Pricing in markets without money

Theory and evidence from home exchanges

Julius Goedde*

October 17, 2024

click here for the latest version

Comments welcome.

Abstract

I examine how to set prices with internal currencies instead of real money. Users earn currency by supplying on a platform and can only spend it by consuming on the platform. A stylized model of an exchange economy with endogenous production shows that reducing the price of attractive goods below market-clearing levels can increase their supply and user welfare. The key insight is that individuals can become satiated with currency because they typically only have a limited demand for the specific goods on the platform. This creates large income effects on supply. I confirm key predictions of the model on a widely used platform for exchanging personal residences for holidays. Combining proprietary data on the universe of transactions with quasi-experimental designs, I demonstrate large income effects and show that a price-compressing reform increased the supply of attractive homes. Remarkably, participation does not decrease. Further evidence suggests that many hosts have a strong preference against for-money rental platforms. Broader insights are that lessons from traditional markets may not easily extend to markets without real money and that the latter may have advantages even in contexts where monetary transactions are commonplace.

Keywords: Market design, digital platforms, scrip systems, sharing economy, short-term rentals, dynamic matching
JEL-Codes: D47, L11, D61

---

*Paris School of Economics (PSE). Email: julius.goedde@psemail.eu. This work has benefited from the support of the Agence Nationale de la Recherche through the program Investissements d'Avenir ANR-17-EURE-0001 and support by the French Government's Investment for the Future Program (PIA)operated by ADEME. The views expressed in this paper are those of the author alone and might not represent the views of ADEME or any other institution. I am grateful to Olivier Tercieux and Liam Wren-Lewis for their excellent guidance, to Yeon-Koo Che and Julien Grenet for their continued advice, to the platform for their willingness to share the data and to Gabrielle Fack and Liam Wren-Lewis for their help in setting up a data-sharing agreement. For helpful comments I thank Sam Altmann, Nick Arnosti, Nicolas Astier, Luc Behaghel, Jean-Marc Bonniseau, Eric Budish, Julien Combe, Laura Doval, Ben Edelman, Glenn Ellison, Gabrielle Fack, Andrey Fradkin, Matt Gordon, Guillaume Haeringer, John Horton, Bruno Jullien, Süleyman Kerimov, Jacob Leshno, Shengwu Li, David Margolis, Eric Maskin, Olivier Meslin, Sébastian Otero, Ricardo Pommer Munoz, Sandra Poncet, Imke Reimers, Bernard Salanié, Nikhil Vellodi and participants of seminars and workshops at PSE and Columbia as well as the CIRM Matchings to Markets conference, EWMES 2023, Exeter MMIW, IIOC 2024, EARIE 2024, Oxford Match-UP and the American UEA meeting 2024. I am also indebted to the numerous people with whom I have had conversations that have informed my thinking about the platform.

1 Introduction

In markets where money is used as a medium of exchange, market-clearing prices are a natural way to allocate goods. Market design studies settings where money is deemed "repugnant" by policymakers or the public (Roth, 2007). A popular solution is using internal currency systems with token money that has no value outside of the system. Previous research has designed internal currency systems with market-clearing prices and has established desirable efficiency and fairness properties of such systems.[1] Yet, many internal currency systems do not set market-clearing prices.[2] Are there good reasons for doing so?

Consider the example of a widely-used platform where members exchange their primary residences during holidays with internal currency. Some homes are more attractive than others. Consider a large, well-equipped apartment in New York City and a less comfortable two-room flat in rural Germany. Suppose the platform set market-clearing internal currency prices across those different homes. Prices of very attractive homes would need to be very high so that no more people want to go there than homes are available, while prices of unpopular homes would need to be very low so that all available homes get filled. As users with popular homes might travel to all kinds of places they spend less per holiday than they earn. However, most users only have time for one or two holidays a year. After some time users with popular homes would have little incentive to host as they still hold enough currency for their next visits and would thus be satiated.

This paper shows theoretically and empirically that reducing the price of attractive assets below market-clearing levels can increase their supply. The key insight is that compressed prices have large income effects and limit the purchasing power that owners of highly demanded assets gain per unit of supply so that they do not become satiated easily. The theoretical results are confirmed using proprietary data from a real-world

[1] Prominent examples of internal currency systems designed with the help of researchers are Feeding America's auction system (Prendergast, 2017, 2022) and course allocation systems in business schools, such as CourseMatch, which determines approximately market-clearing prices (Budish, 2011; Budish et al., 2017). Pseudo-markets have also been proposed for various other contexts, including matching roommates, assigning doctors to hospitals (Echenique et al., 2021) and trading organ donations among hospitals (Agarwal et al., 2019; Ashlagi and Roth, 2021). The idea of using a (competitive) pseudo-market to allocate a fixed set of indivisible objects was pioneered by (Hylland and Zeckhauser, 1979). Miralles and Pycia (2021) prove versions of the first and second welfare theorem in a model with exogenous supply.

[2] For instance, such systems are used to exchange residences during holidays, to trade used goods and to exchange services and favors among local residents. Table D.12 provides an overview of price-setting in various internal currency systems.

internal currency system that enables more than 100,000 users each year to exchange personal residences in 150 countries. The paper also suggests that many peer suppliers have a strong preference against hosting on for-money rental platforms dominated by professional sellers. Even accounting for taxes and fees many users would financially be substantially better off on a for-money rental platform (Airbnb) but continue using the exchange platform. Hence, market designs without real money seem to have advantages even when monetary transfers are not repugnant.

When designing internal currency systems, a natural instinct is to replicate real-market economies and let internal currency prices equilibrate supply and demand as much as possible. Hylland and Zeckhauser (1979) have pioneered the idea of using a pseudo-market to allocate a fix set of indivisible objects, such as positions, to individuals. Recently, economists have helped design real markets with internal currencies. Budish (2011) and Budish et al. (2017) develop a pseudo-market for MBA class seats where internal currency prices minimize the market-clearing error. Thanks to Prendergast (2022) and his collaborators, Feeding America's regional food banks can bid with internal currency for the type of food they need most. Both systems closely mimic market mechanisms and allow as much price dispersion as necessary to (approximately) clear the market. They seem to work quite well in their settings, where most or all of agents' budgets and the supply of goods are exogenous.[3] It is unclear, however, how well market-like mechanisms do in more complex settings where budgets are endogenous to agents' supply.

This paper makes a theoretical and several empirical contributions. First, I demonstrate with a stylized model of a two-period exchange economy with endogenous production that price compression can increase the supply of popular goods, total trade volume and aggregate welfare. There are two types of agents with equal masses: those that own an H-good and those that own an L-good. All types prefer H goods. In each period, individuals can try to consume and try to supply. They maximize their lifetime utility from consuming minus the disutility from supplying subject to a budget constraint with internal currency but do not derive utility from currency itself. I assume a common disutility of hosting, that all homes are acceptable and that preferences are homoge-

[3] In the course allocation case, seats are supplied by the university and students are unable to earn any currency. Students receive equal budgets with a small random noise to ensure the existence of almost-market-clearing equilibria. Regional food banks can earn currency by offering donations they receive from local donors in the national market. In practice, however, the majority of traded food is centrally supplied by Feeding America and most of food banks' currency budget is obtained for free on the basis of local poverty indices (Prendergast, 2022). Furthermore, it seems that food banks are very farsighted and do not easily become satiated with currency (Altmann, 2022). Both systems increased user satisfaction relative to previous mechanisms.

neous within type. I then fully characterize the equilibria of two pricing-strategies the platform can follow: a competitive equilibrium with market-clearing prices and a uniform rationing equilibrium where all goods are equally expensive and excess demand is rationed uniformly at random. I distinguish three regimes. Rationing yields lower utilitarian welfare than the unique competitive equilibrium allocation if and only if agents' preferences are supermodular, that is, H-agents have a stronger preference for H-homes than the L-agents. Under submodular preferences rationing increases welfare. Submodular preferences may be plausible when viewing home types as more or less attractive locations. Third, rationing increases the supply of H homes and welfare if and only if submodularity is sufficiently strong. In the unique competitive equilibrium in this case agents with H goods supply only once but earn enough currency to consume an L-good for two periods. By contrast, in the rationing equilibrium all types supply twice. Intuitively, higher prices exert a strong income effect on users with H goods but create little substitution effect because maximum consumption is limited. Therefore, increasing the price of H reduces the supply of H-goods. A useful corollary of the main result is that if compressed prices increase supply, they also increase welfare.

The second contribution is to confirm the two key predictions of the model empirically on a large platform. First, I estimate that a given user hosts four times less in quarters when she is currency-rich. I also show that she is far less likely to accept a request with identical characteristics. A natural experiment exploiting unexpected wealth increases confirms that the strong income effects are causal. Second, I examine a reform of the pricing algorithm that increased price compression with the stated purpose of increasing trade. I show in a difference-in-difference design that attractive homes increased their supply after the reform reduced their price. Without the large satiation-induced income effects we would expect that lower prices reduce supply. Moreover, simple reduced form evidence suggests that preferences are consistent with submodularity, the theoretical condition for a supply effect. Remarkably, the reform did not reduce participation on the platform. I find little evidence of reduced effort or increased misallocation either. This suggests little side-effects of compressing prices.

The third contribution of this paper is providing evidence that systems without real money can have advantages even in settings where real money is commonly used. Much market design research has focused on settings where monetary transfers seem repugnant or are prohibited (Roth, 2007). I find evidence that many individuals have a strong preference against supplying their personal residence platforms with real money and prefer a platform without. Two complementary approaches show that the prices on the exchange platform are far from market-clearing and much more compressed than on Airbnb. Consequently, many users with attractive homes would face much better terms of trade on Airbnb. I estimate a lower bound of exchange users' willingness to forgo Airbnb income based on a revealed preference argument. I impute Airbnb prices and assume that users could have supplied and consumed on Airbnb what they supplied and consumed on the platform. This analysis indicates that many users are willing to forgo hundreds of dollars of net income that they could have earned with the same host-visit bundles – even after deducting taxes and Airbnb fees. Last but not least, I find that the absolute number of exchange users is surprisingly high. Despite Airbnb's far greater brand recognition, cross-side network effects and monetary incentives, I find some evidence that more family residences are listed on the exchange platform than on Airbnb.

This suggests that for many users the costs of offering personal residences on Airbnb are substantially higher than on the exchange platform. Anecdotal evidence indicates that Airbnb guests are perceived as less respectful and potentially having expectations of a more professional, hotel-like experience. In line with different guest expectations, I find that up to 90% of Airbnb transactions are not with plausibly personal residences. These results suggest that the platform is only competing with for-money rental platforms to a limited extent and consequently wields some pricing power. It can use this power to compress prices without materially reducing participation and thus increase supply.

Finally, these findings also contribute to our understanding of the sharing economy. Previous research has pointed out that non- or semi-professional sellers may incur higher bring-to-market costs than traditional firms (Filippas et al., 2020; Farronato and Fradkin, 2022). To the best of my knowledge, the paper is the first to provide evidence that non-professional sellers may incur substantially lower costs on peer-to-peer platforms excluding professionals than on more commercialized platforms. The results also suggest that platforms without real money are more likely to attract "idle" assets than monetized platforms and can successfully operate at large scale. This has potential implications for debates about regulating supposed "peer-to-peer" platforms.[4]

[4]Numerous papers study the impacts of Airbnb on a variety of outcomes. For instance, Farronato and Fradkin (2022) estimate large benefits for tourists as Airbnb expands supply in peak season and keeps hotel prices low. On the other hand, several papers find that Airbnb reduces the supply of long-term rental units and drives up rents (Garcia-López et al., 2020; Koster et al., 2021; Barron et al., 2021; Duso et al., 2022; Ellison and Ellison, 2024). Among others, Almagro and Domínguez-Iino (2024) and Filippas and Horton (2023) highlight potential neighborhood externalities of Airbnb. Several papers discuss

The rest of the paper is organized as follows. Section 1.1 summarizes additional related literature. Section 2 derives the theoretical results. Section 3 describes the setting and the data. Section 4 tests the first theoretical prediction and demonstrates strong income effects. Section 5 tests the second theoretical prediction - showing that lower prices increase supply. It also finds that they do not reduce participation. Section 6 further investigates why hosts continue using the platform after their price drops and presents evidence of limited competition with Airbnb.

1.1 Further literature

This paper contributes to a variety of literatures, especially on market design, on digital platforms and the sharing economy. In the market design literature, various papers have studied dynamic assignment mechanisms without real money theoretically, often under quite restrictive assumptions on agents' preferences (e.g. Jackson and Sonnenschein (2007); Kurino (2014); Nicolò et al. (2023)). Other papers have studied equilibrium existence (Baldwin et al., 2022; Jagadeesan and Teytelboym, 2022) and established welfare theorems for pseudomarkets with indivisible goods (Miralles and Pycia, 2021), focusing on allocation problems with exogenous supply. Furthermore, the welfare theorems consider Pareto-efficiency and not utilitarian welfare.

Besides Budish (2011) and Prendergast (2022), several recent papers study and propose internal currency systems to be used in practice. For instance, Kornbluth and Kushnir (2021) present a pseudo-market for undergraduate classes with priorities - again with market-clearing prices and exogenous budgets. Agarwal et al. (2019) and Ashlagi and Roth (2021) consider currency systems to trade donor organs among hospitals. Transplant authorities have only recently started to implement currency-like mechanisms and it is too early to assess their success. The famous Capitol Hill baby-sitting co-op offers a cautionary tale about how vulnerable internal currencies may be to satiation. When the co-op injected a large amount of token-money many parents felt they held enough currency for the foreseeable future and stopped baby-sitting (Sweeney and Sweeney, 1977; Krugman, 1994). While a vivid example to teach monetary policy, this story does not

to what extent Airbnb should still be considered a peer-to-peer rather than a business-to-consumer platform (Gyódi, 2019; Quattrone et al., 2020; Demir and Emekli, 2021). This paper adds to the latter by document that a money-less sharing platform, which does not create incentives to own assets primarily for rental, holds a market share in the same order of magnitude as Airbnb in the market for personal residences. It also adds by studying the prevalence of professional hosts in a much larger range of cities, more recent data and by pointing out that the share of personal residences is even smaller when looking at transaction volume (proxied by ratings) rather than the number of listings.

tell us much about how the platform's pricing strategy affects satiation risk and supply. One of the rare papers that studies the wealth evolution in token-based exchange economies is Ashlagi et al. (2024) (see also Ashlagi and Kerimov (2020)). The authors show theoretically that agents' wealth distribution does not diverge when the market is not extremely thin but do not explicitly model agents' supply choices and do not consider different pricing approaches as I do. Furthermore, I find on a very large platform that the strong income effects in token systems can make even moderate levels of wealth accumulation problematic.

