \alpha_h} > \frac{\partial R_{uni}}{\partial \alpha_h}$. We decompose the explanation in two steps. First, we show that for the same host h, guest i exerts less effort under the unilateral reputation system. Second, we demonstrate that, if two hosts who transact with the same guest only differ in their reciprocity weights, the higher the guest's effort is, the larger the effort difference between the two hosts.

The first step is intuitively explained by the fact that the bilateral reputation system introduces reputation concerns for the guest i. Therefore, under a bilateral system, guest i exerts more effort.

We then show that a guest's effort positively affects the effort difference between two hosts transacting with her using an example. Assume two hosts, Ann and Bob. Ann has higher reciprocity weight than Bob, *i.e.*, $\alpha_{Ann} > \alpha_{Bob}$. From the first order condition of the host's optimality problem, we have that, given the same level of e_i, $e_{Ann}(e_i) > e_{Bob}(e_i)$. This means that Ann exerts more effort than Bob if they transact with the same guest i. Moreover, because Ann has higher reciprocity weight, when guest i increases her effort, Ann is more willing to reciprocate, *i.e.*, the difference $e_{Ann}(e_i) - e_{Bob}(e_i)$ increases in e_i. Finally, since the shared experience utility, $u(e_i, e_h)$, depends positively on e_h, the larger difference between the effort of Ann and Bob translates to a larger difference in the shared experience utility, which is then revealed in the rating difference between Ann and Bob.[3]

Our last proposition establishes the positive relationship between the price and ratings.

Proposition 3. *On Airbnb, prices increase after a positive ratings shock, and decrease after a negative shock. Given the same price P_1 at period 1, we have the following relationship for prices posted at period 3 and ratings R_{airbnb} and R'_{airbnb} disclosed in period 2,*

$$if\ R_{airbnb} > R'_{airbnb}\ then\ P_3(R_{airbnb}) > P_3(R'_{airbnb}). \tag{22}$$

Proposition 3 shows that prices respond to ratings. This results from the informativeness of the ratings: a host with higher average rating will look more attractive to future guests, and therefore will have higher expected demand than another host with the same price but lower average rating. Since period 3 price is determined by expected demand, hosts with higher average rating will raise their prices.

The propositions above lead to the following two claims:

Claim 1. *Hosts with higher reciprocity weights have higher average ratings.*

Claim 2. *For a given set of hosts, the one with higher reciprocity weight and lower reputation weight is ranked higher on Airbnb than her counterpart with lower reciprocity weight and higher reputation weight, while the opposite is true on the unilateral reputation system.*

3.3 Relaxing the truth-telling assumption

Next, we relax the truth-telling assumption, allowing selection bias and rating inflation, as observed in Fradkin et al. (2017). Again, we present the main intuition for why our results hold under these relaxed assumptions, providing detailed proofs in Appendix A.4.

Under this assumption, the rating given to a host, denoted by $r_{h,i}$, is a choice variable of the guests, which means that guests do not automatically report their utility in the ratings they leave for their hosts. Instead, guests can choose not to report a rating at all, report a rating not equal to their utility, or truthfully report their utility in the rating.

Formally stated, let $\omega_i(e_i, e_h, v_h) \equiv v_h + \alpha_h u(e_i, e_h)$ be the total utility that guest i obtains from the accommodation. Let ϕ denote the mapping from the guest's utility to ratings, *i.e.*, $\phi : \Omega \to R$, where Ω and R denote full set of ω_i and full set of $r_{h,i}$. In this general setup, not disclosing ratings due to selection bias is modeled as $\phi(\omega_i) = 0$, for some ω_i. Meanwhile, when the guests choose what rating to report, we allow guests to inflate their ratings, *i.e.*,

[3]Following a similar reasoning, we can show that the ratings on both systems positively relate to the host's weight on reputation, and that the unilateral review system responds more to the reputation weight, β_h.

$\phi(\omega_i) > \omega_i$. Also, we assume ϕ to be a weakly increasing function of ω_i. This assumption rules out $\phi(\omega_i)$ being constant or decreasing over ω_i.[4] These two cases are unrealistic since they contradict the empirical observation that Airbnb ratings are informative about listing quality (Fradkin et al., 2017).

Under this extended model, our main result, namely that ratings reveal the reciprocity weight, still holds. The sketch of the proof follows. Formally, since $r_{h,i} = \phi(\omega_i) \equiv \phi(v_h + u(e_i, e_h))$ and ϕ is weakly increasing, $r_{h,i}$ still increases with α_h. Thus, the average rating of the host, also increases with α_h. Therefore, even when we allow for selection bias and rating inflation, the ratings reflect, to some extent, the accommodation experience under weak assumptions on rating informativeness.[5] Finally, since the guest's accommodation experience is determined by the host's reciprocity weight, the guest's rating about the accommodation still reveals the host's reciprocity weight.

4 Airbnb and data

4.1 Airbnb

We use data from Airbnb to motivate our model and test its predictions. Airbnb, which launched in 2008, is a peer-to-peer marketplace for short term accommodation rentals. Airbnb *hosts* offer private or shared accommodation for rent to prospective *guests*. The Airbnb marketplace has seen a dramatic growth over the last few years. At the beginning of 2016 the platform listed approximately 3 million properties from 640,000 hosts in over 150,000 cities and 52 countries.[6] Over 80 million guests have used Airbnb, and with a market valuation of $30B, Airbnb is one of the world's largest accommodation brands.[7]

To build trust among users, Airbnb uses a bilateral reputation system. Hosts and guests can optionally review each other. The text of these reviews is publicly disclosed but their star-ratings are not. Instead, Airbnb only discloses average ratings aggregated across at multiple reviews. Prior to July 2014, Airbnb users had the option of reviewing each other within a 30-day window following the conclusion of each stay. During this 30-day window, reviews were revealed as they were submitted. This sequential revelation mechanism allowed

[4] $\phi(\omega_i)$ decreasing with ω_i implies that guests give hosts providing worse service strictly higher ratings. $\phi(\omega_i)$ constant over ω_i implies guests give all hosts the same rating regardless of quality.

[5] For example, a guest may choose not to disclose her rating after a bad experience. However, a guest does not rate a bad experience better than a good one. Therefore, the rating a guest discloses still weakly reveals the quality of her experience.

[6] See: `http://expandedramblings.com/index.php/airbnb-statistics/`

[7] See `http://qz.com/329735/airbnb-will-soon-be-booking-more-rooms-than-the-worlds-largest-hotel-chains/`

for retaliatory reviewing: the second reviewer could punish the first reviewer by submitting a negative review in exchange for receiving a negative review (Fradkin et al., 2017). In July 2014 Airbnb made a major change to its reputation system by shortening the review window to 14 days and only revealing reviews simultaneously after the review submission deadline. In doing so, Airbnb lessened the possibility of retaliatory reviewing.[8] To reduce the possibility of analyzing ratings that are biased by retaliatory reviewing, we limit our dataset to hosts that entered the Airbnb marketplace from July 2014 onwards.

4.2 Data

We compile a novel dataset of Airbnb listing entry, exit, prices, supply, demand, and reviews. Our dataset is a weekly panel of U.S. Airbnb listings spanning a 17-month period from the beginning of July 2014 to the end of November 2015. During this timeframe, we collected information on all US listings and their hosts from the Airbnb website with weekly frequency. The final panel contains $3,295,188$ listing-week observations for $198,743$ distinct listings and $137,687$ distinct hosts, whose accounts were created on or after July 1, 2014. For each listing, we observe various characteristics including location, listing type (*e.g.*, apartment, house, *etc.*), bed type, number of listing photos, price, star-rating, and number of reviews. Additionally, for each host, we observe reviews left and received, and the number of properties listed by the host on Airbnb.

Airbnb allows hosts to select which days of the year their listings are available for rent without the need to add or remove the listing from the platform. To do so, hosts use a calendar, on which they mark available days and set prices. In addition, Airbnb hosts can make their listings *instantly bookable*, forgoing the opportunity to reject certain guests. We collected calendar information (whether a listing was available for booking, booked, or busy) between September 2014 to September 2015.

Given this data, we define *market participation* as the fraction of days a property was listed for rent (regardless of whether a day was eventually booked or not) during our observation period. The final dataset contains market participation information for $101,596$ listings and $74,909$ hosts. Out of these listings, $51,697$ have a star-rating (Airbnb only assigns a star-rating to listings with at least three reviews.) Figure 1 displays the distribution of star-ratings for the subset of listings for which we know both their market participation and star-rating. As in previous work (Fradkin et al., 2017; Zervas et al., 2015), we find that most of the listings (91.6%) have a star-rating of at least 4.5-stars. In Figure 2, we plot the probability density function of market participation for the same subset of listings, and find

[8]See: `http://blog.airbnb.com/building-trust-new-review-system/`

that, on average, a listing is listed on the platform 85% of the time.

For every listing in our data, we have on average 16 weekly observations. Listings may have fewer or more weekly observations due to entry and exit. The average listing price is \$229, the average number of reviews hosts received from guests is 4.5, and the average star-rating of these reviews is 4.7. On average, hosts write reviews that are shorter than the reviews they receive from guests. Hosts in our data received reviews with an average length of 351 characters, and left reviews 151 characters long. Finally, by November 2015, the instantly bookable feature is enabled for $30,767$ listings.

5 Evidence of Reciprocity on Airbnb

In this section, we use data from the Airbnb platform to provide empirical evidence for the predictions of our model. Our model makes two key predictions. First, hosts that are more reciprocal should have higher ratings. Seconds, higher rated hosts – which includes hosts that are more reciprocal – should be able to charge higher prices.

5.1 The relationship between reciprocity and ratings

Proposition 1 of our model states that a host's ratings on Airbnb are positively related to the host's reciprocity weight, α_h. In other words, hosts who are more reciprocal should have higher ratings. To test this prediction we need to know hosts' reciprocity weights, a_h. However, reciprocity is not directly observable. Instead, we attempt to find proxies in our data that are correlated with reciprocal behavior.

Our first (and preferred) proxy for reciprocity weight is whether the host writes long reviews about her guests. Intuitively, a host that cares more about the overall Airbnb experience will take more time to describe this experience in a review. Note that, as explained in Section 4, Airbnb employs a double-blind review mechanism in which the content (and length) of the host's review is not disclosed to the guest until either the guest submits her own review or 14 days have passed. Therefore, this proxy cannot have a direct impact on guest ratings (*e.g.*, a good review from the host to the guest cannot incentivize a good review from the same guest to the host.) Further, the reviews that hosts leave for guests are not displayed on the hosts' Airbnb pages and, therefore, hosts have little incentive to behave strategically with respect to reviews they leave for their guests.[9] Therefore, we hypothesize

[9]To read reviews a host left for past guests one has to: a) find out who the past guests were by looking at the host profile and checking which guests left a review for the host, b) look up the Airbnb profiles of each of these guests, and c) manually scan each guest profile to locate a review left for the guest by the host in question. We assume that the vast majority of Airbnb users do not engage in this behavior.

that hosts that leave longer reviews about their guests are more likely to have a higher reciprocity weight.

The second proxy we use is whether the host has activated the "Instant Book" feature. Similar to hotel reservations, reservation requests for instantly bookable Airbnb listings do not require explicit host approval. We hypothesize that hosts using the Instant Book feature have more weight on the reputation utility r_j than on the shared experience utility $u(e_i, e_h)$.

We test our hypotheses using the following model:

$$\begin{aligned} \text{Star-rating}_i &= \beta_1 \text{ log Host-to-Guest Review Length}_i \\ &+ \beta_2 \text{ Not Instant Bookable}_i + \mathrm{X}_i\gamma + \epsilon_i, \end{aligned} \tag{23}$$

where the dependent variable is the star-rating of host i. log Host-to-Guest Review Length$_i$ and Not Instant Bookable$_i$, whose coefficients are of interest, are the average (log of the) length of the reviews written to her guests, and whether the host listing is not instantly bookable, respectively. In X_i we include a wide set of controls that can affect the host star-rating. We report the estimates of this regression in Table 1. In the first column, we present our results without any controls. We find that the coefficients of interest are both positive and significant, suggesting that hosts that write longer reviews and hosts that do not use the Instant Book feature have, on average, higher ratings. In the second column, we include a wide array of controls and show that the coefficients are similar to our previous estimates. These results are consistent our hypothesis that more reciprocal hosts have, on average, higher ratings.

5.2 The relationship between ratings and prices

Next, we provide evidence for Proposition 3, which states that listing prices should increase after a positive shock on ratings and decrease after a negative shock on ratings.

Estimating the causal impact of ratings on prices is difficult because unobserved changes in listing quality can simultaneously affect both prices and ratings. For example, consider a host that invests in quality (*e.g.*, upgrades the bed, or installs a new air conditioner.) At the same time, since the listing's quality has improved, the host also raises the listing's price. In this case, a regression of ratings on prices will lead us to mistakenly attribute increased prices to the increase in ratings, when it is in fact driven by unobserved (to us) changes in listing quality.

To overcome this challenge, our identification strategy exploits a unique characteristic of the Airbnb platform: a listing's average rating is only disclosed after the listing accumulates three reviews. Airbnb does not disclose individual review ratings, therefore listing quality can be inferred only by reading the review content until the listing average rating is disclosed. This provides us with a natural experiment to test hosts' reaction to the disclosure of their average rating.

Relying on rating disclosures alone is not sufficient to test the relationship between ratings and prices because unobserved changes to listing quality can happen around rating disclosure. Convincingly controlling for quality changes is difficult especially because Airbnb does not disclose the individual ratings associated with reviews. To minimize endogeneity concerns due to unobserved quality changes correlated with the timing of ratings disclose, we focus on listings whose ratings (which are a proxy for quality) up to and including the third review (when the rating disclosure occurs) are constant.

Despite the fact that listing ratings are not disclosed prior to the third review, we know that for a listing to obtain an average of 5 stars at the time of disclosure, the listing must have received three 5-stars reviews (Airbnb rounds ratings to the nearest half-star, and the only set of three ratings that results in a rounded 5 stars rating is a set of three 5-stars reviews.) Therefore, to implement our identification strategy we limit our data to the subset of listings for which the disclosed rounded average rating after three reviews is 5 stars. Moreover, since we are interested in the immediate effect of rating disclosure on prices, we only consider listing prices up to and including the third review. Thus, the treatment effect we estimate is the average difference in prices between the period when a listing has between zero and two reviews (the pre-disclosure period), and the period when listings has exactly 3 reviews (the disclosure period.)

In addition to limiting our analysis to listings with constant ratings, we further control for changes in quality using a number of time-varying listing characteristics.

The following model implements our identification strategy:

$$\log \text{Price}_{it} = \beta \text{D}_{it} + \gamma X_{it} + \alpha_i + \tau_t + \epsilon_{it}. \tag{24}$$

The dependent variable is the log of the price of listing i in year-week t. D_{it}, whose coefficient is of interest, is an indicator of whether the average rating of listing i has been disclosed at time t. In X_{it} we include a set of time-varying controls to further account for changes in listing quality. Further, we include listing (α_i) and year-week (τ_t) fixed effects, to control for unobserved time invariant listing characteristics and shocks to prices common across listings, *e.g.*, prices are higher during holiday seasons. Finally, to account for serial correlation in our dependent variable, we cluster standard errors at the listing level.

We present the results of this analysis in Table 2. In the first column, we report estimates from a minimum specification without any controls. The coefficient of interest, β, is positive and statistically significant ($p < 0.01$). Our estimates suggest that a 5-star average rating disclosure leads to a 1.6% increase in the listing price.

Next, we test the robustness of our results by including a set of time-varying observable listing characteristics that can potentially affect listing quality and thus both ratings and price. Specifically, we control the number of pictures associated with each listing, the type of cancellation policy, the number of Airbnb listings in the same ZIP code, and several other observable listing attributes that can vary over time. We report the estimates of this specification in the second column of Table 2. The coefficient of interest is positive and similar in magnitude to our previous estimate.

Overall, the empirical evidence provided in this section is consistent with the predictions made in our analytical framework: hosts with higher reciprocity weights have higher ratings. Further, listing prices respond to changes in reputation, implying that Airbnb ratings are informative. Taken together, these findings suggest that hosts who are more reciprocal can charge higher prices.

5.3 Why do casual hosts have higher ratings than professionals?

We conclude our analysis by presenting a novel empirical observation, and showing how reciprocity can help explain it. Specifically, we show that professional Airbnb hosts have lower ratings than casual hosts. Because ratings are informative about listing quality (Fradkin et al., 2017), it would appear that professional hosts offer lower quality. We argue this pattern can be explained by reciprocity.

We begin by characterizing hosts by how often they participate in the market. Specifically, we define market participation as the fraction of days over the entire year hosts list their property for rent. This quantity is intended to capture the willingness of hosts to participate in the market, and it is not the same as listing occupancy rate, which is the fraction of booked nights over available nights and is a measure of demand. While market participation is a continuous measure, for convenience we will refer to hosts with relatively high market participation as professional hosts and those with relatively low market participation as casual hosts.

We demonstrate that casual hosts have higher ratings than professional hosts by regressing the star-rating of Airbnb listings on market participation. We report this analysis in the first column of Table 3. We find that increased market participation is significantly associated with lower ratings – roughly, a unit increase in average rating predicts one percentage point decrease in market participation. In column 2, we repeat the same analysis including a large set of controls that could affect rating. For example, we include proxies for face-to-face interaction such as whether a listing is a private accommodation and whether the host has more than one property listed (Fradkin et al., 2017). We also include review counts and time trends to allow for the possibility that the ratings of professional hosts decay faster because they serve a wider selection of guests, akin to the findings of (Godes and Silva, 2012). We obtain similar results.