The conclusion that non-market prices can increase welfare resembles recent results in mechanism design theory (Che et al. (2013); Dworczak et al. (2021), see also Weitzman (1977)). The channel in my theoretical analysis is different, however, because supply is fixed or increasing in prices in their models (where buyers do not supply) while supply is *decreasing* in prices in my model due to satiation-induced income effects. I complement (Akbarpour et al., 2022) by analyzing a setting with endogenous multi-unit supply and demand and present empirical evidence from a large real-world market with internal currency. The idea of a downward-sloping supply curve is also present in a small literature on child labor, pointing out that a ban may reduce wages and therefore - perversely - increase supply (Basu and Van, 1998; Bharadwaj et al., 2020). Thus, the phenomenon observed on the home exchange platform connects to an old idea in economics, but in a very different setting and with very different welfare implications.

Finally, a nascent literature investigates the benefits and drawbacks of centralized vs. decentralized platform pricing. Several empirical studies of services examine the value of sellers' price-setting choice for eliciting private information (Filippas et al., 2022; Gaineddenova, 2022). By contrast, I analyze if (largely) centralized prices should be market-clearing or compressed. However, a natural prediction is that decentralized prices would approach market-clearing prices and thus risk quickly satiating owners of popular goods. Hence, my findings suggest a possible new rationale for limiting sellers' price setting power, namely to avoid satiation. I also broaden the scope of platforms studied as existing papers have focused almost exclusively on commercial ride-hailing platforms.

2 Model

This section demonstrates with a simple model of an exchange economy with endogenous production that compressed prices can increase the supply of popular goods, trade volume and utilitarian welfare. I fully characterize the relationship between agents' preferences and those outcomes in competitive equilibrium and a rationing equilibrium with excess demand. I distinguish three regimes. In the first regime rationing reduces utilitarian welfare, in the second it increases it and in the third, rationing increases both utilitarian welfare, supply and total trade. In this regime, "rich" agents supply only 1 out of two periods in the competitive equilibrium while they supply both periods in a rationing equilibrium with compressed prices and excess demand.

2.1 Setup

We study a two-period exchange economy with two types of indivisible objects (homes). The set of homes types is $\mathcal{N} = \{H, L\}$. There is a continuum of agents. We assume that each agent owns one home and that a unit mass of agents owns each type of home. Formally, agents are represented as points on the interval $\mathcal{I} = [0, 2]$. For simplicity, let $\mathcal{I}_H = (0, 1]$ be the set of agents with an H-home and $\mathcal{I}_L = (1, 2]$ the set of agents with an L-home. Denote by η the Lebesgue measure over $\mathcal{I}$ with $\eta(\mathcal{I}_H) = \eta(\mathcal{I}_L) = 1$ and $\eta(\mathcal{I}) = 2$. We define a function $h : \mathcal{I} \to \mathcal{N}$ that maps each agent in $\mathcal{I}$ to the type of home they own. Hence $h(\mathcal{I}_j) = j$ for $j = H, L$.

Each period, agent i can choose to supply her house and demand a type-j house. Let $d^i_{jt} = 1$ indicate that agent i demands a type-j home in period t, and $d^i_{jt} = 0$ otherwise. Similarly, s^i_t indicates whether agent i supplies their house, with $s^i_t = 1$ if they do and $s^i_t = 0$ otherwise. As we allow for excess demand and supply agents may not be able to buy and sell what they desire. Let $x^i_{jt} = 1$ if agent i obtains a type-j home in period t and $x^i_{jt} = 0$ otherwise. Similarly, $y^i_t = 1$ if agent i hosts and $y^i_t = 0$ otherwise. We assume $x^i_{jt} \leq d^i_{jt}$ and $y^i_t \leq s^i_t$ for all i, j and t, so agents cannot get more than they ask for. We assume that each agent can visit at most once each period and host at most one other agent in each period. Formally, $\sum_j x^i_{jt} \leq 1$ for $i \in \mathcal{I}$ and $t = 1, 2$, ensuring that each agent either visits one house or is unassigned. Similarly, $\int_{i \in \mathcal{I}} x^i_{jt} d\eta(i) \leq 1$ for $t = 1, 2$ and $j = H, L$, indicating that each house type is visited by at most a unit measure of agents per period.

Agents derive disutility c from hosting and utility v^i_j from visiting a type-j home. Visits are paid with an internal currency. Agents pay p_{jt} units of internal currency for visiting a j-home in t, that is if and only if $x^i_{jt} = 1$, while demanding a house is cost-less. Conversely, agent i earns $p_{h(i)t}$ units of currency if and only if $y^i_t = 1$. The lifetime utility of agent i is given by

$$u(x^i, y^i) = \sum_{t=1,2} \left(\sum_{j=H,L} x^i_{jt} v^i_j - c y^i_t \right) \tag{1}$$

Crucially, the internal currency does not enter agents' utility directly. Agents simply need to satisfy a life-time budget constraint. Hence, tokens earned in one period carry over to the next, as on the exchange platform studied empirically.[5]

$$\sum_t \sum_j x^i_{jt} p_{jt} \leq \sum_t p_{h(i)t} y^i_t \tag{2}$$

We denote the vector of prices by $p \equiv (p_j)_{j \in \{H,L\}, t \in \{1,2\}}$ and $v \equiv (v^i_j)_{i \in \mathcal{I}, j \in \{H,L\}}$, $d^i \equiv (d^i_{jt})_{j \in \{H,L\}, t \in \{1,2\}}$, $s^i \equiv (s^i_t)_{t \in \{1,2\}}$, $d \equiv (d^i)_{i \in \mathcal{I}}$, $s \equiv (s^i)_{i \in \mathcal{I}}$, and x^i, y^i, x and y analogously. To represent aggregate demand for and supply of each home type j in period t, we define $D_{jt} \equiv \int_{i \in \mathcal{I}} d^i_{jt} d\eta(i)$ and $S_{jt} \equiv \int_{i \in \mathcal{I}} s^i_{jt} d\eta(i)$. Analogously, we define define consumption and sales respectively as $X_{jt} \equiv \int_{i \in \mathcal{I}} x^i_{jt} d\eta(i)$ and $Y_{jt} \equiv \int_{i \in \mathcal{I}} y^i_{jt} d\eta(i)$.

2.2 Equilibrium concepts

We now introduce the first key equilibrium concept.

Definition 1 (CE)**.** *Call (p, d, s) a competitive equilibrium (CE) iff all agents choose a sequence of consumption-supply bundles (d^i, s^i) that maximizes their lifetime utility among bundles in their budget set and all markets clear. Formally,*

$$\begin{aligned}
&\textit{For all agents } i, \\
(d^i, s^i) \in \arg\max_{d^i, s^i} \quad & \sum_t \sum_j x^i_{jt} v^i_j - c y^i_t && (3)\\
s.t. \quad & \sum_t \sum_j x^i_{jt} p_{jt} \leq \sum_t p_{h(i)t} y^i_t && (4)\\
& \sum_j x^i_{jt} \leq 1 \textit{ for all } t && (5)
\end{aligned}$$

[5] On the platform agents cannot borrow. While the lifetime budget constraint 2 does permit borrowing, all results go through in an equivalent overlapping generations model where the budget constraint rules out borrowing.

The markets for all goods (home types) clear in all periods

$$D_{jt} \leq S_{jt} \qquad \textit{for all } t \textit{ and } j \tag{6}$$

$$D_{jt} < S_{jt} \implies p_{jt} = 0 \textit{ for all } t \textit{ and } j. \tag{7}$$

and, hence,

$$x^i_{jt} = d^i_{jt} \qquad \textit{for all } t \textit{ and } j \tag{8}$$

$$y^i_t \leq s^i_t \qquad \textit{for all } t \tag{9}$$

$$y^i_t < s^i_t \implies p_{jt} = 0 \textit{ for all } t\ . \tag{10}$$

This is the familiar definition of competitive equilibrium, which requires all agents to behave individually optimal and respect their budget constraints while the resulting aggregate demand and supply clear all markets. Note that agents' choice variables are their demand and supply, (d^i, s^i), but thanks to the market-clearing condition these coincide with consumption and realized supply (x^i, y^i).[6]

We now turn towards the second equilibrium concept, namely a uniform rationing equilibrium. This generalizes competitive equilibrium by allowing for rationing. We assume that excess demand and supply are rationed using uniform tie-breaking numbers TB^i where lower is better. Uniform rationing also features in Glaeser and Luttmer (2003); Che et al. (2013); Dworczak et al. (2021) among others. Agent i's tie-breaking number TB^i is a realization of $\widetilde{TB}^i \sim U[0,1], \forall i$. To simplify the exposition, we assume that tie-breaking numbers are drawn ex ante and known by agents. Hence, every agent who demands a given good has the same (ex ante) probability of being served and every agent who supplies a given good has the same (ex ante) chance of selling.

Definition 2 (URE)**.** *Call (p, d, s) a uniform rationing equilibrium (URE) iff excess demand and supply are rationed strictly based on random tie-breaking numbers and all agents i choose demand-supply bundles (d^i, s^i) with associated (x^i, y^i) that maximize their lifetime utility among feasible bundles in their budget set - given prices and the*

[6]For goods with $p_{jt} = 0$ s^i_t and y^i_t can differ but this does not affect revenues.

rationing rule. Formally,

$$\begin{aligned}
&\textit{For any agent } i \textit{ and any } TB^i, \\
&(d^i, s^i) \in \arg\max_{d^i, s^i} \quad \sum_t \sum_j d^i_{jt} v^i_j - cs^i_t && (11)\\
&\qquad\qquad\qquad s.t. \quad \sum_t \sum_j x^i_{jt} p_{jt} \leq \sum_t p_{h(i)t} y^i_t && (12)\\
&\qquad\qquad\qquad\qquad \sum_j x^i_{jt} \leq 1 \textit{ for all } t && (13)
\end{aligned}$$

and all x^i_{jt} and y^i_t are determined in a rationing process such that

- *if $D_{jt} = S_{jt}$, then $x^i_{jt} = d^i_{jt}$ for all i, and $y^i_t = s^i_t$ for all i s.t. $h(i) = j$*
- *if $D_{jt} > S_{jt}$, then*
  $y^i_t = s^i_t$ *for all* i *s.t.* $h(i) = j$, *and* $X_{jt} = S_{jt}$ *and*
  $x^i_{jt} = d^i_{jt}$ *iff* $TB^i \leq S_{jt}/D_{jt}$
- *if $D_{jt} < S_{jt}$, then*
  $x^i_{jt} = d^i_{jt}$ *for all* i *and* $Y_{jt} = D_{jt}$ *and*
  $y^i_t = s^i_t$ *iff* $TB^i \leq D_{jt}/S_{jt}$

The rationing rule is intuitive. When demand for a good equals supply of that good, all buyers and sellers of that good consume and sell. When demand exceeds supply all hosts who want to host do host and all guests who are selected have lower tiebreak numbers than any applicant who is not selected. When supply exceeds demand, all guests who want to visit do visit and all hosts who are selected have lower tiebreak numbers than any candidate host who is not selected.[7]

Note that we impose $\sum_j x^i_{jt} \leq 1$ but not $\sum_j d^i_{jt} \leq 1$. Hence, Agents' realized consumption and supply have to satisfy their budget-constraint (12) but agents can request more than one house but only one that they actually obtain. Agents obtain their most preferred feasible house but can also demand every more preferred affordable house.[8]

[7]Because agents demand only homes they weakly prefer to their assignment, all agents do want to visit an j-home if they are selected by tiebreak. Note that under the notation above there could be a measure zero market-clearing error. This can easily be ruled out by rewriting the rationing rule as $ess\,sup\{TB^i | x^i_{jt} = 1\} < ess\,inf\{TB^i | x^i_{jt} = 0, d^i_{jt} = 1\}$.

[8]Note that the definition here admits both equilibria where agents demand all preferred homes that

Letting agents demand all their preferred homes allows us to track excess demand. An intuitive interpretation is that agents sequentially request homes until one accepts them (their tie-breaking number is low enough) or no other homes are affordable, which seems to be what users do on the exchange platform. Again, the assumption that tie-breaking numbers are drawn ex ante is only to simplify the exposition. All results can be obtained in an equivalent model where tie-breaking numbers are drawn in an interim stage but this complicates notation.[9]

It is straightforward to show that the competitive equilibrium is a special case of the rationing equilibrium that additionally imposes the market-clearing conditions (Remark 1 in the appendix). As all markets clear, consumption and revenues become independent of tie-breaking numbers and agents' choices satisfy the individual optimality conditions of the URE for any tie-breaking realization.

2.3 Comparison of equilibria

We will derive the set of competitive equilibria, contrast it with a rationing equilibrium that entails a single price for all goods and then compare the aggregate supply and utilitarian welfare of both equilibria.

We assume that preferences do not vary within type, i.e. $h(k) = h(i)$ implies $v_j^k = v_j^i$. Abusing notation we will denote v_j^H and v_j^L type H and L's valuations for type-j homes. Throughout, we make further two assumptions on agents preferences: $v_j^i > c, \forall (i,j)$ and $v_H^i > v_L^i, \forall i$. The first assumptions implies that it is acceptable to supply one unit if this allows to consume one unit of some house (thus, all homes are acceptable). The second assumption implies that agents have the same ordinal preferences. We also assume that agents demand good L when indifferent.[10] To simplify the exposition we

they obtain with zero probability and equilibria where they do not. This is just to simply exposition. In an alternative definition, agents maximize the "pseudo-utility" objective $\sum_t \sum_j d_{jt}^i v_j^i - cs_t^i$. This admits only equilibria where agents do demand all preferred unfeasible homes, because homes that reject them increase the objective without tightening their budget or consumption constraints.