These results suggest that professional hosts, on average, have lower rating than casual hosts. Next, we show that this difference in ratings could be due to different reciprocity preferences. Specifically, we test whether the reciprocity proxies described in Section 5.1 are negatively correlated with market participation by estimating the following model:

$$\begin{aligned}\text{Market Participation}_i &= \beta_1 \log \text{Host-to-Guest Review Length}_i \\ &+ \beta_2 \text{ Not Instant Bookable}_i + \mathrm{X}_i\gamma + \epsilon_i,\end{aligned} \tag{25}$$

where the dependent variable is the market participation of listing i and the coefficients of interest β_1 and β_2 correspond to our two reciprocity weight proxies. In X_i we include a set of controls that could be potentially correlated with the dependent variable, while at the same time affecting the host-to-guest review length and the host decision to activate the instantly bookable feature.

We report our results in Table 4. In the first column, we present the results without including any controls. The estimated coefficients are in line with our hypothesis: longer reviews and not being instantly bookable are negatively correlated with market participation. In the second column, we incorporate a wide set of controls. In both cases, we find a negative relationship between our reciprocity proxies and market participation.

Overall, these results suggest that casual hosts are more reciprocal, which could explain their higher ratings.

6 Alternative explanations

Our analytical model of reciprocity can explain the rating patterns observed on Airbnb, including the fact that casual hosts have higher ratings than professional hosts. However, reciprocity is not the only way to explain this phenomenon. In this section, we discuss various plausible alternative models and explain why we believe that these alternative models do not easily fit the patterns we observe in our data. Specifically, we show that allowing for biased ratings, risk aversion, endogenous quality, or altruistic behavior cannot easily explain the observation that casual hosts have higher ratings than professional hosts. We discuss the intuition behind these alternative models next, and we refer the reader to Appendix B for a formal analysis.

Exogenous service quality We start by considering a simple model where service quality is private information of the host, and guests infer quality from publicly available ratings. Quality is exogenously given and fixed across transactions. We assume a monopoly host who maximizes expected profits by choosing prices. All players report their true utility in their ratings. Under this model, for casual hosts to have higher ratings than professional hosts, we need to assume that market participation and service quality are negatively correlated. While in principle this relationship could hold, in practice we should expect professional hosts to strive to maintain higher ratings than casual hosts because they are more reliant on Airbnb revenue.

Relaxing the truth-telling assumption So far we have assumed that players truthfully report ratings. However, on many peer-to-peer platforms negative ratings are under-reported (Fradkin et al., 2017; Dellarocas and Wood, 2008). We can extend our model to allow for selection in reporting ratings, and show that this extension does not help to explain the observed patterns. Under selection, guests incur a cost associated with giving low ratings, and hence, only high ratings are reported. In this variant of the simple model, to explain that casual hosts have higher ratings than professional hosts we need to assume that the former are more likely to interact with guests who face higher costs from leaving low ratings. This assumption would drastically reduce the informativeness of ratings, contrary to findings in the literature (Fradkin et al., 2017) and what we see in our data.

Relaxing the risk neutral assumption As a final attempt to use the simple model to explain the empirical facts, we relax the risk neutrality assumption and assume that the guests behavior affects the host's welfare. We still assume that the service quality is exogenously given and guests truthfully report their utility in their ratings. In this model, the risk aversion towards misconduct of the guests induces the host to increase her price and lower the transaction volume. However, since the host cannot endogenously choose the service effort, the higher ratings of casual hosts are indicative of their better service quality. This means that to explain that casual hosts have higher ratings than professionals, we still need to assume that casual hosts have exogenously higher service quality than professional hosts.

Endogenous service quality So far, we argued that under the exogenous quality assumption it is difficult to explain why professional hosts have lower ratings. Next, we assume that service quality — the effort exerted by the host in each transaction — is endogenously chosen, and can vary between transactions. In this model, guests report the host's effort in their ratings. Therefore, since future guests infer the host's service quality from prior ratings, higher ratings generate higher expected demand. Thus, ratings incentivize hosts to exert effort. Because professional hosts should rely more heavily on Airbnb revenue than casual hosts, they should also have stronger incentives to exert higher levels of effort. This, would result in higher ratings for professional hosts, contradicting our empirical observations. The models we have considered so far suggest that when reputation is the sole incentive for exerting effort, it is difficult to explain the fact that professional hosts have lower ratings.

Introducing interdependent preferences As a final attempt to explain the observed rating patterns without introducing reciprocity, we allow the host to have interdependent preferences, *i.e.*, the utility of guest i, denoted as U_i, enters into the host's utility function. This approach is similar to the one used in our analytical framework, but in this case we do not allow hosts and guests to be reciprocal. Under interdependent preferences, casual hosts have higher ratings than professional hosts if casual hosts have a higher weight on the guests' welfare than professional hosts. In other words, casual hosts would have to be more altruistic than professional hosts. While in principle this could be true, we claim that altruistic behavior on Airbnb, which requires that a host behaves well independently of a guest's behavior, is a stronger assumption than reciprocity, which requires that a host behaves well in response to a guest's good behavior.

7 Conclusion

Two salient characteristics of the sharing economy make these marketplaces unique and interesting to study. First, during a transaction, buyers and sellers are likely to interact closely. Second, because transactions are between peers rather than firms and customers, sellers have a much lower market power than they do in traditional markets. Because of these characteristics, the behavior of buyers and sellers in peer-to-peer markets can be substantially different from traditional marketplaces.

In this paper, we develop an analytical framework that, by introducing reciprocity — the tendency to increase effort in response to others' increased effort — can explain how buyers and sellers interact in these marketplaces. We show that reciprocity can improve the welfare of reciprocal peers, who can obtain higher ratings and charge higher prices. We test the key predictions of our analytical framework using data collected from Airbnb, a popular peer-to-peer rental accommodation website.

We contribute to the existing literature on peer-to-peer markets by deepening our understanding of how trust is generated in these marketplaces. Specifically, we provide four insights.

First, we show that reciprocity is likely to arise in peer-to-peer platforms such as Airbnb. This is important because reciprocity can informally regulate players' behavior and promote cooperation. Because of this, more reciprocal sellers (and buyers) are rewarded with higher ratings.

Second, our theory predicts that ratings on bilateral reputation systems are more responsive to reciprocity than unilateral reputation systems (while unilateral reputation systems are more responsive to reputation concerns than bilateral reputation systems.) This means that being reciprocal is a behavior that is rewarded more in bilateral reputation systems than unilateral ones. Thus, in contrast to the existing literature that has primarily focused on the shortcomings of bilateral reputation systems, our work highlights an important practical advantage of allowing reviews from both sides of the market. This may explain the choice of bilateral reputation systems by platforms such as Airbnb and Uber.

Third, we show that reciprocity indirectly affects equilibrium prices by affecting ratings. Because ratings are informative about seller quality, by earning higher ratings, more reciprocal sellers are able to charge higher prices.

Fourth, our findings have implications for market design. A good matching mechanism that matches reciprocal hosts and guests can induce positive reciprocity which, in turn, increases the welfare of both hosts and guests. By contrast, a poor matching mechanism may trigger negative reciprocity and decrease welfare. Our work suggests that reciprocity preferences should be considered among other matching criteria. While reciprocity is difficult to measure outside experimental settings, the reciprocity proxies we develop in this paper, such as the host-to-guest review length, offer a starting point.

Overall, our paper represents a first step towards understanding the drivers of peer behavior in the sharing economy by combining a theoretical model incorporating reciprocity with data from an established peer-to-peer marketplace.

References

Anderson, Michael, Jeremy Magruder. 2012. Learning from the crowd: Regression discontinuity estimates of the effects of an online review database. *The Economic Journal* **122**(563) 957–989.

Andreoni, James. 1990. Impure altruism and donations to public goods: A theory of warm-glow giving. *The Economic Journal* **100**(401) 464–477.

Blanchard, Simon J, Kurt A Carlson, Jamie D Hyodo. 2016. The favor request effect: Requesting a favor from consumers to seal the deal. *Journal of Consumer Research* **42**(6) 985–1001.

Charness, Gary. 2004. Attribution and reciprocity in an experimental labor market. *Journal of Labor Economics* **22**(3) 665–688.

Dellarocas, Chrysanthos, Charles A Wood. 2008. The sound of silence in online feedback: Estimating trading risks in the presence of reporting bias. *Management Science* **54**(3) 460–476.

Edelman, Benjamin. 2017. The market design and policy of online review platforms. *Oxford Review of Economic Policy* **33**(4) 635–649.