[9] The model is an overlapping generations (OLG) model where every generation (except the first) is born without wealth and lives three periods. Agents have to finance current consumption with revenues from previous periods ($d_{jt}^i p_{jt} \leq \sum_{s=1}^{t-1} y_s^i p_{h(i)s} - \sum_j x_{js}^i p_{js}, \forall t$ and d_{jt}^i, as it is the case on the platform studied in section 3). Therefore, agents can only consume in the last two periods and only have incentives to supply in the first two periods. As period-budgets only depend on past tie-breaking realizations agents face no uncertainty which homes are affordable in the current period. Therefore, demanding their favorite affordable house first and requesting another house if rejected is a straightforward strategy, which individuals plausibly follow in practice.

[10] This assumption resembles Hylland and Zeckhauser (1979) who assume that agents buy the cheapest

consider agents' lifetime supply and consumption choices. Since utility does not depend on which period a good is consumed or supplied this is without loss of generality.[11] Abusing notation slightly, we hereafter denote agent i's lifetime consumption of good j as $x_j^i \equiv \sum_t x_{jt}^i \in \{0, 1, 2\}$ and her lifetime (realized) supply as $y^i \equiv \sum_t y_t^i \in \{0, 1, 2\}$. Analogously, $d_j^i \in \{0, 1, 2\}$ and $s^i \in \{0, 1, 2\}$ denote (desired) lifetime demand and supply. Utilitarian welfare is given by $W(x, y; v, c) \equiv \sum_{i=H,L}(\sum_j x_j^i v_j^i - cy^i)$.

2.3.1 Preference regimes

Naturally, the equilibrium outcomes depend on agents' preferences. Three particular sets of preference profiles will be important. Figure 5 illustrates them graphically.

Definition 3 (Supermodular preferences)**.** *Call preferences supermodular iff*

$$v_H^H - v_L^H > v_H^L - v_L^L \tag{14}$$

Definition 4 (Submodular preferences)**.** *Call preferences submodular iff*

$$v_H^H - v_L^H < v_H^L - v_L^L \tag{15}$$

Definition 5 (SSB)**.** *Call preferences "strongly submodular" iff*

$$v_H^H - v_L^H \leq c/2 < c < v_L^L < v_H^L - v_L^L \tag{SSB}$$

If preferences are supermodular, agents that own a better house themselves have a stronger preference for better homes than agents owning a lower-type house. If preferences are submodular, the opposite is true. If preferences are strongly submodular, lower type agents have a substantially stronger preference for better homes than higher-type

lottery that yields a given level of utility. Agents thereby behave "altruistically" when possible at no cost. Here, all agents prefer H goods, so consuming L goods when indifferent weakly increases aggregate welfare.

[11] The period in which a good is consumed or supplied does not affect agents' lifetime utility or budget constraint (because the population and preferences are identical across periods, prices will be the same across periods). We only need to ensure that aggregate consumption X_{jt} is feasible each period given Y_{jt}, but as agents are indifferent in which period to consume and to supply, there always exists a set of period-wise choices s.t. $X_{jt} \leq Y_{jt}$ if $\sum_t X_{jt} \leq \sum_t Y_{jt}$. Similarly, it is without loss of generality to focus on time-invariant prices. Consider the competitive equilibrium in Proposition 1 under SSB. If in, say, $t = 1$ $p_{A1} = 2p_L$ and $p_{A2} > p_L$, we obtain the same equilibrium allocation as stated in Proposition 1. $d_H^L = 1$, $d_L^L = 0$, $s_1^H = 1, s_2^H = 0, d_L^H = 2$ clear the market and are individually optimal. If $p_{A2} = p_L$, then $D_{A2} > S_{A2}$ and markets do not clear. A similar argument can be made for the non-SNSM case.

agents. What type of preferences do we expect in reality? If valuations vary primarily with the size of a home, supermodular preferences seem more plausible since parties who live in larger homes themselves presumably value large holiday homes more than smaller parties. Yet, if heterogeneity mostly reflects different locations, preferences may well be submodular. The wide range of museums, concerts and international restaurants in a city like New York may be more valued by individuals living in the country-side than by families living in a city like Paris with similar amenities. As we will see, homes in certain locations are much more popular than others. Hence, it would not be surprising to find submodular preferences.

2.3.2 Competitive equilibrium

We now construct and characterize both types of equilibria.

Proposition 1. *In competitive equilibrium aggregate supply of H and utilitarian welfare are uniquely given by*

$$S_H^{CE}(v,c) = \begin{cases} 1 & \textit{if } SSB \\ 2 & \textit{else} \end{cases}$$

and

$$W^{CE}(c,v) = \sum_{i=H,L} \Big(\sum_{j=H,L} x_j^i v_j^i - c y^i \Big) = \begin{cases} 2v_L^H + \ v_H^L - 3c & \textit{if } SSB \\ 2v_H^H + 2v_L^L - 4c & \textit{else} \end{cases}$$

The proof is in Appendix A.1. The hosting and visiting patterns in the two allocations are summarized in figure 2. An arrow from L to H with mass 1 indicates that a mass 1 of L-agents visits type-H homes. Figure 3 shows the prices supporting the respective competitive equilibria.

The key pieces of intuition are the following. Since all homes are acceptable and agents can always afford a house of their own type by hosting once, all agents will visit at least once. They would like to visit twice if it is not too costly. As all types prefer H-homes these homes need to be more expensive than L-homes to clear the market.[12] At prices

[12]Note that with indivisible goods the existence of competitive equilibria is not guaranteed in general. In this model it happens to exist - partly because agents' maximum possible consumption equals the maximum supply.

$p_H < 2p_L$ the H-agents need to host twice to afford two visits of any type and thus prefer H-homes since $v^i_H > v^i_L, \forall i$. Agents want to visit twice since $v^i_j > c, \forall i, j$. Hence, the market does not clear. At $p_H = 2p_L$ Hs can afford two L visits by hosting once. They prefer this over two H if and only if $2v^H_H - 2c \leq 2v^H_L - c \Leftrightarrow v^H_H - v^H_L \leq c/2$. This is precisely the left part of the SSB condition and the horizontal separating line in figure 3. Conversely, L-agents prefer one H over two L if and only if $2v^L_L < v^L_H \Leftrightarrow v^L_L < v^L_H - v^L_L$. This is the right part of the condition and the vertical separating line in figure 3. Note that the middle part of SSB is true by the assumption that all homes are acceptable.

We can understand the supply response of H-types in terms of a standard labor supply model. As p_H increases to $p_H = 2p_L$ from below there are two effects. The first is a substitution effect. The higher price increases the marginal revenue of supplying, but because lifetime consumption is limited to two units this does not increase the marginal utility of supplying as much as it would without the capacity constraint. At the same time the higher price creates an income effect. H-agents can reduce their supply and still afford two units. Hence, reducing the price of H homes - as the following rationing equilibrium does - can increase their supply.

2.3.3 Rationing equilibrium

Proposition 2. *There exists a rationing equilibrium with $p_H = p_L$, $D_H = 4 >= 2 = S_H, S_L = 2$ and rationing of H at rate $2/4 = 1/2$. Agents that do not obtain H then demand L. The resulting allocation is s.t. $E_{TB}[x^i_H] = E_{TB}[x^i_L] = 1$ and $y^i = 2$ for all i. Welfare is*

$$W^{RE}(v, c) = v^H_H + v^H_L + v^L_H + v^L_L - 4c. \tag{16}$$

Proof. Since $X_H = Y_H = X_L = Y_L = 2$ consumption is feasible and rationing of H-homes occurs uniformly at rate $S_H/D_H = 2/4$. Supply and demand for L homes is not rationed. The individually optimally choice functions derived in the Appendix (equations 22 and 21) imply immediately that both types' preferred choice at $p_H = p_L$ is $(s^i, d^i_H, d^i_L) = (2, 2, 0)$. Since $v^i_j > c, \forall(i, j)$ the second-best choice for both types is $(s^i, d^i_H, d^i_L) = (2, 0, 2)$, which is preferred over $(1, 0, 1)$ and $(0, 0, 0)$. Thus, demanding H and then L iff rejected is individually optimal. □

The demand and trade patterns are summarized in figure 4. As both types prefer H homes to L homes and both cost the same all agents try to obtain an H house. Thus there is excess demand for H which is rationed uniformly at random. Hence, half the mass of each type gets a type H house in each period, the others get L. Since all homes are acceptable and each visit has to be financed by one hosting, agents do not have incentives to reduce their supply and agents cannot do better than requesting an H house each period. Thus, the equilibrium choices are individually optimal.

2.3.4 Main result

The following proposition summarizes the comparison of both equilibrium types. This is the main result.

Theorem 1. *Supply of H-homes in the rationing equilibrium is higher than in the competitive equilibrium iff preferences are strongly submodular (SSB). Utilitarian welfare is higher in RE than in CE iff preferences are submodular. That is,*

$$S_H^{RE}(.) - S_H^{CE}(.) = \begin{cases} 2-1=1 & \textit{if SSB} \\ 2-2=0 & \textit{else} \end{cases}$$

$$W^{RE}(.) - W^{CE}(.) = \begin{cases} (v_H^L - v_L^L) - (v_H^H - v_L^H) < 0 & \textit{if supermodular preferences} \\ (v_H^L - v_L^L) - (v_H^H - v_L^H) = 0 & \textit{if identical preferences} \\ (v_H^L - v_L^L) - (v_H^H - v_L^H) > 0 & \textit{if submodular and not SSB} \\ v_H^H - v_L^H + v_L^L - c \qquad > 0 & \textit{if SSB} \end{cases}$$

Proof. The expressions follow directly from the previous propositions. We can immediately determine the signs of the first three cases. For the fourth case, $v_H^H - v_L^H > 0$ and $v_L^L - c > 0$ because we assumed $v_j^i > c, \forall(i,j)$ and $v_H^i > v_L^i, \forall i$. □

Figure 5 display the four regimes graphically. The intuition for the supermodular and submodular cases is simple. Popular goods need to be expensive in CE to clear the market. Only agents that own expensive goods themselves are able to afford them. Thus, the CE-allocation is determined by ability to pay while the optimal allocation depends only on agents' valuations. When preferences are supermodular, the agents with highest ability to pay also have the highest valuations for high-quality goods and the market allocation is optimal. When preferences are submodular, ability to pay diverges from valuations so that the CE is not welfare-optimal. Compressing prices and rationing equalizes ability to pay across agents and thus enables some agents with high valuations and low ability to pay to consume otherwise expensive goods.[13] Note that the competitive equilibrium is Pareto-efficient. Thus, the First Welfare Theorem holds. H-agents are worse off in the rationing equilibrium (H's allocation in the rationing equilibrium was a feasible choice in competitive equilibrium) but this is more than compensated by the gain of L-agents.[14]

The key new insight is that compression can also induce large income effects for agents with popular goods, increase their supply and thus total trade. In other models of rationing, supply is either exogenous or increasing in prices (Glaeser and Luttmer, 2003; Che et al., 2013; Dworczak et al., 2021). In this case, rationing is merely a tool for redistribution, which has to be traded off against reduced supply. Theorem 1 makes clear that there may be no such trade-off in internal currency systems. This is reflected in the following corollary.

Corollary 1. *If there is a "supply effect", compressing prices increases welfare. Formally,*

$$\frac{p_H^{CE}}{p_L^{CE}} > \frac{p_H^{RE}}{p_L^{RE}} \text{ and } S_H^{CE} < S_H^{RE} \implies W^{RE} > W^{CE}$$

The corollary follows immediately from Theorem 1. The nesting structure of the welfare and supply result can also be seen in figure 5. A large section of the empirical part of the paper focuses on clean reduced-form evidence of this supply effect.

2.4 Discussion of model assumptions

Uniform random rationing In assuming random rationing I follow Glaeser and Luttmer (2003); Dworczak et al. (2021) among others. It is worth pointing out, however, that random rationing is a conservative assumption. As agents' chances of obtaining a good are not at all increasing in their valuations, this is a worst case scenario for

[13] This part of the intuition resembles Che et al. (2013) and Dworczak et al. (2021).

[14] In practice, one could imagine that even high-type agents benefit from the additional supply of attractive homes by increasing market thickness. This might help high-type agents find better matches if homes' availability for particular weeks is highly stochastic or agents have strong idiosyncratic preferences for particular homes and thus benefit from variety. While this model abstracts from such considerations this is an interesting avenue for future research.

rationing. Assuming the rationing allocation be correlated with agents' valuations would both increase welfare and incentives to participate under rationing. In practice, there are reasons to believe that rationing works better than under a uniform policy. Appendix B.1 shows that, besides token wealth, hosts' acceptance choices depend large on the timing of requests and characteristics of the trips, e.g. if the party size matches the home size (see figure D.21). Fack et al. (2024) demonstrate that even within a narrow time window hosts on this platform are much more likely to accept the first request they receive. Hence, actual rationing on the platform quite resembles a first-come-first-serve policy. It seems plausible that users who particularly value certain types of homes try to request these homes early and thus tend to have higher chances of obtaining them.

For-money platform as outside option The analysis so far assumes a monopolistic platform so that agents do not have the outside option of hosting or visiting elsewhere. A natural concern is that under compressed prices some H agents might prefer another platform. Appendix A.2 describes a simple extension with a competing for-money platform. Each period, agents can supply on the exchange platform or on the for-money platform but not both. Similarly, they can choose between visiting on one of the platforms and not visiting at all. Hostings and guestings on the exchange platform need to satisfy the same token-money budget constraint as before. Hostings and guestings on the for-money platform are paid for with real money, which linearly enters agents' utility. The disutility of hosting differs across platforms. Hosting on the for money platform creates disutility γc with $\gamma > 0$. Section 6.3 discusses this assumption. The upshot is that guests might behave better on the exchange since all guests host themselves and are expecting a personal residence rather than a professional listing. In this case we expect $\gamma > 1$. It is easy to show that for any prices on the for-money platform and any utility from money, there exists a threshold $\bar{\gamma}$ such that neither type hosts on the for-money platform if $\gamma \geq \bar{\gamma}$. In this case, all previous results go through. Section 5.3 shows empirically that increased price compression does not materially reduce participation on the exchange platform. Furthermore, the willingness-to-pay estimates in section 6.2 suggest that for many platform users the costs of hosting on Airbnb, the largest for-money alternative, are hundreds of dollars higher than the costs of hosting on the exchange platform. Hence, it seems plausible that γ is high in practice and competition with for-money platforms is limited.