Einav, Liran, Chiara Farronato, Jonathan Levin. 2016. Peer-to-peer markets. *Annual Review of Economics* **8** 615–635.

Falk, Armin, Urs Fischbacher. 2006. A theory of reciprocity. *Games and Economic Behavior* **54**(2) 293–315.

Farronato, Chiara, Andrey Fradkin. 2018. The welfare effects of peer entry in the accommodation market: The case of Airbnb. Tech. rep., National Bureau of Economic Research.

Fehr, Ernst, Simon Gächter. 2000. Fairness and retaliation: The economics of reciprocity. *Journal of Economic Perspectives* **14**(3) 159–181.

Fradkin, Andrey, Elena Grewal, David Holtz. 2017. The determinants of online review informativeness: Evidence from field experiments on Airbnb.

Godes, David, José C Silva. 2012. Sequential and temporal dynamics of online opinion. *Marketing Science* **31**(3) 448–473.

Gouldner, Alvin W. 1960. The norm of reciprocity: A preliminary statement. *American Sociological Review* 161–178.

Horton, John, Joseph Golden. 2015. Reputation inflation: Evidence from an online labor market.

Kahneman, Daniel, Jack L Knetsch, Richard Thaler. 1986. Fairness as a constraint on profit seeking: Entitlements in the market. *The American Economic Review* 728–741.

Levine, David K. 1998. Modeling altruism and spitefulness in experiments. *Review of Economic Dynamics* **1**(3) 593–622.

Luca, Michael. 2016. Reviews, reputation, and revenue: The case of Yelp.com.

Malmendier, Ulrike, Vera L te Velde, Roberto A Weber. 2014. Rethinking reciprocity. *Annu. Rev. Econ.* **6**(1) 849–874.

Rabin, Matthew. 1993. Incorporating fairness into game theory and economics. *The American Economic Review* 1281–1302.

Resnick, Paul, Richard Zeckhauser. 2002. Trust among strangers in internet transactions: Empirical analysis of ebay's reputation system. *The Economics of the Internet and E-commerce*. Emerald Group Publishing Limited, 127–157.

Rotemberg, Julio J. 2006. Altruism, reciprocity and cooperation in the workplace. *Handbook of the economics of giving, altruism and reciprocity* **2** 1371–1407.

Seiler, Stephan, Song Yao, Georgios Zervas. 2018. Causal inference in word-of-mouth research: Methods and results.

Sobel, Joel. 2005. Interdependent preferences and reciprocity. *Journal of Economic Literature* 392–436.

Tadelis, Steven. 2016. Reputation and feedback systems in online platform markets. *Annual Review of Economics* **8**(1).

Yoganarasimhan, Hema. 2013. The value of reputation in an online freelance marketplace. *Marketing Science* **32**(6) 860–891.

Zervas, Georgios, Davide Proserpio, John Byers. 2015. A first look at online reputation on airbnb, where every stay is above average.

Zervas, Georgios, Davide Proserpio, John W Byers. 2017. The rise of the sharing economy: Estimating the impact of airbnb on the hotel industry. *Journal of Marketing Research* **54**(5) 687–705.

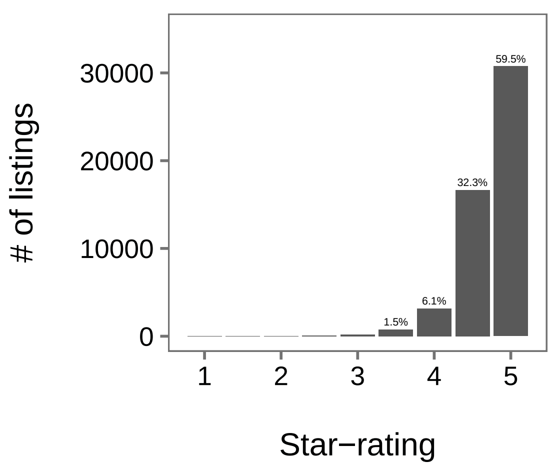

Figure 1: Listing star-rating distribution.

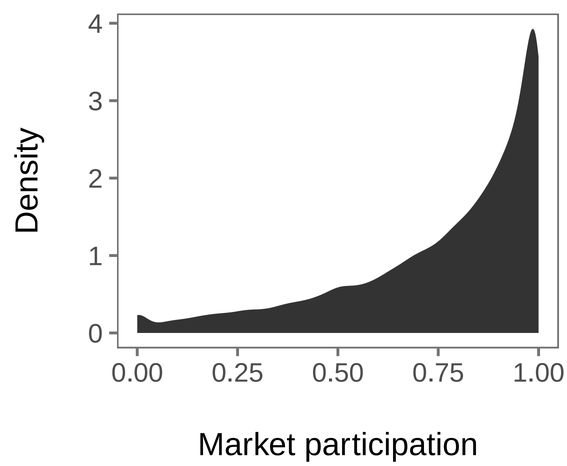

Figure 2: Probability density function of market participation.

Table 1: Correlation between host star-rating and reciprocity weight proxies.

	(1)	(2)
log Host-to-Guest Review Length	0.144***	0.123***
	(0.003)	(0.003)
Not Instant Bookable	0.030***	0.047***
	(0.004)	(0.004)
Private Accommodation		−0.006
		(0.004)
Host Has 2+ Listings		−0.019***
		(0.003)
log Price		0.048***
		(0.003)
log Reviews		0.048***
		(0.002)
Number of Bedrooms		0.007**
		(0.003)
Number of Beds		−0.003
		(0.002)
Number of Bathrooms		0.026***
		(0.003)
log Minimum Booking Nights		0.002
		(0.003)
log Number of Photos		0.033***
		(0.003)
Has Complete Description		0.029***
		(0.003)
Person Capacity		−0.007***
		(0.002)
Number of Guests Included		−0.002*
		(0.001)
log Extra Guest Price		−0.003***
		(0.001)
Is Superhost		0.217***
		(0.006)
Host Subscribed Since (Months)		0.007***
		(0.000)
N	59868	59868
R^2adj.	0.051	0.093

Note: The dependent variable is the average star-rating of host i. Robust standard errors are shown in parentheses.

Significance levels: * p<0.1, ** p<0.05, *** p<0.01.

Table 2: The impact of rating disclosure on listing price.

	(1)	(2)
Rating Disclosed	0.016***	0.016***
	(0.001)	(0.001)
log Reviews		0.002
		(0.002)
log Number of Photos		−0.008*
		(0.004)
log Extra Guest Price		−0.004**
		(0.002)
Number of Guests Included		0.009***
		(0.003)
Is Instant Bookable		−0.021***
		(0.003)
log Minimum Booking Nights		0.011*
		(0.006)
log Airbnb Zipcode Supply		0.003
		(0.006)
Cancellation Policy Dummy	No	Yes
Bed Type Dummy	No	Yes
N	352281	352281
R^2 within	0.0044	0.0090

Note: The dependent variable is the (log) price of listing i at time t. All the models include year-week and listing fixed effects. Cluster-robust standard errors (at the individual listing level) are shown in parentheses.

Significance levels: * p<0.1, ** p<0.05, *** p<0.01.

Table 3: Correlation between listing star-rating and market participation.

	(1)	(2)
Market participation	−0.081***	−0.092***
	(0.007)	(0.007)
Private Accommodation		−0.033***
		(0.005)
Host Has 2+ Listings		−0.047***
		(0.004)
log Price		0.072***
		(0.004)
log Reviews		0.045***
		(0.002)
Number of Bedrooms		0.005*
		(0.003)
Number of Beds		−0.002
		(0.002)
Number of Bathrooms		0.024***
		(0.003)
Not Instant Bookable		0.056***
		(0.004)
log Minimum Booking Nights		0.016***
		(0.003)
log Number of Photos		0.053***
		(0.003)
Has Complete Description		0.058***
		(0.003)
Person Capacity		−0.011***
		(0.002)
Number of Guests Included		0.001
		(0.001)
log Extra Guest Price		−0.001
		(0.001)
Is Superhost		0.249***
		(0.006)
Host Subscribed Since (Months)		0.005***
		(0.000)
Cancellation Policy Dummy	No	Yes
Bed Type Dummy	No	Yes
N	51697	51697
R^2adj.	0.0028	0.081

Note: The dependent variable is the average star-rating of listing j. Robust standard errors are shown in parentheses.

Significance levels: * p<0.1, ** p<0.05, *** p<0.01.

Table 4: Correlation between host market participation and reciprocity proxies.

	(1)	(2)
log Host-to-Guest Review Length	−0.006***	−0.016***
	(0.002)	(0.002)
Not Instant Bookable	−0.063***	−0.036***
	(0.003)	(0.003)
Private Accommodation		−0.084***
		(0.003)
Host Has 2+ Listings		0.047***
		(0.003)
log Price		−0.016***
		(0.003)
log Reviews		0.053***
		(0.002)
Number of Bedrooms		−0.010***
		(0.002)
Number of Beds		0.001
		(0.001)
Number of Bathrooms		0.008***
		(0.003)
log Minimum Booking Nights		−0.029***
		(0.002)
log Number of Photos		0.019***
		(0.002)
Has Complete Description		−0.003
		(0.003)
Person Capacity		0.002**
		(0.001)
Number of Guests Included		0.003***
		(0.001)
log Extra Guest Price		0.007***
		(0.001)
Is Superhost		−0.042***
		(0.006)
Host Subscribed Since (Months)		−0.004***
		(0.000)
Has Star Rating		0.195***
		(0.018)
Has Star Rating × Star Rating		−0.035***
		(0.004)
N	52966	52966
R^2adj.	0.0068	0.11

Note: The dependent variable is the market participation of host i. Robust standard errors are shown in parentheses.
Significance levels: * p<0.1, ** p<0.05, *** p<0.01.

A Proofs

A.1 Proof of proposition 1

Proposition 1. *The host's average rating on Airbnb is positively related to her reciprocity weight, i.e.,* $\forall \alpha_i \geq 0$ *and* $\alpha_i \neq 0$,

$$\frac{\partial R_{airbnb}}{\partial \alpha_h} > 0 \tag{26}$$

where

$$R_{airbnb} \equiv \int_i r_{h,i}(e_h, e_i) di \tag{27}$$

is the average rating of the host.

Proof: We solve the subgame equilibrium at Period 2. Note that at Period 2, both players choose their optimal effort level as best response to the other's effort. Then each player reports their utility of accommodation in the rating.

Plugging $r_{h,i}(e_i, e_h) = v_h + \alpha_i u(e_i, e_h)$ and $r_i(e_i, e_h) = v_i + \alpha_h u(e_i, e_h)$ into the ex-post utility of host and guest i, we have

$$U_i(e_i|e_h) = \max_{e_i}\{v_h + \alpha_i u(e_i, e_h) - P - C(e_i) + \beta_i[v_i + \alpha_h u(e_i, e_h)]\} \tag{28}$$

$$U_h(e_h|e_i) = \max_{e_h}\{v_i + \alpha_h u(e_i, e_h) + \beta_h[v_h + \alpha_i u(e_i, e_h)]\} \tag{29}$$

Combining the first order conditions of the two optimality problems, we have

$$\begin{aligned} e_i^* &= A(\alpha_h + \beta_h\alpha_i)^{\frac{1-k}{2}}(\alpha_i + \beta_i\alpha_h)^{\frac{1+k}{2}} \\ e_h^* &= B(\alpha_h + \beta_h\alpha_i)^{1-\frac{k}{2}}(\alpha_i + \beta_i\alpha_h)^{\frac{k}{2}} \end{aligned}$$

where $A \equiv (\frac{k}{c_i})^{\frac{1-k}{2}}(\frac{1-k}{c_h})^{\frac{1+k}{2}}$ and $B \equiv (\frac{k}{c_i})^{1-\frac{k}{2}}(\frac{1-k}{c_h})^{\frac{k}{2}}$.

Thus,

$$\begin{aligned} r_{h,i}(e_i, e_h) &= v_h + \alpha_i e_i^k e_h^{1-k} && (30) \\ &= v_h + \alpha_i[A(\alpha_h + \beta_h\alpha_i)^{\frac{1-k}{2}}(\alpha_i + \beta_i\alpha_h)^{\frac{1+k}{2}}]^k[B(\alpha_h + \beta_h\alpha_i)^{1-\frac{k}{2}}(\alpha_i + \beta_i\alpha_h)^{\frac{k}{2}}]^{1-k}. \end{aligned}$$

From the above, we have

$$\frac{\partial r_{h,i}}{\partial \alpha_h} > 0 \tag{31}$$

and since $R_{airbnb} \equiv \int r_{h,i} di$

$$\frac{\partial R_{airbnb}}{\partial \alpha_h} > 0. \tag{32}$$

$\square$

A.2 Proof of proposition 2

Proposition 4.2. *The ratings on both review systems depend positively on the weights of reciprocity and reputation, α_h and β_h. However, given the same pool of guests, the host's average rating on Airbnb responds more to the reciprocity weight, while the rating on the unilateral review system responds more to reputation weight,* i.e,

$$\frac{\partial R_{airbnb}}{\partial \alpha_h} > \frac{\partial R_{uni}}{\partial \alpha_h} > 0 \text{ and } \frac{\partial R_{uni}}{\partial \beta_h} > \frac{\partial R_{airbnb}}{\partial \beta_h} > 0. \tag{33}$$

where $R_{airbnb} \equiv \int_i r^{air}_{h,i}(e_h, e_i)di$ is the host's average rating on Airbnb, and $R_{uni} \equiv \int_i r^{uni}_{h,i}(e_h, e_i)di$ is the host's average rating on an unilateral review system.

Proof: From previous results, we have

$$e_i^* = A(\alpha_h + \beta_h \alpha_i)^{\frac{1-k}{2}}(\alpha_i + \beta_i \alpha_h)^{\frac{1+k}{2}}$$
$$e_h^* = B(\alpha_h + \beta_h \alpha_i)^{1-\frac{k}{2}}(\alpha_i + \beta_i \alpha_h)^{\frac{k}{2}}.$$

For the same host and the same guests pool on both the unilateral review system platform and Airbnb, and given identical values of the parameters, except $\beta_i = 0$ under the unilateral review system, we have

$$\frac{\partial r^{air}_{h,i}}{\partial \alpha_h} > \frac{\partial r^{uni}_{h,i}}{\partial \alpha_h} > 0 \tag{34}$$

$$0 < \frac{\partial r^{air}_{h,i}}{\partial \beta_h} < \frac{\partial r^{uni}_{h,i}}{\partial \beta_h}. \tag{35}$$

Since $R_{airbnb} \equiv \int_i r^{air}_{h,i}(e_h, e_i)di$ and $R_{uni} \equiv \int_i r^{uni}_{h,i}(e_h, e_i)di$, we have

$$\frac{\partial R_{airbnb}}{\partial \alpha_h} > \frac{\partial R_{uni}}{\partial \alpha_h} > 0 \text{ and } \frac{\partial R_{uni}}{\partial \beta_h} > \frac{\partial R_{airbnb}}{\partial \beta_h} > 0. \tag{36}$$

Thus, under a unilateral reputation system, hosts that care more about reputation are ranked higher than hosts that care more about the shared experience utility. The opposite is true on bilateral reputation systems like Airbnb.[10] $\square$

[10] An alternative and equivalent way to prove proposition 2 would be to remove the term $\beta_i r_i$ from the ex-post utility of guest i in equations 11 and then solve the optimization problem.

A.3 Proof of proposition 3

Proposition 3. *On Airbnb, prices increase after a positive shock on ratings, and decrease after a negative shock on ratings. Given the same price P_1 at period 1, and ratings R_{airbnb} and R'_{airbnb} disclosed at period 2, we have the following relationship for prices posted at period 3,*

$$if\ R_{airbnb} > R'_{airbnb}\ then\ P_3(R_{airbnb}) > P_3(R'_{airbnb}). \tag{37}$$

Proof: Ex-ante, guests determine whether to enter the market according to $V_i(P)$. Guest i enters the market if and only if $V_i(P) \geq 0$, given P. Therefore, the marginal guest is guest i^* where $V_{i^*}(P) = 0$. For each guest i at period 3, the expected utility $V_i(P)$ is positively determined by her inference of α_h, β_h.

From previous results, we have $\frac{\partial u(e_i,e_h)}{\partial \alpha_h} > 0$ and $\frac{\partial u(e_i,e_h)}{\partial \beta_h} > 0$. Together with $r_i = v_i + \alpha_h u(e_i, e_h)$, we have

$$\frac{\partial r_i}{\partial \alpha_h} > 0 \tag{38}$$

$$\frac{\partial r_i}{\partial \beta_h} > 0 \tag{39}$$

Then, since $U_i(e_i, e_h, r_i, v_h) = v_h + \alpha_i u(e_i, e_h) + \beta_i r_i$, we have that $\forall(\alpha_i, \beta_i) > 0$,

$$\frac{\partial U_i(e_i, e_h, r_i, v_h)}{\partial \alpha_h} > 0 \tag{40}$$

$$\frac{\partial U_i(e_i, e_h, r_i, v_h)}{\partial \beta_h} > 0, \tag{41}$$

where $u(e_i, e_h)$ is the shared experience utility and $U_i(e_i, e_h, r_i)$ is ex-post utility of guest i during the accommodation stay.