2.5 Testable implications

The model makes two key predictions. First, it implies that there are strong income effects. Hence, agents should significantly reduce their supply when they unexpectedly receive a sizable quantity of currency. Second, reducing the price of popular homes can increase their supply. While I do not observe a competitive equilibrium in the empirical setting, I do observe a pricing reform that seems to have moved prices substantially further away from market clearing. Finding that this increased the supply of attractive homes would provide strong support for the model because we would usually expect that lower prices reduce supply.

Prediction 1. *Users' supply reacts strongly to wealth shocks (users can get satiated).*

Prediction 2. *Reducing the price of popular homes can increase their supply.*

In the model, observing that lower prices increase supply implies that preferences are not supermodular and, when prices are not fully compressed, the matching should not be perfectly positive assortative. Instead, agents with attractive homes should visit cheaper homes at least some of the time and thus temporarily accumulate wealth. Properly estimating preferences in a two-sided matching market requires a structural model and is beyond the scope of this paper.[15] Nevertheless, section B.3 presents reduced-form evidence that preferences of the platform users do not increase much in the type of their own home, and might well be decreasing, and that many users visit homes substantially cheaper than their own. Hence, the empirical setting might be in a regime where lower prices can increase supply.

3 Empirical setting and data

Internal currency systems around the world use a variety of price-setting approaches that vary in their degree of price compression. At one extreme are systems that do not limit price dispersion and try to clear the market as much as possible. The other extreme are systems with maximal compression that assign all objects the same price. In between are some systems that allow prices to vary but less than market-clearing prices presumably would. Table D.12 provides an overview of the degree of price compression in various internal currency systems in different sectors. The surprising diversity makes it an important live question if compressed price have any benefits.

[15] Altmann et al. (2024) develop and estimate such a model and study further important design questions of exchange economies with token-money.

3.1 Setting

The platform I study is in between the two extreme pricing approaches. This makes it an interesting setting because we may observe some rationing as well as some (remaining) satiation. Another important reason for choosing this platform is that it is one of the largest internal currency systems in the world and has been online for over a decade. Hence, its users have had some time to accumulate currency. Today, the platform lists about 200,000 completed homes in 187 countries and facilitates over 100,000 trips per year.[16]

The search and matching process is decentralized. Users see the homes of others users that are available at a given time and place and send requests to homes they like. When a host accepts a request and both parties finalized the trip, the guest pays the host a per-night amount of internal currency (IC).[17] The price-setting is described below. When users sign up and complete their profile they receive an initial endowment of IC which allows them to stay several nights in an average home. After this, the main way of earning IC is to host, though users occasionally receive some free IC thanks to promotions or for inviting others to the platform.[18] IC cannot be converted to money and the platform prohibits transfers to friends and checks for such "fraudulent" transactions. Hence, the only way to use IC is to visit others' homes.

Users can join the platform for free but need to pay an annual subscription fee of 150€ to finalize trips but no variable charges per trip. While the search and matching process bears some resemblance with Airbnb, it is important to stress that the users are clearly not professionals. While commercial listings now constitute the majority of offers on Airbnb (Gyódi, 2019), 80% of homes on this platform are primary residences, 90% of users list only one home and 95% of users host and visit no more than three weeks per year. The absence of professional hosts is not surprising. As hosts earn nothing but IC and IC cannot be cashed out, there is no incentive for commercial rental. Sections 6.4 and 6.2 compare the share and number of personal residences on this platform to Airbnb and investigate users' incentives to rent out for money.

[16] Section 6.3 discusses potential motivations for exchanging homes without real money.

[17] Besides normal one-way exchanges of homes for IC users can do "reciprocal" two-way exchanges, where both swap their homes and can optionally pay the price difference between both homes in IC. Due to the double "coincidence of wants" swaps are harder to organize and account only for around 20% of exchanges.

[18] To mitigate liquidity issues, the platform allows users to buy IC when they are about to finalize a trip but do not have enough currency. The cost per IC increases with the share of IC bought. However, the platform does not advertise this possibility and it is used for less than 2% of trips.

Price-setting Prices are largely set by the platform. For each home the platform sets a default price and recommends users to keep this price. It allows adjustments to increase the price by up to 30 points but since the default price is around 150 for the median home and about 300 for the more attractive homes, the possible relative markup is quite low for expensive homes. Figure D.13 shows that the actual prices closely follow the default prices. The prices are not set to clear the market and seem unlikely to do so. An algorithm maps reported home characteristics to the recommended price. Users report the number and types of beds they have, their amenities (for instance an A/C, a garden or a microwave) and how attractive their location is on a range from "*in the heart of international attraction*" to *"away from any tourist site"*. Users do not know the algorithm but only the broad types of inputs. While the information provided is not directly verifiable users receive ratings for the accuracy of their home description and are thus discouraged from misreports. In practice, the default price of an apartment in the best possible location is less than twice as high as the price of the same apartment in the worst location. We may suspect already that market-clearing prices would vary more. Note also that the recommended price does not vary with other obvious features like season or users' reputation. Indeed, section 6.1 shows with two complementary approached that prices are substantially compressed below market-clearing prices.

3.2 Data

3.2.1 Exchange platform data

The analyses are based on proprietary data on the universe of transactions from the creation of the platform until the end of 2023. It covers about 40 million requests and 800,000 completed trips. We observe two kinds of data: First, we observe the complete history of requests and transactions on the platform with details of each request including the time it was sent, the proposed travel dates, number of guests, price, approval decisions and ratings if any were given. Furthermore, we observe the time and length of all messages and indicators of keywords suggesting if the guest preferred a one-way or two-way exchange. Second, we observe snapshots all of profiles in 2021 and in 2023, which include basic information on users (e.g. age, languages, dates of sign-up, ID verification and profile completion as well as reported interests) and very comprehensive information on users' homes.[19]

[19]Similarly, the platform only keeps the most recent version of users' availability calendar. As a consequence, I do not know exactly at what time a host created a calendar entry if part of the indicated period was subsequently booked. Nonetheless, we know with high confidence if a given period was ever

3.2.2 Airbnb data

I additionally collect data on Airbnb listings and their prices between 2021Q3 and 2022Q3 for all 107 locations available on the commonly used "Inside Airbnb" website. These locations cover a third of all homes registered on the exchange platform platform. Figure C.7 shows a map comparing all locations on both platforms. Every three months, Inside Airbnb scrapes all listings in the 107 locations and records their headline price per night, their reviews, as well as home and host information. The home characteristics declared on Airbnb are very similar to those on the exchange platform site and include for instance the number of beds, bedrooms, bathrooms, the max number of guests the house can accommodate and indicators for various amenities, for example garden, balcony, TV, dishwasher, washing machine, swimming pool and air-conditioning. Both platforms also indicate if host and listing are verified.[20]

4 Test of prediction 1: Strong income effects

4.1 Empirical strategy

The first important prediction of the model is that users can get satiated and reduce their supply once they have accumulated sufficient wealth. I test both if hosts' are less likely to accept a request with given characteristics and if they host fewer total nights. For the latter, I construct a home-quarter panel and study if a given host finalizes fewer nights when she is IC-rich at the start of a quarter. To focus on active users with a demonstrated willingness to host the main sample are users that have hosted at least two times in the past, but I also report results for the full sample of completed and verified homes.[21] Let $Supply_{it}$ be a measure of supply of home i quarter t, such as the number of hostings or the number of nights finalized in that quarter. t refers to the quarter when a trip is finalized and not the quarter where the travel takes place. $Wealth_{it}$ is the host's

marked as available in the calendar.

[20] The few characteristics that do not match are the following. Homes on the exchange platform website often indicate their size in square meters, and some preference indicators - such as if the host would like the guest to water her plants. Homes also contain somewhat more detailed information how suitable homes are for children. Airbnb listings sometimes indicate detailed information on small items, for instance, if soap, towels are available, which are standard for exchange platform and thus not mentioned.

[21] Completing the listing 100% implies that the host provided a proof of address to verify the home and uploaded at least 5 pictures. Yet, not all such homes have hosted by the end of our sample period.

wealth at the very beginning of the quarter. I estimate the following regression

$$Supply_{it} = \mu_i + \beta_1 Wealth_{it} + X_{it}\theta + \varepsilon_{it} \tag{17}$$

The home fixed effects μ_i account for the fact that the owners of some homes may be able or willing to host more often and thereby earn more currency, which creates a spurious positive correlation between wealth and supply in a cross-section. X_{it} are other important determinants of supply. All specifications include at least indicators for the number of quarters the user is signed up on the platform and location-by-quarter fixed-effects to account for the fact that locations may be more popular in certain quarters and that users might reduce or increase their supply over the life-cycle as changes in their life circumstances become more likely.[22] To allow for non-linear effect, I also estimate a specification with indicators for wealth deciles $D_d, d = 1, ..., 10$. Deciles are defined pooling all homes and quarters. I drop users whose wealth decile never changes.

$$Supply_{it} = \mu_i + \sum_{d=2}^{10} \beta_d \mathbb{1}(D_{d-1} < Wealth_{it} \leq D_d) + X_{it}\theta + \varepsilon_{it} \tag{18}$$

We expect increasingly negative coefficients β_d. Variations in wealth within host arise from three main sources: First, how many nights a user visits and at what price, which depends partly on the available homes. Second, how many nights she gets to host, which depends on requests received. Third, free "promotional" IC obtained without hosting, which depend on the platform's policy. Thus, users only have partial control over these factors.[23] Section 4.5 isolates exogenous wealth variations more explicitly in a quasi-experimental design that exploits variation in promotional internal currency.

[22] As a robustness check, I also include lagged number of nights finalized and "pending" nights as users may not want to host in two consecutive quarters. The main specification does not include lagged outcomes as this may induce Nickell-Bias in the presence of individual fixed effects (Nickell, 1981). Reassuringly however, the estimates with and without lagged outcomes are very similar.

[23] A remaining concern might be that hosts are deliberately going to accumulate wealth in advance when they anticipate not being able to host next year or quarter. This would imply over-estimating the wealth effect (in absolute size). However, the opposite bias seems just as likely. Latent availability may well be positively auto-correlated: when users renovate their house, they might not host multiple consecutive periods but keep visiting. When a family member temporarily moves out for a year, she might host more. If wealth had no causal effect on hosting, this auto-correlation would create a positive relation between wealth and supply as high wealth proxies episodes with high availability. Another reason that we might under-estimate the negative effect of wealth somewhat is that the location-month fixed effects pick up increases in average wealth in demanded locations and are thus potentially a "bad control".

4.2 Benchmark results

Table 1 shows the results of estimating equation 17. Increasing a given users' wealth by 1000IC reduces the number of times she hosts by around 50 percent of the mean (the standard error is less than one percentage point). The effect has a similar size for the number of nights a user hosts per quarter. Thus, satiated users do not host the same number of parties for fewer nights but host fewer parties. This makes sense as the costs of hosting a largely fixed - notably cleaning and arranging the exchange of keys. Indeed, the probability that a users hosts at least two or three times in the same quarter is reduced even more, by about 65 percent and 90 percent respectively (Figure D.17 shows additional outcomes). Similarly, the effects are larger for hosting trips taking place in periods not marked as available in users' calendar ("cold call" requests).

The estimated effects are very robust and very precise. Estimates are similar with and without controls for lagged supply (Table 1). When including users with no or little hosting record in the sample, the effects naturally become smaller but remain large and highly significant (Table D2) and Figure D.18). Figure 6 shows the results of estimating the equation 18. While an average host is predicted to finalize about 0.7 hostings in a quarter when she holds between 0 and 1000IC at the beginning of the quarter, the same host is predicted to host only 0.2 times at wealth 4000IC and only 0.1 times when her wealth approaches 6000IC.

Interestingly, higher wealth does not reduce hosting "reciprocal" two-way exchanges (Table 1).[24] One interpretation is that some users' preference for reciprocal visits is very inelastic, perhaps because of trust. Another interpretation is that users keep accepting guests in hosting two-way exchanges as this allows them to visit better apartments. Suppose that many hosts with prime Manhattan apartments stopped hosting for IC because they already hold enough currency. Yet, if somebody proposes them a swap to a house in central Paris they might accept. Thus, two-way exchanges could be a fall-back option once the market for currency-based exchanges collapses. However, two-way exchanges are not able to replace one-ways as the total number of nights hosted clearly decreases with wealth as discussed aove. This makes sense, as two-way exchanges are harder to organize and require many more requests.

[24] This is similar for self-initiated and other-initiated two-way exchanges.

4.3 Additional results

Additional analyses in the appendix provide further evidence that the negative effect of wealth is driven by the host. Appendix B.2 presents results using availability marked in homes' calendars as alternative measure of supply. Although many users do not use the calendar feature and the data on historic availability has some limitations I find clear negative effects of wealth on calendar supply. Appendix B.1 examines hosts' decision to accept or reject requests. I find that a given host is much less likely to approve a given request (controlling for almost all observable characteristics) the wealthier she is at the time of the request. The probability that a given host accepts a request with given characteristics decreases monotonically from ca. 22% at almost 0 wealth to only 5% at 8000 IC- a reduction by more than four times. Furthermore, a flexible XGBoost machine learning exercise shows hosts' wealth at the time of a request as one of the strongest predictors of approval.

4.4 Wealth dynamics

Figure D.14 shows the resulting dynamics of users' wealth. I estimate specification 18 with indicators for increasing or decreasing wealth by the next quarter as outcomes. When users' wealth is close to zero at the start of a quarter, almost 80 percent increase their wealth until the next quarter. This probability falls to around 20 percent at wealth 3000IC and to almost zero at even higher wealth levels. Conversely, the probability to decrease wealth increases monotonically in users' wealth. They are very unlikely to reduce their wealth below 1000IC but are more likely to decrease their wealth than to increase it as they reach 3000IC. Thus, users clearly avoid accumulating too much currency.