At period 3, $V_i(P_3)$ is the ex-ante utility of the guest i if she decides to transact with the host. The guest i forms her expectation based on the publicly observed price, P_3, and the average rating of the host, $R_h \equiv \int r_{h,j} dj$, *i.e.*,

$$V_i(P) = \int U_i(e_i, e_h, r_i, v_h) dF^{updated}(\Psi_h), \tag{42}$$

where $F^{updated}(\Psi_h)$ is the posterior distribution of the host's characteristic parameter vector (α_h, β_h, v_h). The Bayesian-updated guests update the prior distribution $F(\Psi_h)$ upon the signal R_h to derive $F^{updated}(\Psi_h)$.

From Equations 40, 41, and 42, we have that $\forall(\alpha_h, \beta_h, v_h) \in \Psi_h$, the expected value of $U_i(e_i, e_h, r_i, v_h)$ are positively related to α_h and β_h, i.e.,

$$\frac{\partial V_i(P_3)}{\partial \alpha_h} > 0 \tag{43}$$

$$\frac{\partial V_i(P_3)}{\partial \beta_h} > 0 \tag{44}$$

The two conditions above imply that the ex-ante utility to transact is higher if the host is perceived to have higher weight on shared experience or reputation respectively.

Meanwhile, from previous results we have

$$\frac{\partial r_{h,i}}{\partial \alpha_h} > 0 \tag{45}$$

$$\frac{\partial r_{h,i}}{\partial \beta_h} > 0 \tag{46}$$

Conditions 45 and 46 show that the hosts with higher value of α_h and β_h have higher rating $r_{h,i}$ given the same guest i. Then, from $R_h \equiv \int r_{h,i} di$, we have that, given the same guests pool, the average rating, R_h, reveals higher value of (α_h, β_h).

Suppose that two hosts are identical, except for their ratings, *i.e.*, they have the same price P_1, same location, and similar property, but $R_{airbnb} > R'_{airbnb}$. Then the higher rating R_{airbnb} is a positive signal relative to R'_{airbnb}, *i.e.*, the expected value of the host's characteristics is better for the host with R_{airbnb}. Thus, if the two hosts post identical P_3, the expected demand towards the host with higher average ratings would be higher, *i.e.*, $Q(P_3, R_{airbnb}) > Q(P_3, R'_{airbnb})$.

Since each period-3 guest is more willing to transact with a host having R_{airbnb} than with a host having R'_{airbnb}, the host with higher rating posts higher price P_3 in this monopolistic pricing setting (assuming that the hosts are faced with the same pool of guests).

Formally, let's suppose P_1 and Q_1 are the same for Host A and Host B, and, without loss of generality, assume the mass of guests enter the market have the same set of parameters, (α_i, β_i). Let the Host A have an average rating R_{airbnb} and Host B an average rating R'_{airbnb}. If $R_{airbnb} > R'_{airbnb}$, then from Equation 38 and Equation 39, we have at least one of the following conditions holds:

$$\alpha_A > \alpha_B$$

$$\beta_A > \beta_B$$

Then, for any positive P_3,

$$V_i(transact\ with\ A|P_3) > V_i(transact\ with\ B|P_3) \tag{47}$$

The marginal guest is defined as the one with $V_i(transact|P_3, R_{airbnb}) = 0$. Then, from Equation 47, we know that the marginal guest transacting with Host A can bear a higher P_3 compared to that trading with Host B, since Host A is expected to have higher value of α or β or both.

Since P_3 is determined by

$$\max_{P_3}\{P_3Q(P_3) + \int U_h(e_i, e_h, v_h^*, r_{h,i})dF(\Psi_g)\}, \tag{48}$$

we have that, given that Ψ_g is the same for Host A and B, $P_3(R_{airbnb}) > P_3(R'_{airbnb})$. □

A.4 Relaxing the truth-telling assumption

As defined in Section 3.3, we have that

$$r_{h,i} \equiv \phi(v_h + u(e_i, e_h)) \tag{49}$$

Plugging the equation above into $R_{airbnb} \equiv \int r_{h,i}di$ we have

$$R_{airbnb} = \int \phi(v_h + u(e_i, e_h))di \tag{50}$$

Now, assume all hosts are faced with an identical pool of guests. Let $\omega_i \equiv v_h + u(e_i, e_h)$. From previous results we have:

$$\frac{\partial \omega_i}{\partial \alpha_h} > 0 \tag{51}$$

Since ϕ is a weakly increasing mapping, we have

$$\frac{\partial \phi(\omega_i)}{\partial \omega_i} \geq 0 \tag{52}$$

The inequality is strict for some ω_i.

From the Equations 51 and 52, we have $\frac{\partial \phi(\omega_i)}{\partial \alpha_h} \geq 0$ and the inequality is strict for some ω_i, i.e.,

$$\frac{\partial r_{h,i}}{\partial \alpha_h} > 0 \tag{53}$$

□

B Alternative models

B.1 A Simple model

We start with a simple model where we assume the host to be a risk-neutral profit maximizer. The service quality offered by the host is exogenously given and fixed across transactions. Further, the service quality is private information of the host and only revealed to the guests during their stay in the host's property. The distribution of service quality is common knowledge. We assume a monopolistic host and a continuum of guests with heterogeneous tastes. The guests are located on a line and the host is located at the center. The heterogeneous taste of guest i is modeled as x_i, which denotes the *distance* between the host and the guest i, and it follows a uniform distribution over $[0, 1]$. Without loss of generality, we assume the effort cost of publishing a review is zero. The timing of the game is as follows:

- in Period 1, the host posts price P_1 and a continuum of guests – the early guests – enter the market. The early guests can only observe P_1. As stated above, the expected value of service quality, $E[v_h]$, is common knowledge. Guests decide whether to request accommodation. The transaction volume for this set of guests, Q_1, is realized.

- in Period 2, the accommodation stay takes place. Service quality is now revealed to the guest who requests the accommodation. The utility of guest i is $v_h - x_i$. At the end of this period, the guest i publishes a rating $r_{h,i}$ for the host.

- in Period 3, the host observes the ratings received in period 2 and post a new price P_3. A continuum of guests – the late guests – enter the market. Their heterogeneous taste parameter x_i follows the same distribution of the early guests. They observe the average rating for host h disclosed in Period 2 and the price P_3. They make their accommodation decision accordingly.

We assume that guests truthfully report their utility in the ratings, *i.e.*, $r_{h,i} = v_h - x_i$, where v_h denotes the service quality and x_i denotes the heterogeneous taste of guest i. We don't consider this truthtelling assumption is a strong one here, since under this scenario, an agent does not have incentive to collude with host in inflating ratings. Meanwhile, if agents truthfully report service quality in the ratings, informative ratings have value to other users on the platform. Thus, the truthtelling assumption is justified by the phenomenon of "warm glow" discussed in (Andreoni, 1990) [11]

[11] Note that since P_1 is common knowledge, assuming $r_{h,i} = v_h - x_i$ or $r_{h,i} = v_h - x_i - P_1$ is exactly the same for our analysis.

Moreover, we assume that the distribution of v_h and x_i is common knowledge and, hence, $E[v_h]$ is known ex-ante to the guests and $E[x_i]$ is known ex-ante to the host.

In this setting, the only choice variable of the host is price. The objective function of the host is given by:

$$V_h(P) = \max_P \{E[PQ - C(Q)|P]\} \tag{54}$$

where P denotes price and Q denotes the transaction volume. $V_h(P)$ denotes the ex-ante utility contingent on choosing price P. Since the host is risk-neutral and only interested in expected profits, the utility function coincides with expected profits.

Similarly, the only choice variable of a guest is whether to request accommodation. A guest requests accommodation from the host if and only if her expected utility from the accommodation is non-negative. The guest's ex-ante utility prior to the accommodation is given by:

$$V_i(transact|P) = E[v_h] - x_i - P \tag{55}$$

where $V_i(transact|P)$ is the ex-ante utility of the guest i if she books the accommodation. P denotes the price set by the host, v_h denotes the service quality, and x_i denotes the heterogeneous taste of guest i. we assume x_i to have uniform distribution over $[0, 1]$).

After staying in the host's property, the guest i publishes a rating $r_{h,i} = v_h - x_i$. The host's rating $R \equiv \int r_{h,i} di$ depends only on v_h and $E[v_h]$, as $E[v_h]$ determines the pool of guests requesting accommodation from the host.

In the solution below, we solve backwards for Perfect Bayesian Equilibrium.