4.5 Unexpected wealth shocks

Institutional details While the main source of internal currency is hosting other members, the platform occasionally transfers some users additional, "promotional" IC. In particular, users receive IC for referring new members. If user j, the "godson", enters the referral code of an existing user i, the "godmother", both users will receive 50IC when j completes her profile on the platform. Recently, the platform introduced an additional 200IC bonus for i when j pays the 160EUR membership fee, which typically happens upon finalizing a trip (the paid membership is only required at the final step of booking trips). While i might try and recruit new members when she needs additional internal currency, it is hard for her to predict or even control if and when j completes her listing and even more when and if j wants to book a trip and manages to find a match. In fact, over 57 percent percent of referred users never complete a listing and 74 percent never become paying members. If they do, this often happens months after completing their profile. Furthermore, godmothers may not even be aware of the second bonus, because for a long time there used to be only a single bonus upon profile completion. Hence, when i eventually receives the 200IC this likely represents an unexpected wealth shock that is uncorrelated with her intentions of finalizing a hosting in that period. I focus on this second, less predictable and larger bonus.

Empirical strategy A regression discontinuity design leverages the resulting variation. I construct a monthly panel of "godmothers" around the time when they receive the referral bonus, t_0, and test if users reduce their supply just after receiving the bonus internal currency.[25] The identifying assumption is that other determinants of users' supply change smoothly around the discontinuity. Finding a statistically significant effect is challenging because the wealth transfer is quite small, just about the cost of one night in an average home, and the referral bonus upon membership was introduced only recently, limiting the number of events. Our estimates in the previous section suggest that 1000IC additional wealth reduce supply by about 50 percent. Hence, we would expect a short-term reduction of supply by around 10 percent. Two further complications are that users may enter referral codes retroactively after signing up and that some users might try to fraudulently obtain the referral bonus by creating fake profiles. In both cases the membership bonus is likely transferred soon after the profile completion bonus. To rule out such cases, our main specification omits the periods immediately before and after the event ("donut RDD") and we exclude events where the subscription occurs within 30 days of the godson completing his profile. This way, we restrict attention to events where the godson's membership subscription was plausibly unexpected for the godmother. As the previous results show that wealth affects only 1-way hostings, we focus on this outcome to reduce noise. I follow the approach of Calonico et al. (2014, 2017) and estimate the first stage and reduced form discontinuities using local linear functions, optimal bandwidths and valid standard errors.

[25] The monthly level offers an acceptable compromise between measuring supply reasonably well and maintaining exogeneity of the wealth shocks. At finer level it is hard to measure supply accurately as average monthly supply in this sample is only 0.13. As many godsons pay the membership fee 3-10 months after creating their profile, godmothers likely cannot predict well or even control which month the godson will book a trip and pay the membership fee.

Results The top panel of figure 7 shows the expected first stage effect. Users' wealth sharply increases at the time when the godson pays the membership fee. The lower panel shows supply. Looking at the evolution of supply before the event provides support for the identifying assumption. If godmothers were able to strategically obtain the membership referral bonus, they would likely do so at times where they are unable to earn internal currency by supplying. Hence, we would expect a downward trend in supply before the transfer. This is not the case, however. Instead, supply sharply falls after the event from about 0.14 monthly hostings to 0.125 - an 11 percent reduction. Table 2 and figure D.15 show that the discontinuity persists for different bandwidths and polynomials as well as with and without "donut" design. Hence, even the users in this sample, who host more frequently than average users and are enthusiastic enough about the platform to recruit new members, respond to an income shock as small as this.

In summary, increasing the wealth of a given user makes her much less likely to approve a given request, less likely to mark her house as available and, consequently, much less likely to host overall. That confirms implication 1 of the model. Users clearly can get satiated with internal currencyand experience strong income effects.

5 Test of prediction 2: Effect of prices on supply

5.1 Empirical Strategy

The second key prediction of the model is that reducing the price of popular goods can increase their supply, because it limits the purchasing power of "rich" users and reduces the risk of satiation. To assess the effect of prices on supply I exploit a reform of the pricing algorithm. After the reform, the new pricing algorithm was applied to all homes but affected homes differently. It strongly reduced the price of many expensive homes but induced only small or no changes for other homes.[26]

I examine in a difference-in-difference design if the expensive homes that became cheaper during the reform host more often after the reform. I define a treatment group that experienced a large price reduction and a control group where prices did not change substantially. For the main specification, I consider homes to be treated if their price was reduced by at least 30 percent and as control if their price changed by at most

[26] Figure 10 displays home prices before and after the reform relative to imputed Airbnb prices. Overall, the reform substantially compressed prices, creating exactly the kind of variation needed to test the theoretical predictions.

10 percent. I discard other homes but report estimates for alternative definitions of treatment and control group. I estimate changes in supply of the treatment group relative to the control group before and after the reform. To analyze treatment effects over time, I construct a home-year panel. The effect is likely to appear only gradually because users host infrequently, because many held some wealth at the time of the reform and because many might not notice the price change immediately.[27] I estimate a difference-in-difference event-study via the following two-way fixed effects specification

$$Y_{it} = \mu_{h(i)} + \sum_{s \neq -1} \lambda_s \mathbb{1}(s = t) + \sum_{s \neq -1} \beta_s \mathbb{1}(s = t) D_i + \epsilon_{ijst} \tag{19}$$

where Y_{it} is an outcome of home i in year t, $D_i \equiv \mathbb{1}(p^*_{i0}/p^*_{i-1} \leq 0.7)$ is a treatment indicator, λ_s are time fixed effects and $\mu_{h(i)}$ are fixed effects for cells of home types, which determine the price change in the reform. Some specifications additionally include location-year fixed effects to account for differential changes in demand across locations over time. Note that the reform happened at a single point in time for all homes. Hence there are no concerns about dynamic contamination of the control group. Following Farronato and Fradkin (2022), our main measure of supply Y_{it} is the number of hosting nights in home i that were booked in year t.

I restrict the sample to homes with 1-15 regular beds where both the home profile and the owner's user profile are at least 85% completed (which implies that the home has least 5 pictures and is verified). To capture the total effect on supply, including participation, I do not condition on ever hosting in the main specification and keep all homes in the sample until the last period, even if the user becomes inactive at some point. Effects on participation are very likely concentrated on the "intensive" margin, after creating a listing. Since home owners do not know the price of their home until they create the listing, signing up on the platform itself should not be affected by the price change. Including users who never host leads to noisier estimates. To increase power, the main sample therefore includes homes that joined at different points in time. In robustness checks I also use a strongly balanced panel of homes that all joined a certain number of years before the reform. Again to increase power, I also estimate a simple before-after

[27] The large wealth effects observed in section 4 suggest that users chose their supply somewhat myopically and might not react immediately to changes in expected future revenue. In fact, most hosts presumably only noticed months after the reform that the price of their home changed when visiting the website to look for a holiday or post a calendar entry, which users do infrequently. For the same reason, and because on average users host in less than one out of four quarters. I aggregate to the annual level.

difference-in-difference specification as follows

$$Y_{it} = \mu_{h(i)} + \alpha After_{it} + \beta D_i \times After_{it} + \varepsilon_{ijst} \tag{20}$$

where $After_{it}$ equals 1 for all homes following the price reform and 0 before.

5.2 Results

5.2.1 Main results

Figure 8 shows the evolution of prices in the treatment group and the control group of homes. Prices in the treatment group dropped sharply in the reform year and do not change in other years. Figure 9 shows the results of estimating the difference-in-difference event-study specification (eq. 19). The figure shows that the number of nights that treatment and control groups follow very similar trends before the reform.[28] We expect that homes whose price decreased during the reform start hosting more nights after the reform relative to control homes. Indeed, the supply of nights gradually increases after the reform and is about 20-25% higher (relative to $t = -1$) three years after the reform.

Table 3 shows results for further outcomes. Given the absence of pre-trends, I focus on the before-after difference-in-difference specification. The effects on the number of times hosted are slightly smaller than the effects on nights hosted but very similar. Hence, treated hosts mainly react by hosting more often and host only slightly longer exchanges. Treated hosts particularly host more guests who sent unsollicited requests for periods not marked as available in the host's calendar. This makes sense, since hosts could more easily afford turning down spontaneous requests at the high pre-reform prices. Relative to the outcome mean effects are also larger for "reciprocal" two-way exchanges than for one-way hostings. This suggests that affected hosts substitute somewhat to reciprocal exchanges, which are typically not paid with internal currency. This response seems rational. For treated users the reform deteriorates terms of trade with internal currency while two-way exchanges are unaffected. Thus, two-way exchanges become relatively more attractive. It is important to emphasize, however, that this is just a marginal shift as only a quarter of hostings are reciprocal.

[28]Implementing the approach of Rambachan and Roth (2023) shows that results are robust to differential trends much larger than the worst pre-treatment violations.

5.2.2 Robustness

The results are robust to a range of alternative specifications. Table 3 shows that including location-year fixed effects barely changes estimates. Table D3 displays results for various alternative definitions of treatment and control groups. As we would expect, the effects tend to be stronger for homes that experienced larger price reductions. Also as expected, the effects are stronger in highly demanded locations where hosts can get satiated more easily (table D5). As an alternative measure of supply I study the number of nights hosts offer their house in the calendar. This measure has two drawbacks. First, only half of all hosts uses the calendar feature and even those who do sometimes accept unsollicited requests for periods not offered in the calendar. Second, the data on historical calendar availability have some measurement error. Nevertheless, I find that treated hosts offer significantly increase the number of nights offered in their calendar (table D4 and figure D.22).

Table D6 and figures D.24 and D.25 show that the results are very similar and clearly significant in a strongly balanced sample of homes created multiple years before the reform, although the smaller sample leads to less precise estimates. I also assess the sensitivity of the event-study estimates to differential trends. Following the "honest difference-in-differences" approach of Rambachan and Roth (2023) I calculate the confidence sets for various values of M. A value of $M = 2$ ($M = 1.5$) means that differential post-treatment are allowed to be two times (one and a half times) as large as the worst violation in the pre-treatment period. Figure D.30 shows that the estimates for period two break down only at $M = 1.5$ and the estimates for periods three and four are robust violations up to $M = 2.5$.

5.3 Effect on participation

As noted before, the supply effects above already include the dampening effect of potentially reduced participation. When users become inactive or end up never hosting after creating a profile their supply is recorded as zero. Nonetheless, it is interesting to study participation directly. One might expect users with attractive homes to substitute to Airbnb when their prices are compressed substantially and their terms of trade deteriorate. After all, most users are presumably aware of Airbnb as an outside option and many have likely used it as guests.

I create several measures of participation. The first two track users' survival on the platform. I define a user as still active in a given year if and only if she sents any message on the platform in or after that year (user did not exit until this year). Similarly, I define an indicator if the user ever hosts in after the current year. I focus on these survival outcomes because users rarely explicitly declare their departure from the platform by deleting their account. Furthermore, many users exchange homes infrequently and often finalize exchanges after several years without exchanges. Sending messages is by far the most frequent user activity observed in the data and necessary to finalize a trip. In principle, the reduced prices could also deter candidate users from entering the platform. Since users do not know their price until creating a listing, this seem unlikely. The only plausible way this could happen would be through word-of-mouth. Hence, I also study if treated users refer fewer new users after the reform.

Table 4 shows the difference-in-difference estimates for all participation outcomes. Figure D.29 shows the corresponding event-study. There are no significant negative effects on any outcome. The effects on survival are small, positive and precisely estimated. The positive estimates are likely an artifact of the imperfect survival proxy.[29] Nonetheless, this finding speaks against a sizable drop in user participation. We can rule out a 10 percent reduction in referrals. To rule out any effects on participation not captured by the survival measures, I also study how the total stock of active homes with treated and control characteristics evolves.[30] Figure D.28 suggests that the number of homes with treated characteristics grew minimally slower than the control group in the first year after the reform. Hence, we cannot rule out that some dissatisfied users left immediately after the reform. After two years, however, this gap completely disappears and the number of treated homes evolves exactly parallel to control homes.

Overall, the results in this section confirm the prediction of the model that reducing the price of popular goods increases their supply. The fact that few if any users leave following the much stronger price compression is notable. Section 6 investigates users' incentives to participate and shows that many users are willing to forgo large amounts of after-tax Airbnb income.

[29] As the treatment induces the average user to host more frequently, some users who leave exogenously (independently of treatment status) are observed hosting a year longer if treated than if untreated.

[30] Recall that the pricing algorithm is a deterministic function of home characteristics. Hence, I can define the treatment group in exercise as all homes with characteristics that induced a 30% price reduction during the reform.

5.4 Other side-effects of compressed prices

5.4.1 Moral hazard

The effect of compressed prices on hosts' effort is ambiguous. On many platforms sellers try hard to obtain good ratings as this allows them to charge a higher price. Here, reputation does not enter the recommended price. What's more, hosts in over-demanded locations might find guests quite easily even without perfect ratings. This could allow them to reduce their effort. At the same time, hosts need to host more often when their price is reduced and might thus have incentives to increase effort.

I test in the same difference-in-difference design if the reform affected the ratings of treated hosts. Table 5 shows that the price reduction lead to slightly worse ratings given to treated hosts. Yet, the effects are very small and not significant. The point estimates suggest a 0.3 percent deterioration in overall ratings (column 1) and a 0.5% deterioration in cleanliness ratings (column 2). We can rule out a 1 percent decrease. Similarly, the ratings that hosts give their guests changed less than 1 percent (column 3, the point estimate is positive but not significant). Appendix B.4.1 shows in an additional cross-sectional analysis that the same guest gives slightly worse ratings to hosts in over-demanded places, using the excess demand measure developed in appendix C.1. Again, the effects are very small however.

5.4.2 Misallocation

Last but not least, a concern may be misallocation. Glaeser and Luttmer (2003) argue that under rent control homes are not allocated to the tenants with the highest valuation and present evidence that apartment sizes are less related to family size in US cities with rent control than those without. Intuitively, only (large) families, who strongly value large apartments, request those apartments when they cost significantly more than small places. When the price difference is small, smaller parties might be tempted to request big places, too. On the other hand, there could also be small parties with upscale properties (say, older couples) that become satiated when prices are less compressed and, subsequently, visit large homes even when they do not need them, while some large families with lower quality properties might be budget-constrained as guests. Hence, the effect of price compression is again ambiguous.

Column 4 of table 5 shows that the reform slightly increased the average party size of groups requesting the large, expensive places that became over 30 percent cheaper. The increase is 0.2 guests relative to a mean of 3.1 (7 percent) and significant. Furthermore, appendix B.4.2 shows that the utilization of beds does not decrease with excess demand in the cross-section either. These findings mitigate concerns about misallocation at the current prices.[31]

Overall, we find little evidence of sizable side-effects.

6 Why do users not leave to Airbnb?

The last section found that few if any users with attractive homes left the platform after the reform heavily reduced their prices. This section examines users' incentives to participate in more detail. I start by comparing prices and terms of trade on both platforms before discussing user motives and studying the types of hosts and properties on both platforms.

6.1 Degree of price compression

I start by showing that prices are much more compressed than on Airbnb and confirm with a second, complementary approach that prices are far from market-clearing. I also find that a reform of the pricing algorithm increased the degree of compression.

6.1.1 Comparison with Airbnb prices

The first approach examines how close the relative prices of homes on the exchange platform platform are to what their relative prices would be on Airbnb. It seems natural to expect that home characteristics that are valued by Airbnb guests - such as a good location, more bedrooms and amenities - are also valued by users on the exchange platform platform, leading to similar relative prices. To impute the hypothetical Airbnb prices of homes on the exchange platform platform I use a hedonic regression model. That is I estimate the price $Price_{it}$ of an Airbnb listing i at time t as a flexible function $g(.)$ of its characteristics X_{it} and then apply this model to homes on exchange platform.

$$Price_{it} = g(X_{it}) + \epsilon_{it}$$

[31]In practice, the platform could also use the search algorithm to prioritize large groups for big homes and assign them a low rank for small groups.

X_{it} contains home characteristics that are declared on both sites, such as the number of beds, bedrooms, bathrooms, the maximum number of guests the house can accommodate and indicators for various amenities, for example garden, balcony, TV, kitchen, dishwasher, washing machine, swimming pool and air-condition. It also includes the average and median Airbnb listing price per bedroom, per bed and per bathroom in the home's neighborhood in that quarter. Finally, X_{it} contains the number of ratings and the average star rating the listing obtained, and indicators if the host's ID is verified, if she is a superhost and if the listing is instantly bookable.I keep only verified full apartments that have received some ratings since offering rooms is very rare for exchange platform . To estimate $g(X_{it})$ I train a flexible Extreme Gradient Boosting (XGB) machine learning model. The model has a good fit on a hold-out sample of over 80,000 Airbnb listings not used in the prediction, shown in figure C.8. I then apply this model to the characteristics of exchange platform homes and predict their Airbnb price. I only use homes in city districts where I observe Airbnb prices.[32]

Figure 10 plots IC prices against imputed Airbnb prices. If both platforms set the same relative prices, all data points should be on a straight line. The figure shows a line going through median Airbnb and IC price and the origin. The IC prices turn out to be much more compressed than homes' hypothetical Airbnb prices.[33] For instance, houses whose imputed Airbnb price is around 800 USD, roughly four times the median imputed Airbnb price of platform homes, have an average IC price that is less than twice the median IC price.[34] In addition to the current IC prices figure 10 also shows the prices that were set before a reform of the pricing algorithm several years ago. Even more remarkably than the previous result, I find the pre-reform IC prices were closer to Airbnb prices than the current prices. Thus, over time the platform compressed prices more - with the stated purpose of increasing trade.

[32] I set the indicators for superhosts and instant booking to zero for all exchange platform homes.

[33] Relative prices might differ across platforms due to differences in supply. We would expect that individuals whose homes are strongly under-priced on the exchange platform platform are less likely to use the platform. Such selection would increase the scarcity of very attractive homes and thus increase their market-clearing price on the exchange platform platform. Hence, we would underestimate the distance from current prices to market-clearing prices.

[34] This gap is estimated very conservatively as I abstract from some characteristics that Airbnb likely differentiates prices on and the exchange platform platform ignores, such as ratings, and I the Inside Airbnb data cover almost exclusively large cities. Hence, the comparison misses the price differential between urban and rural areas that is almost certainly larger on Airbnb. This is another reason why I underestimate the distance from market-clearing prices.

6.1.2 Excess demand

The second approach looks for signs of excess demand and rationing. It is described in detail in Appendix C.1 and only summarized here. In a competitive equilibrium we would expect an acceptance rate of 100% everywhere. Market-clearing requires that the number of users willing to pay the posted price for a given type of home equals the supply of such homes. In practice, there are substantial search frictions even on Airbnb and more than 40% of requests get rejected (Fradkin, 2017). Yet, systematic differences in rejection rates across home types suggest differences in demand relative to supply. Indeed, requests to homes that are under-priced relative to Airbnb are much more likely to be rejected (see figure C.9).

Since the Airbnb data cover only large cities I separately analyze rejection rates across all locations on the platform. The chances of being accepted vary dramatically across locations. For instance, users looking for a stay in New York City are 5-10 times less likely to be successful than in less popular places. To reduce this massive excess demand prices would need to rise substantially. Exploiting the reform of the pricing algorithm, I estimate the price-elasticity of demand in a difference-in-difference-event-study. The resulting estimates suggest that prices would need to increase by a factor of 2-4 in the most demanded locations. Taken together, these two analyses clearly suggest that current IC prices are far from market-clearing.

6.2 Willingness-to-pay for exchange platform

6.2.1 Revealed preference argument

This section tries to measure how strong users' preference for the platform is. The idea is a simple revealed choice argument. It seems plausible that most users are aware of Airbnb as an outside option. Each time users host on the platform they presumably could have rented out on Airbnb and earned real currency. Similarly, given the size of Airbnb, users presumably could have rented an Airbnb listing with comparable features each time they visit a home on exchange platform. Given the stark price compression observed in section 6.1 some users with attractive homes would face substantially more favorable terms of trade on Airbnb.

For concreteness, imagine a user with a home worth 300 USD per night on Airbnb and worth 170 IC. Say she hosts 9 nights on exchange platform. Around the same time visits an apartment worth 150 IC and but only 100 USD on Airbnb for 10 nights. On net, she earns $9 \cdot 170 - 10 \cdot 150 = 30$ IC. On Airbnb her hosting would have cost 2700 USD, her visit only 1000 USD. Airbnb charges a service fee of 14-16 percent. Bibler et al. (2021) estimate that taxes are paid on at most 24 percent of Airbnb transactions. As many countries have allowances for occasional hosts, most exchange platform would presumably pay little to no taxes on their Airbnb revenue. If we nevertheless assume a tax rate of 15 percent, the user could have earned $(1 - 0.16) \cdot (1 - 0.15) \cdot 2700 - 1000 \approx 930$ USD on Airbnb while hosting the same duration and visiting a similar apartment. As the user gained 30 IC during the exchange transactions, we additionally deduct the replacement value of the internal currency, approximately 40 USD.[35] Hence, she could have obtained the same supply-consumption bundle on Airbnb and would have additionally received a net income of around 890 USD. We interpret the fact that she forgoes these earnings as the user being willing to sacrifice at least 890 USD for using exchange platform over Airbnb.

6.2.2 Results

Following this idea, I estimate the forgone Airbnb income for all users with a home in a city with Airbnb data. I consider their hostings as well as all visits to homes whose Airbnb price we could estimate and group host-visit transaction pairs at the user-week level. Table 6 shows the willingness-to-pay (WTP) estimates. For 52 percent of users we observe at least one transaction with forgone Airbnb income of 100 USD or more. 38 percent of users forgo at least 500 USD, 21 percent 1,500 USD or more. If we restrict to users with a listing worth at least 300 USD per night (row 2), who are the most affected by price compression and potentially more concerned about damage to their property, we find that 81 percent sacrifice 100 USD or more. About 53 percent forgo at least 1000 USD in net income.

As a robustness check, the lower part of the table restricts to pairs of exchanges where the user earned 0 IC or lost IC. For these transactions, we do not need to specify a replacement value and can get an extremely conservative lower bound of willingness-to-pay by not counting the dollar value of lost IC as benefit of using Airbnb for these transactions. In reality, users would have gained dollars *and* IC by using Airbnb. Even under these overly conservative assumptions, we still find that 40 percent of users with a home worth at least 300 USD per night could have made over 500 USD of additional

[35]To mitigate liquidity issues, the platform allows users to buy IC when they are about to finalize a trip but do not have enough currency. The cost of 1 IC is between 1 USD and 1.25 USD depending on the country of residence.

net-income while simultaneously saving IC.

6.2.3 Discussion of assumptions

These estimates provide lower bounds for several reasons. First, in many of the transactions users *lose* IC and thus need to host more at other points in time than they would on Airbnb. We do not count the dollar value of lost IC as benefit of using Airbnb, while we deduct the value of earned IC from Airbnb revenues. Second, the data only cover popular Airbnb destinations, which are almost exclusively large cities or metropolitan areas and more expensive than many rural destinations. Therefore, we likely miss transactions with more extreme imbalances where owners of urban homes visit rural homes that are very cheap on Airbnb. Third, Airbnb is only one outside option and in principle, consumers might achieve even better terms of trade by multi-homing on other sites for commercial vacation rental or by staying in cheap hotels. Overall, the results suggest that a sizable share of users has a strong preference for exchange platform over for-money rental platforms.

6.3 Possible sources of pricing power

It is remarkable that many consumers choose to rent out their homes on a platform that uses a far less fungible, internal currency and offers worse terms of trade than Airbnb. It suggests that many individuals are reluctant to list their primary residences on Airbnb. Anecdotal evidence point towards different norms on both platforms. For instance, in an Airbnb forum discussion about house exchange a user familiar with both platforms writes: "Not the same mentality. Much nicer and respectful relations between members". We can understand this as guests behaving better on the house exchange platform, which reduces costs of hosting.[36]

There are several plausible explanations for better guest behavior. First, the internal currency system forces all users to act on both sides. We can imagine that users who have hosted themselves and know what it is like to let a stranger sleep in their own bed are more mindful when staying at someone else's place. In line with this idea, Klein

[36]Consumers might also have a preference for visiting more "authentic" primary residences. Yet, this seems unlikely to be the main explanation. In 2022, 150 million consumers used Airbnb as guests (see here), but there are only 6 million Airbnb listings globally. Given the estimated 11% share of personal residences in section 6.4, not much more than one million of those are primary residences. Hence, less than 1 percent of Airbnb guests are willing to host themselves.

et al. (2017) reveal in interviews with users who hosted both on Airbnb and Couchsurfing that Airbnb shifts perceived power from hosts to guests and that hosts trust their guests more on Couchsurfing. Second, house exchange guests might have different expectations. Section 6.4 shows that only a third of rated Airbnb listings are plausibly personal residences. Since professionals supply more frequently, personal residences account for less than 10 percent of all ratings given. As the vast majority of users' interactions are with quasi-professional hosts guests presumably expect a more seamless, hotel-like experience. It may be demanding for non-professional hosts to provide such an experience in their primary residence.[37] By contrast, it is pointless for professionals to host house exchanges precisely because the internal currency cannot be cashed out. Therefore, exchange platform guests likely expect a less professional experience. Last but not least, guests and hosts alike might value the more personal interactions arising on this platform. Alvin Roth points out that "in matching markets you care who you are matched with. And there's something impersonal about buying things in a commodity market." (Roberts, 2015).

6.4 Share of personal residences on each site

To support the idea of different norms on both platforms we try to measure the share of personal residences on each platform. There are no signs of professional hosts on the exchange platform, which is unsurprising given the limited use for IC. Hence, almost all listings are very likely primary or secondary residences. 80 percent of homes are indicated as primary residences.

6.4.1 Proxy for personal residences on Airbnb

As Airbnb does not elicit or even display which listings are personal residences, we have to construct a proxy. Following Chen et al. (2022), who use proprietary Airbnb data, I assume that non-professionals never operate more than one entire-home property in a single city in any quarter. Like Almagro and Domínguez-Iino (2024) I require that listings do not receive more than ten reviews per year on average (the scraped Airbnb data do not show bookings directly). This definition seems very conservative and is consistent with hosting up to 30 times per year - much higher than almost any exchange

[37]Conversations with numerous users support this idea. It is also in line with the high bring-to-market costs described by Filippas et al. (2020) and Farronato and Fradkin (2022) finding higher Airbnb seller costs in cities with more families and fewer single households.

platform user. [38] I also report results for more realistic definitions assuming 3-5 ratings per year. I only consider verified listings offering entire apartments or homes as renting out individual rooms is very uncommon for exchange platform . I also discard listings without any ratings, listings available only to stays longer than 28 nights and hosts with a response rate below sixty percent.

6.4.2 Results

Figure 11 shows that 97 percent of exchange listings in the 107 destinations with Airbnb data satisfy this definition. By contrast, only a third of Airbnb listings in the same locations are plausibly personal residences. The difference becomes even starker when studying transaction volumes. I compute the share of ratings given to non-residence listings among the ratings given to all listings. Naturally, the quasi-professional hosts offering non-residence listings typically rent out more frequently than residences. As a result, residences receive less than 9 percent of all Airbnb ratings (figure 11).[39] Hence, Airbnb guests rarely visit personal residences and predominantly interact with quasi-professional hosts. Therefore, Airbnb guests likely form different expectations than exchange users who always visit personal residences.