Proof. In the first period, early guests make a decision based on $E[v_h]$ which is common knowledge. When the host chooses P_1, the marginal guest i^* who is indifferent about reserving the accommodation or not, is given by

$$E[v_h] - x_{i^*} - P_1 = 0. \tag{56}$$

Guests with $x \in [0, x_{i^*}]$ book the accommodation. Hence, the first period transaction volume is

$$Q_1 = x_{i^*} = E[v_h] - P_1. \tag{57}$$

Now let's consider the rating a host receives. The average rating, denoted by R, is given by

$$R \equiv \frac{\int_0^{Q_1} v_h - x_i dx_i}{Q_1} = v_h - \frac{Q_1}{2}. \tag{58}$$

Since both R and Q_1 are common knowledge in Period 3, late guests can infer the value of v_h from $v_h = R + \frac{Q_1}{2}$. Therefore, the marginal guest in Period 3 is given by $v_h - x_{j^*} - P_3 = 0$, *i.e.*,

$$Q_3 = x_{j^*} = v_h - P_3. \tag{59}$$

Hence, the host maximizes profits in Period 3 by solving the following optimality problem:

$$\max_{P_3} P_3 Q_3 = \max_{P_3}\{P_3(v_h - P_3)\}. \tag{60}$$

From the first order condition (FOC), we have $P_3^* = \frac{v_h}{2}$. Then in Period 1, assuming hosts discount future revenue at rate β, a host solves:

$$\max_{P_1}\{P_1(E[v_h] - P_1) + \beta P_3^*(v_h - P_3^*)\}. \tag{61}$$

From the FOC, we have

$$P_1 = \frac{E[v_h]}{2} \tag{62}$$

Then from Equations 57 and 62, we have $Q_1 = \frac{E[v_h]}{2}$. Thus, we have

$$R = v_h - \frac{Q_1}{2} = v_h - \frac{E[v_h]}{4}. \tag{63}$$

For a pool of hosts with $v_h \in [\underline{v}_h, \bar{v}_h]$, the average value of v_h is $E[v_h]$, hence, the average ratings of hosts is $M = \int R_j dj = \frac{3}{4}E[v_h]$.

Let R_c denote the rating casual hosts and R_p denote the rating of professional hosts. For R_c to be systematically higher than R_p, we have to assume that the average service quality of casual hosts is higher than that of professional hosts, *i.e.*, $E[v_h^c] > E(v_h^p)$, where $E[v_h^c]$ and $E[v_h^p]$ denote the average service quality of casual hosts and of professional hosts, respectively. Note that casual and professional hosts differ in their market participation frequency. In this simple model, the service quality of the host is not endogeneously chosen by the host, hence we lack a mechanism to link a host's market participation frequency with their ratings. Thus, in order for $E[v_h^c] > E[v_h^p]$, we need to assume that market participation negatively correlates with the exogenously given service quality. However, this does not seem to be a natural assumption to make. As higher service quality can translate into higher ratings which attracts future business, hosts choosing to participate more frequently should not have less incentive to work hard towards achieving high ratings. If anything, this rationale suggests that market participation frequency should be positively correlated with service quality. We conclude that this simple model cannot easily explain why professional hosts have lower ratings. □

B.2 Relaxing the truth-telling assumption

Next, motivated by the observation that the guests who do not report ratings are likely to have had a worse experience (Fradkin et al., 2017), we relax the *truth-telling* assumption. In doing so, we allow selection bias in ratings. To allow for selection bias, we assume that guests provide a rating to a host only if the rating is above a threshold θ_i, *i.e.*,

$$\mathbf{1}\{rating\} = 1 \text{ iff } r_{h,i} > \theta_i \tag{64}$$

Then the ex-ante utility of the guest i in period 1 is given by

$$V_i(transact|P) = E[v_h] - x_i - P + E[r_i - \mathbf{1}\{rating\}\alpha_i|(v_h - x_i) - r_{h,i}|], \tag{65}$$

where $r_{h,i}$ denotes the rating guest i gives to the host, and r_i denotes the rating the host gives to the guest i. The term α_i denotes the weight of the host's reputation in guest i's utility.

Compared with the guest's ex-post utility in the simple model (Equation 55), the current utility function has two new parts. The first part includes $\mathbf{1}\{rating\}|(v_h - x_i) - r_{h,i}|$ and the rating threshold θ_i. The difference between the disclosed rating and the true level of the guest's utility captures the guest's disutility derived from reporting a rating different from the true value of service quality and deteriorating rating informativeness. The term θ_i captures the cost (*e.g.*, psychological cost) for a guests to give a low rating. Together, these terms capture the trade-off faced by the guest when choosing whether to rate a host and what rating to disclose. While psychological costs may encourage guests to inflate ratings, guests also have an incentive to provide informative ratings as a contribution to other users on the platform. The two forces work in opposite direction and together determine the rating guests report.

The second part is r_i, the rating guest i receives from the host, which captures the fact that the guest i has reputation concerns. The reason for r_i to be part of the guest's utility is to match the Airbnb setting, where a host can rate the guest, and this rating affects whether future hosts will accept the guest's accommodation request.

In all our analyses, we only consider ratings produced under the new simultaneous Airbnb reputation mechanism. Therefore, strategic rating manipulation of ratings is not a concern, *i.e.*, hosts cannot strategically collude with guests to exchange high ratings which implies that r_i is not a function of $r_{h,i}$.

In the equilibrium of this model, a host receives a higher rating than in the simple model, and the inflated ratings depend on the average level of θ_i associated with the pool of the guests. While θ_i may be affected by the interaction between guests and hosts, it seems unlikely that it is also correlated with market participation, since this information is not revealed to guests. Unless we are willing to assumme such a correlation, we cannot easily explain the difference in ratings observed in Section 5.3. The formal proof of this statement is as follows.

Proof. Under this model, the only choice variable of the guest during Period 2 is $r_{h,i}$. If guest i decides to request the accommodation, her ex-ante utility in Period 1 is given by

$$\begin{aligned} V_i(transact|P) &= \max_{r_{h,i}} E[v_h] - x_i - P + E[r_i - \alpha_i \mathbf{1}\{rating\}|r_{h,i} - (v_h - x_i)|] \\ \text{where} \quad & \mathbf{1}\{rating\} = 1 \; iff \; r_{h,i} > \theta_i. \end{aligned} \tag{66}$$

The ex-post utility of guest i at the end of Period 2 is given by

$$\begin{aligned} u_i(r_{h,i}) &= \max_{r_{h,i}} \{v_h - x_i + r_i - \alpha_i \mathbf{1}\{rating\}|r_{h,i} - (v_h - x_i)|\} \\ \text{where} \quad & \mathbf{1}\{rating\} = 1 \; iff \; r_{h,i} > \theta_i. \end{aligned} \tag{67}$$

With respect to the best response of guest i, two options exist:

1. If $r_{h,i} \geq \theta_i$, then:

$$V_i(transact|P) = E[v_h] - x_i - P + E[r_i - \alpha_i|r_{h,i} - (v_h - x_i)|] \tag{68}$$

$$u_i(r_{h,i}) = [v_h - x_i + r_i - \alpha_i|r_{h,i} - (v_h - x_i)|] \tag{69}$$

   where $u_i(r_{h,i})$ is the ex-post utility after the accommodation stay. Then from Equation 69 and $\alpha_i > 0$, we have:

$$r_{h,i} = v_h - x_i. \tag{70}$$

2. If $r_{h,i} < \theta_i$, then, the guest i does not publish a rating.

Then the ex-ante utility for a guest to enter the market is given by:

$$\begin{aligned} V_i(transact|P) &= E[v_h] - x_i - P + E[r_i] + \alpha_i[Prob(v_h - x_i > \theta_i) * [v_h - x_i - (v_h - x_i)] \\ &+ (1 - Prob(v_h - x_i > \theta_i)) * 0] \end{aligned} \tag{71}$$

Note that under simultaneous ratings, guest i cannot directly influence r_i by choosing the rating she gives to the host. Thus, the only choice of guest i in Period 1 is to transact if and only if $E[v_h]-x_i-P \geq 0$. Therefore, the following conditions still hold: $Q_1 = x_i^* = E[v_h]-P_1$ and $P_1 = \frac{E[v_h]}{2}$.

Then, the average rating of a host under this model, denoted as R^{biased}, is:

$$\begin{aligned} R^{biased} &\equiv \frac{\int_{i\,rate} r_{h,i}di}{\int_{i\,rate} 1di} = \frac{\int_{i\,rate}(v_h - x_i)di}{\int_{i\,rate} 1di} \\ &= v_h - \frac{\int_{i\,rate} x_i df(x_i)}{\int_{i\,rate} di} \\ &= v_h - \frac{\int_0^{x^{**}} x_i df(x_i)}{\int_0^{x^{**}} df(x_i)} \end{aligned}$$

where $x^{**} \equiv \min\{v_h - \theta_i, E[v_h] - P_1\}$, and f is the density function of x_i.

Then, denoting the average rating in the simple model as $R^{unbiased} \equiv \frac{\int_i r_{h,i}di}{\int_i idi}$, we have that:

$$R^{biased} - R^{unbiased} = g(\theta) > 0 \tag{72}$$

i.e., the average rating under selection bias is higher than the average rating without selection bias, and their difference is a function of θ. Under this condition, to observe systematically lower ratings for professional hosts we need to assume that v_h^p is systematically lower than v_h^c, or that the distribution of v_h^c first order stochastically dominates (FOSD) the distribution of v_h^p.