6.5 Absolute size of both platforms

Last but not least, I show that the exchange platform is not completely niche but appeals to a large population. We would expect Airbnb to hold a much larger market share than the exchange platform even in the market for personal residences for a variety of reasons. First, Airbnb's earlier entry into the market provides it with a first-mover advantage, which likely contributes to stronger network effects. Second, as one of the world's largest providers of vacation rentals Airbnb benefits from greater brand recognition and cross-

[38] Fradkin et al. (2018) reports that 67 percent of guests write a review after visiting a listing that does not have any reviews yet. Arguably, incentives to review are lower when a listing already has many ratings. Comparing publicly shown ratings to subpoenaed Airbnb booking data the Budget and Legislative Analyst's Office of San Francisco estimates that only 30.5 percent of booking were rated (Brousseau, 2015). Thus, listings with an average of 10 reviews likely host 15 to 33 times per year, listings with 5 reviews per year 7 to 16 times and listing with 3 reviews 4.5 to 10 times, which is still higher than 95% of listings on the exchange platform. At the same time, 85% exchange users rate their visits. Hence, the personal residence criterion is more conservative for Airbnb than the platform.

[39] It is possible that the share of non-professionals is larger in the country-side than in the largely urban destinations covered by the InsideAirbnb data. Yet, the share of plausible personal residences is actually smaller in Ireland and New Zealand, where the InsideAirbnb data cover the entire country. Furthermore, guests are likely more committed to rate casual hosts than professionals (Proserpio et al., 2018) so that the share of professionals among transactions is presumably even higher than among ratings.

side network effects. Consumers' previous guest experiences increase familiarity with the platform and encourages their participation as hosts. Third, Airbnb hosts receive monetary compensation instead of internal currency. The latter is far less fungible, because it can only be spent on home rentals on a single platform. Furthermore, search costs on the exchange platform platform are likely considerably higher than on Airbnb. The average acceptance rate is less than 25 percent and thus substantially lower than on Airbnb. Most users rent out their properties only once or twice a year. This infrequent usage likely leads to hosts checking their messages less often and not regularly updating their availability calendars, which can result both in outdated vacancies and slower response times. These increased search costs further reduce the fungibility of internal currency.

To measure the popularity of both platforms among personal residences I compare the absolute numbers of residences applying the definition from before. Figure 12 shows that around 112,000 rated Airbnb listings in the 107 destinations satisfy the most generous definition of personal residences, allowing for 10 ratings per year on average. There are about a third as many exchange platform listings satisfying the same criterion. Using the more plausible threshold of three ratings per year (compatible with hosting 4.5 to 10 times), there are about 61,000 rated Airbnb listings and 27,000 rated exchange listings. Inspecting some of the selected Airbnb listings suggests that the even this criterion includes a significant number of professional or semi-professional rentals,[40] presumably because many Airbnb guests do not leave ratings. Providing a hotel-like experience in their personal residence may be particularly costly for families. Hence, I also consider several proxies for "family homes". Figure 12 shows that there are more personal residences with a garden on exchange platform than on Airbnb. The same is true for listings with dishwashers, listings with three or more rooms and listings with some facilities for children (toys or beds for children or babys). It is worth noting that Inside Airbnb covers only locations with substantial Airbnb activity and thus likely has a high market share.

Consistent with the dichotomy between personal residences and professional holiday rentals, table C1 shows that homes on the exchange platforms generally have much more equipment and amenities typical for personal residences, such as a garden, balcony, dishwasher, dryer and a bathtub. Interestingly, Airbnb listings are more likely to have a pool and an air-condition - consistent with trying to offer a more hotel-like experience.

These results suggest that the exchange platform platform has a surprisingly large market

[40]For instance, a listing displaying a "No. 1 wedding villa" award.

share in the market for personal residences and especially for family homes.

7 Conclusion

Key findings This paper examines if platforms with an internal currency should try to clear the market or instead compress the spread of prices and tolerate excess demand. I first demonstrate with a simple theoretical model that compressed prices can increase the supply of popular goods, total trade and user welfare. The key insight is that compressed prices exert large income effects and prevent owners of highly demanded assets from becoming easily satiated. Using comprehensive data from one of the world's largest internal currency systems, I then confirm the predictions of the model. First, I demonstrate large income effects. I then show that reducing the price of attractive homes indeed increased their supply. I do not find evidence of sizable side-effects of compressing prices. There are no signs of increased misallocation and at most very small effects on hosts' effort. Participation is not affected either. Hence, users may be better off without market-clearing prices.

Implications for internal currency systems My findings have important practical implications for the operation of internal currency systems. Thanks to smartphone apps and examples of very successful tech companies these systems are becoming more and more common. While price-setting in settings with monetary transfers has been extensively studied (Bergemann and Välimäki, 2019), research on pricing with internal currencies is still very limited. A broader insight of my paper is that lessons from traditional markets may not easily extend to markets without real money. Existing research has typically studied market-clearing approaches and actively helped design such systems (Budish, 2011; Budish et al., 2017; Prendergast, 2017, 2022). Yet, internal currency systems are currently using various pricing approaches ranging from market-clearing prices to full compression with a single price for all goods. What kind of prices are best is thus a crucial live question. When the platform collects a commission per transaction, as some platforms do, its incentives are aligned with maximizing trade. With a membership model, the prices maximizing revenue might not exactly coincide with maximizing trade or user welfare, but plausibly, it could still be in the platform's interest to leverage some of the large additional supply of attractive assets that compressed prices can generate. Internal currency systems are also discussed in a variety of institutional settings, such as exchanging kidney donations among hospitals, for worker and student exchange and in education. Here, the designer is often interested in increasing participants' welfare.

Internal currency systems in other contexts My results highlight that satiation can be a major issue in practice and creates the large income effects we observe here. To the best of my knowledge this the first paper demonstrating the phenomenon empirically on a large platform. Ashlagi et al. (2024) show theoretically that the distribution of agents' wealth does not diverge when the market is not extremely thin. I find in a very large and relatively thick market that even moderate levels of wealth accumulation can create under-supply.[41] Satiation could arise in other settings where token money is used and has been a concern, for instance, for kidney exchanges among hospitals (Ashlagi and Roth, 2021; Ashlagi et al., 2024). Income effects are likely stronger the more limited agents opportunities to spend token money are. This depends both on how much agents can consume per period and how farsighted they are. In the context of home exchanges, most individuals can only do a small number of holidays per year and may not want to save tokens for the too distant future. In the case of Feeding America, most of the large regional food banks seem to be quite patient and can often absorb sizable amounts of food (Altmann, 2022). Even under these favorable circumstances, however, some persistent imbalances emerge, which are mitigated only by a very comprehensive taxation scheme. My findings show that compressing prices can be a workable solution. While this paper has focused on the high-level question of using market-clearing prices or not, future research should investigate which exact prices are optimal. Structural models, such as the one of Altmann et al. (2024), allow researchers to simulate varying degrees of price compression along different dimensions, for example location and home size.

Broader implications Another notable finding is that many non-professional hosts appear reluctant to offer their personal residences on for-money rental platforms like Airbnb and are willing to forgo hundreds and sometimes thousands of dollars. This creates a somewhat separate market for personal residence rentals, limits competition with commercial vacation rentals and gives the platform some pricing power. Further investigating users' motives to use various platforms without real money is one interesting avenue for future research. The findings suggest that platforms without real money

[41] In their model it is sufficient that two agents can supply any service that another agent demands. On the exchange platform there are typically dozens and often hundreds of homes available at a given location and month.

are more likely than for-money platforms to facilitate the sharing of idle resources that was originally envisioned for the sharing economy. Together with the fact that platforms with token money can successfully operate at quite large scale, this has potential implications for regulating supposed "peer-to-peer" platforms. Finally, much market design research has focused on contexts where money is “repugnant” or prohibited (Roth, 2007). The results here indicate that market designs without real money could be attractive in a wider range of settings, suggesting exciting opportunities to explore further non-traditional market designs.

References

Agarwal, N., Ashlagi, I., Azevedo, E., Featherstone, C. R., and Karaduman, Ö. (2019). Market Failure in Kidney Exchange. *American Economic Review*, 109(11):52.

Akbarpour, M., Combe, J., He, Y., Hiller, V., Shimer, R., and Tercieux, O. (2022). Unpaired Kidney Exchange: Overcoming Double Coincidence of Wants without Money. Working Paper.

Almagro, M. and Domínguez-Iino, T. (2024). Location sorting and endogenous amenities: Evidence from amsterdam. Technical report, National Bureau of Economic Research.

Altmann, S. M. (2022). Choice, Welfare, and Market Design: An Empirical Investigation of Feeding America's Choice System. Job Market Paper.

Altmann, S. M., Goedde, J., and Wren-Lewis, L. (2024). A structural model of internal currency systems with endogenous supply.

Ashlagi, I. and Kerimov, S. (2020). Scrip Systems with Minimal Availability. Working Paper.

Ashlagi, I., Kerimov, S., and Tamuz, O. (2024). The Power of Two in Token Systems. *arXiv preprint arXiv:2405.12414*.

Ashlagi, I. and Roth, A. E. (2021). Kidney exchange: An operations perspective. *Management Science*, 67(9):5455–5478.

Baldwin, E., Jagadeesan, R., Klemperer, P., and Teytelboym, A. (2022). The Equilibrium Existence Duality. *Journal of Political Economy*.

Barron, K., Kung, E., and Proserpio, D. (2021). The Effect of Home-Sharing on House Prices and Rents: Evidence from Airbnb. *Marketing Science*, 40(1):23–47.

Basu, K. and Van, P. H. (1998). The Economics of Child Labor. *The American Economic Review*, 88(3):412–427.

Bergemann, D. and Välimäki, J. (2019). Dynamic Mechanism Design: An Introduction. *Journal of Economic Literature*, 57(2):235–274.

Bharadwaj, P., Lakdawala, L. K., and Li, N. (2020). Perverse consequences of well intentioned regulation: Evidence from India's child labor ban. *Journal of the European Economic Association*, 18(3):1158–1195.

Bibler, A. J., Teltser, K. F., and Tremblay, M. J. (2021). Inferring tax compliance from pass-through: Evidence from Airbnb tax enforcement agreements. *Review of Economics and Statistics*, 103(4):636–651.

Brousseau, F. (2015). Analysis of the impact of short-term rentals on housing. Technical report, Budget and Legislative Analyst's Office.

Budish, E. (2011). The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103.

Budish, E., Cachon, G. P., Kessler, J. B., and Othman, A. (2017). Course Match: A Large-Scale Implementation of Approximate Competitive Equilibrium from Equal Incomes for Combinatorial Allocation. *Operations Research*, 65(2):314–336.

Calonico, S., Cattaneo, M. D., Farrell, M. H., and Titiunik, R. (2017). Rdrobust: Software for Regression-discontinuity Designs. *The Stata Journal: Promoting communications on statistics and Stata*, 17(2):372–404.

Calonico, S., Cattaneo, M. D., and Titiunik, R. (2014). Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs: Robust Nonparametric Confidence Intervals. *Econometrica*, 82(6):2295–2326.

Che, Y.-K., Gale, I., and Kim, J. (2013). Assigning resources to budget-constrained agents. *Review of Economic Studies*, 80(1):73–107.

Chen, W., Wei, Z., and Xie, K. (2022). Regulating Professional Players in Peer-to-Peer Markets: Evidence from Airbnb. *Management Science*.

Demir, E. and Emekli, G. (2021). Is Airbnb no longer a sharing economy platform? Evidence from Europe's top 10 Airbnb destinations. *Anatolia*, 32(3):470–488.

Duso, T., Michelsen, C., Schaefer, M., and Tran, K. D. (2022). Airbnb and Rental Markets: Evidence from Berlin.

Dworczak, P., Kominers, S. D., and Akbarpour, M. (2021). Redistribution Through Markets. *Econometrica*, 89(4):1665–1698.

Echenique, F., Miralles, A., and Zhang, J. (2021). Constrained pseudo-market equilibrium. *American Economic Review*, 111(11):3699–3732.

Ellison, G. and Ellison, S. F. (2024). Effects of Home Rental Sites on Residential Real Estate: Evidence from New Hampshire.

Fack, G., Goedde, J., and Wren-Lewis, L. (2024). Misplaced trust? Analyzing property owners' responses in an online market.

Farronato, C. and Fradkin, A. (2022). The Welfare Effects of Peer Entry: The Case of Airbnb and the Accommodation Industry. *American Economic Review*, 112(6):1782–1817.

Filippas, A. and Horton, J. J. (2023). The Tragedy of Your Upstairs Neighbors: The Negative Externalities of Home-Sharing Platforms. Technical report.

Filippas, A., Horton, J. J., and Zeckhauser, R. J. (2020). Owning, Using, and Renting: Some Simple Economics of the "Sharing Economy". *Management Science*.

Filippas, A., Jagabathula, S., and Sundararajan, A. (2022). The Limits of Centralized Pricing in Online Marketplaces and the Value of User Control. Technical report, Working Paper.

Fradkin, A. (2017). Search, Matching, and the Role of Digital Marketplace Design in Enabling Trade: Evidence from Airbnb. Technical report.

Fradkin, A., Grewal, E., and Holtz, D. (2018). The determinants of online review informativeness: Evidence from field experiments on Airbnb. Working Paper.

Gaineddenova, R. (2022). Pricing and Efficiency in a Decentralized Ride-Hailing Platform. Job Market Paper.

Garcia-López, M.-À., Jofre-Monseny, J., Martínez-Mazza, R., and Segú, M. (2020). Do short-term rental platforms affect housing markets? Evidence from Airbnb in Barcelona. *Journal of Urban Economics*, 119:103278.

Glaeser, E. L. and Luttmer, E. F. (2003). The misallocation of housing under rent control. *American Economic Review*, 93(4):1027–1046.

Gyódi, K. (2019). Airbnb in European cities: Business as usual or true sharing economy? *Journal of Cleaner Production*, 221:536–551.

Hylland, A. and Zeckhauser, R. (1979). The Efficient Allocation of Individuals to Positions. *Journal of Political Economy*, 87(2):293–314.

Jackson, M. O. and Sonnenschein, H. F. (2007). Overcoming incentive constraints by linking decisions 1. *Econometrica*, 75(1):241–257.

Jagadeesan, R. and Teytelboym, A. (2022). Matching and prices. Technical report, Working paper, Stanford University.

Klein, M., Zhao, J., Ni, J., Johnson, I., Hill, B. M., and Zhu, H. (2017). Quality Standards, Service Orientation, and Power in Airbnb and Couchsurfing. *Proceedings of the ACM on Human-Computer Interaction*, 1(CSCW):1–21.