Alternatively, to explain the rating patterns observed, we could assume that guests of professional and casual hosts differ systematically in their θ, the psychological cost of leaving a bad review. Specifically, guests of professional hosts must have a systematically lower psychological threshold than guests of casual hosts. However, guests cannot observe hosts market participation and, thus, they cannot discern between professional or casual hosts; and hosts cannot infer the guests' θ, so they cannot select a specific type of guest. Thus, this self-selection of guests depending on the host type (and vice versa) is unlikely to occur in practice.[12] □

B.3 Relaxing the risk neutral assumption

In this section, we relax the risk-neutral assumption, and allow the behavior of the guests to enter into the utility function of the host. We propose this modification because of the nature of Airbnb transactions. Since Airbnb hosts share their own properties, it is natural

[12]Note that even under the assumption that hosts are able to infer guests' θ, all hosts, and in particular professional hosts, who have a higher weight on reputation, would select those guests with higher θ in order to reduce the probability of receiving a lower rating.

to expect them to be risk-averse towards guest misconduct.

Formally speaking, we assume the effort of the guest i, denoted as e_i, to enter into the utility of the host, and the host's utility is assumed to be concave with respect to e_i. The utility function of the host is given by:

$$V_h(P) = \max_P \{E[PQ - C(Q)|P] + E[(r_{h,i} + u_h(e_i))|P]\} \tag{73}$$

where, as before, P denotes price, Q denotes the transaction volume, and $v_h(P)$ denotes the ex-ante utility contingent on choosing price P. $r_{h,i}$ is the rating the host receives from the guest i. The term $u_h(e_i)$ shows that the guest's effort e_i affects the host's welfare. The concavity of u_h captures the risk aversion of the host. Then, hosts trade off expected profits against the possibility of guest misconduct in accepting guests. While this assumption reduces the number of guests a host will accept, the host's rating do not affect service quality since it is still exogenously given and fixed. Formally, the only choice variable is still P and the optimality problem of the host is reduced to:

$$\max_P \{E[PQ - C(Q)|P] + E[r_{h,i}]\}. \tag{74}$$

Independently of which assumptions are invoked on $r_{h,i}$, the absence of strategic manipulation of $r_{h,i}$ makes it impossible to alter the optimality problem, which therefore is analogue to the simple model discussed above.

Moreover, since service quality is exogenously given in this scenario, to explain the systematically lower ratings to professional hosts we would need to assume that the service quality of casual and professional hosts follow different distributions, and that the distribution of v_h^c FOSD the distribution of v_h^p.

B.4 Endogenous service quality

Next, we relax the exogenous service quality assumption. We allow service quality to vary between transactions. This assumption is consistent with the high heterogeneity of service quality on Airbnb. In this scenario, we consider three alternative models.

B.4.1 Model I: Hosts only care about profit and reputation

First, we endogenize service quality without changing the assumption that hosts care only about profit and reputation. The optimality problem of the host in Period 1 is given by:

$$U_{h,1}(P) = \max_P \{E[PQ - C(Q)|P] + E[U_{h,2}(e_h, r_{h,i})]\}. \tag{75}$$

Demand is realized in Period 1, while only reputation concerns and effort cost determine effort levels in Period 2. Then, the host's optimality problem in Period 2 is:

$$U_{h,2}(e_h, r_{h,i}) = \max_{e_h}\{\beta_h r_{h,i} - C_h(e_h)\}. \tag{76}$$

Because of their higher market participation, professional hosts have a higher weight on reputation concerns, *i.e.*, professional hosts have systematically higher β_h. With the assumption that guests report hosts' effort levels in ratings $r_{h,i}$, *i.e.*, $r_{h,i} = e_h$.

From the FOC, we have

$$C_h'(e_h)^* = \beta_h. \tag{77}$$

The above condition solves for the optimal level of e_h^*.

If the effort cost function is identical across hosts, *i.e.*, the function $C_p(e_h) = C_c(e_h) \equiv C(e_h)$ for all e_h, and effort costs increase with effort level, *i.e.*, $C'(e_h) > 0, \forall e_h$, then professional hosts exert higher levels effort due to their higher β_h, *i.e.*, from $\beta_p > \beta_c$ and $C_h' > 0$, we have $e_p^* > e_c^*$.

Even if the effort cost function of professional hosts is systematically different from that of casual hosts, because of economies of scale, it is unlikely that the latter have systematically lower marginal effort cost than professional hosts.That is, $C'(e_p) \leq C'(e_c)$ for each e_h. Even if we assume $C'(e_p) > C'(e_c)$ for some e_h, the difference in marginal effort cost has to be large enough to offset the effect of β_h so that casual hosts can exert systematically higher effort under this model. Therefore, this model cannot easily explain why professional hosts have lower ratings.

B.4.2 Model II: Hosts also care about guest behavior

In this model hosts also care about guest behavior. Hosts on Airbnb are relatively "small" service providers compared to firms in the accommodation industry, such as Hilton or Marriott. Therefore, a natural assumption is that hosts are risk- averse with respect to possible guest misconduct. We model this by letting guest conduct enters hosts' utility function. However, as we show, simply introducing guest conduct in the utility function of a risk-averse host does not suffice to explain why professional hosts have lower ratings than casual hosts. Formally speaking, if we assume e_i and e_h are separable in the host's utility, guest behavior cannot alter the host's effort. Thus, we have to further assume a non-separable function of e_h and e_i as we do in our main model.

As U_h takes a seperable functional form of e_i and e_h, without loss of generality we assume

U_h is given by:

$$U_h(e_h, r_{h,i}) = \max_{e_h} g(e_i) - C_h(e_h) + \beta_h r_{h,i} \tag{78}$$

where $g(e_i)$ captures the effect of guests' conduct on the host's utility, and the concavity of $g(\cdot)$ models the host's risk averse attitude towards the guest misconduct.

Assuming $r_{h,i} = e_h$, we have the same FOC as the previous model, *i.e.*, $C'_h(e_h)^* = \beta_h$. This means that e_i does not determine the optimal level of e_h and therefore it cannot explain the difference in ratings between casual and professional hosts.

B.4.3 Model III: Hosts have interdependent preference

Next, we assume that the host has interdependent preference (without assuming the reciprocity feature discussed in our main theoretical framework), *i.e.*, the utility of guest i, denoted as U_i, enters into the host's utility function.

Under this assumption, the optimality problem of the host is given by:

$$U_h(e_h, e_i, r_{h,i}) = \max_{e_h} g(e_i) + \alpha_h U_i(e_i) - C_h(e_h) + \beta_h r_{h,i} \tag{79}$$

where U_i is the ex-post utility of guest i given by $U_i(e_i, e_h, r_i) = f(e_h) - C_i(e_i) + \beta_i r_i$. Also, e_i and e_h are the efforts of guest i and the host h, respectively. The function $g(\cdot)$ denotes how much hosts care about guest's behavior, and $f(\cdot)$ denotes guest utility derived from host's effort.

To derive a closed-from solution of the optimal effort level, we assume $f(e_h) = e_h^{1-k}$ and $g(e_i) = e_i^k$, and the effort cost function to be quadratic, *i.e.*, $C(e_h) = \frac{c_h}{2} e_h^2$.

After invoking $r_i = g(e_i)$, $r_{h,i} = f(e_h)$, we have:

$$(\alpha_h + \beta_h) f'(e_h) = C'_h(e_h) \tag{80}$$

Then, we have $e_h^* = (\frac{(1-k)(\alpha_h+\beta_h)}{c_h})^{\frac{1}{k+1}}$ and $\frac{\partial e_h^*}{\partial \beta_h} > 0$.

If casual hosts differ from professional hosts only in β_h, and professional hosts have a higher level of β_h, we have $e_p^* > e_c^*$. Therefore, $r_{h,i}^c < r_{h,i}^p$, $\forall i$. Thus, the average rating of professional hosts is expected to be higher than casual hosts, contradicting the what we observe in the Airbnb data.

In order to explain the systematically lower ratings of professional hosts, we need an additional assumption. Specifically, we need to assume that casual hosts are systematically more altruistic than professionals. This assumption can be implemented by adding the parameters α^g in front of the interdependent utility term in the gross utility of the host, and assuming that the difference on the altruism weights are larger than the difference between reputation weights. Formally stated:

$$\alpha_c - \alpha_p > \beta_p - \beta_c > 0,$$

where α_c and α_p are the weights on guests' utility of the non-professional hosts and the professional hosts, respectively.

By invoking this assumption, we allow casual hosts to have systematically higher intrinsic altruism, and the difference between altruistic attitude is large enough to offset the difference between reputation concerns. However, we consider altruistic behavior to be a larger departure from standard assumptions, and our explanation based on reciprocity to be more natural for the setting of Airbnb.