Kornbluth, D. and Kushnir, A. (2021). Undergraduate Course Allocation through Competitive Markets.

Koster, H. R. A., van Ommeren, J., and Volkhausen, N. (2021). Short-term rentals and the housing market: Quasi-experimental evidence from Airbnb in Los Angeles. *Journal of Urban Economics*, 124:103356.

Krugman, P. (1994). *Peddling Prosperity: Economic Sense and Nonsense in the Age of Diminished Expectations.* WW Norton & Company.

Kurino, M. (2014). House allocation with overlapping generations. *American Economic Journal: Microeconomics*, 6(1):258–289.

Miralles, A. and Pycia, M. (2021). Foundations of pseudomarkets: Walrasian equilibria for discrete resources. *Journal of Economic Theory*, 196:105303.

Nickell, S. (1981). Biases in dynamic models with fixed effects. *Econometrica: Journal of the Econometric Society*, pages 1417–1426.

Nicolò, A., Salmaso, P., and Saulle, R. (2023). Dynamic One-Sided Matching. SSRN Scholarly Paper 4352130, Rochester, NY.

Prendergast, C. (2017). How food banks use markets to feed the poor. *Journal of Economic Perspectives*, 31(4):145–162.

Prendergast, C. (2022). The Allocation of Food to Food Banks. *Journal of Political Economy*, 130(8):000–000.

Proserpio, D., Xu, W., and Zervas, G. (2018). You Get What You Give: Theory and Evidence of Reciprocity in the Sharing Economy. *Quantitative Marketing and Economics*, 16(4):371–407.

Quattrone, G., Nocera, A., Capra, L., and Quercia, D. (2020). Social Interactions or Business Transactions?What customer reviews disclose about Airbnb marketplace.

In *Proceedings of The Web Conference 2020*, WWW '20, pages 1526–1536, Taipei, Taiwan. Association for Computing Machinery.

Rambachan, A. and Roth, J. (2023). A more credible approach to parallel trends. *Review of Economic Studies*, 90(5):2555–2591.

Roberts, R. (2015). Alvin Roth on Matching Markets (EconTalk).

Roth, A. E. (2007). Repugnance as a Constraint on Markets. *Journal of Economic perspectives*, 21(3):37–58.

Sweeney, J. and Sweeney, R. J. (1977). Monetary theory and the great Capitol Hill Baby Sitting Co-op crisis: Comment. *Journal of Money, Credit and Banking*, 9(1):86–89.

Weitzman, M. L. (1977). Is the Price System or Rationing More Effective in Getting a Commodity to Those Who Need it Most? *The Bell Journal of Economics*, 8(2):517–524.

Figures and Tables

Figure 1: Preference regimes

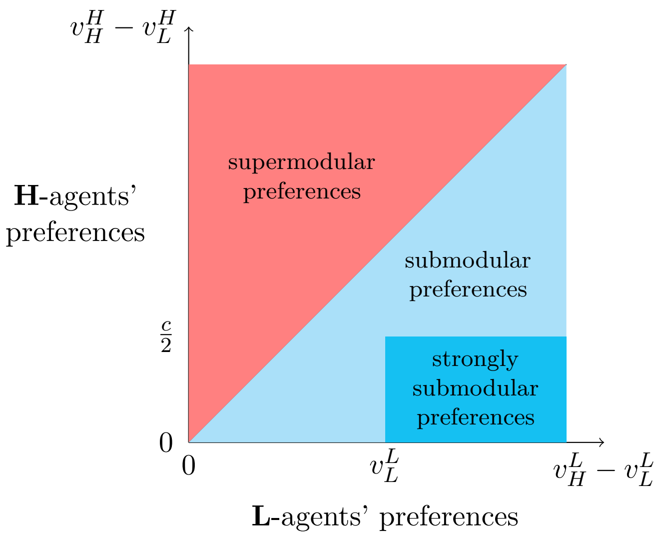

The figure summarizes the three preference regimes defined in section 2.3.1. If preferences are supermodular, H-agents have a stronger preference for H-goods than L-agents. If preferences are submodular, the opposite is true. If preferences are strongly submodular, L-agents have a substantially stronger preference for H-homes than H-agents.

Figure 2: Exchange patterns in competitive equilibrium

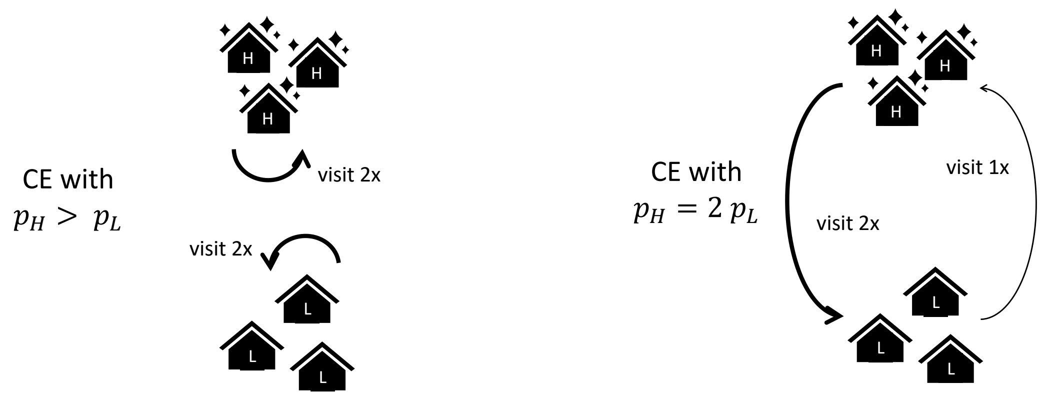

The figure summarizes proposition 1. When preferences are not SSB (left), the type H agents always visit type H homes. When preferences are SSB (right), H agents host once and visit an L homes twice.

Figure 3: Prices in competitive equilibrium

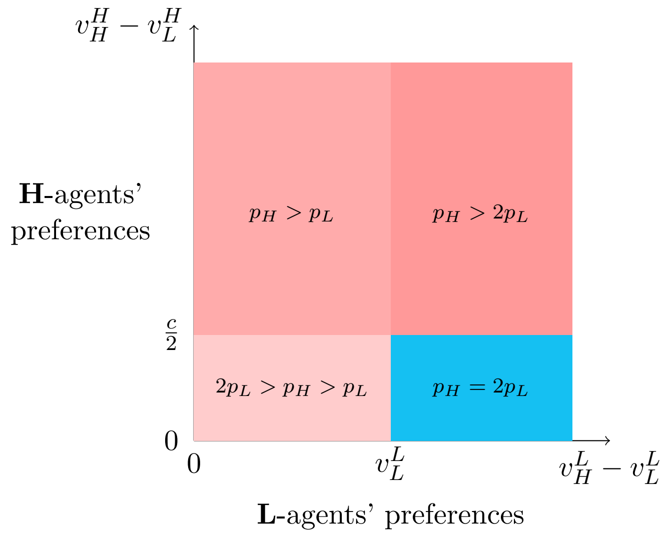

The figure shows the prices supporting the competitive equilibria characterized in proposition 1. When preferences are SSB (blue area), the unique market-clearing price is $p_H = 2p_L$. When preferences are not SSB, there is a set of prices which clears the market and supports the unique competitive equilibrium allocation. The proof is in Appendix A.1.3.

Figure 4: Exchange patterns in uniform rationing equilibrium

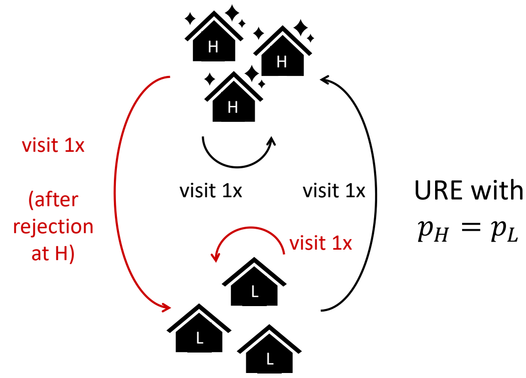

The figure summarizes proposition 2. When prices are equal, all types demand H homes. Demand exceeds supply and is rationed uniformly at random. Each type manages to obtain an H home half the time and visits an L home if rejected.

Figure 5: Summary of Theorem 1

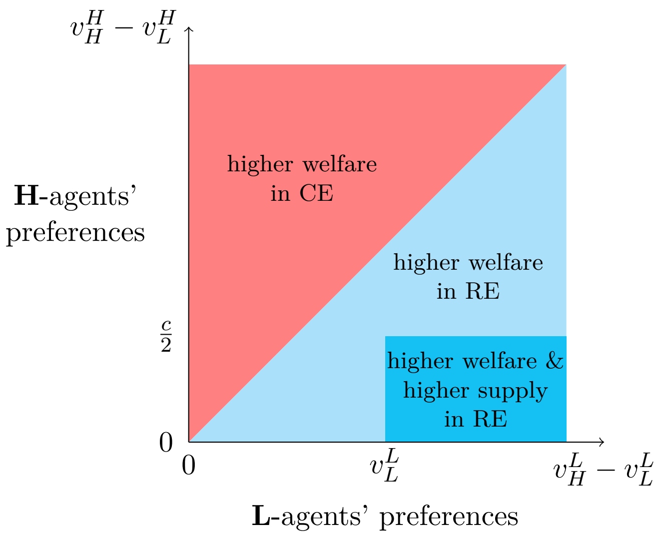

The figure summarizes Theorem 1. The effect of price compression on total welfare is ambiguous. If preferences are supermodular, welfare in the competitive equilibrium is higher than in the rationing equilibrium (red regime). If preferences are submodular, compression increases welfare. In the dark blue regime, when preferences are strongly submodular, compression increases the supply of **H** type homes and welfare.

Figure 6: Effect of wealth on times hosted in quarter

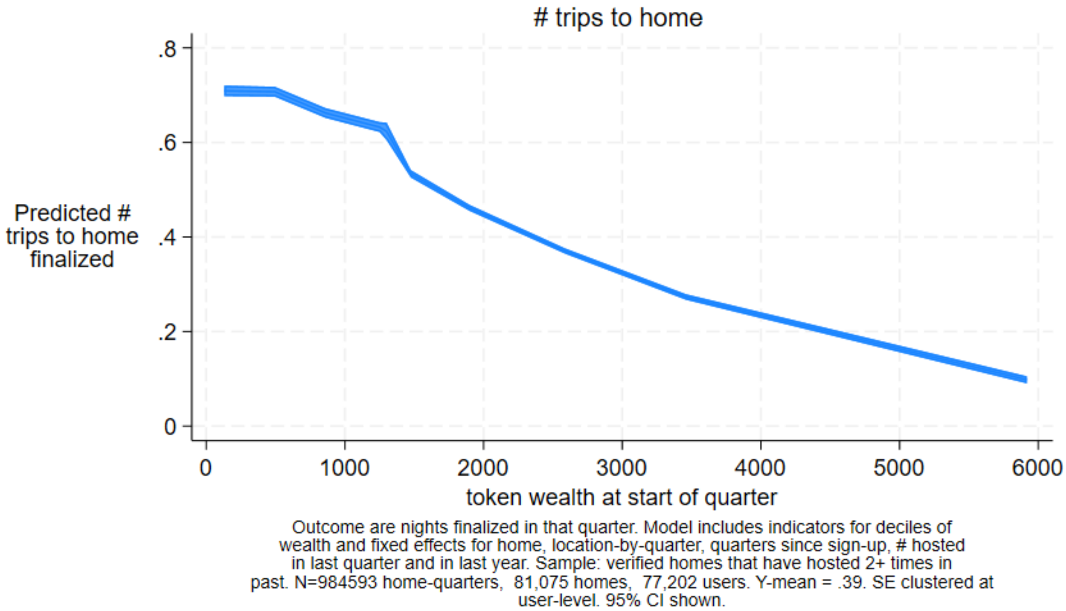

Figure D.16 shows the corresponding results for the number of nights finalized.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	# hosted	nights	# hosted	nights	host 3+	cold calls	reciprocal
wealth (1000s) at start of quarter	-65.02***	-70.48***	-49.40***	-54.06***	-90.91***	-61.23***	13.06***
	(0.680)	(0.763)	(0.668)	(0.744)	(1.659)	(0.941)	(0.744)
Benchmark controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
FE lagged hostings	No	No	Yes	Yes	Yes	Yes	Yes
Mean of Y	.46	2.6	.39	2.3	.034	.16	.086
R^2	0.38	0.32	0.34	0.29	0.22	0.27	0.23
N	1,048,925	1,048,925	984,575	984,575	984,575	984,575	984,575
Number of hosts	81,217	81,217	77,201	77,201	77,201	77,201	77,201

SE clustered at user-level in parentheses. Effects are normalized as percent of the outcome mean. All models include fixed effects for home, location-by-quarter and quarters since sign-up. Some model additionally include fixed effects # hosted in last quarter and in last year. In column 'cold calls' the outcome is the number of hostings finalized for periods not marked as available in the calendar. Reciprocal are 2-way exchanges which are typically done without currency transfers. All other columns include both 1-way and 2-way hostings. Sample: verified homes that have already hosted 2+ times.

Table 1: Effect of wealth on supply

	(1)	(2)	(3)	(4)	(5)	(6)
	wealth	wealth	wealth	# hosted	# hosted	# hosted
RD_Estimate	133.5*	143.2*	168.5**	-0.0212*	-0.0227*	-0.0198
	(55.55)	(56.02)	(60.66)	(0.00932)	(0.0109)	(0.0178)
Bandwith left	10	8	10	10	8	10
Bandwith left	10	8	10	10	8	10
Polynomial	1	1	2	1	1	2
Optimal polynomial	Yes	Yes		Yes	Yes	
Optimal bandwidth	Yes		Yes	Yes		Yes
Number of events	6927	6927	6927	6927	6927	6927
Number of hosts	5784	5784	5784	5784	5784	5784
Y mean	2985	2985	2985	.129	.129	.129

Columns labeled *wealth* show the estimated effect on IC wealth. Columns labeled *# hosted* refers to the number of 1-way hostings finalized in a given month. Bandwidth and polynomial are automatically selected using the default mserd method and triangular kernel of the *rdrobust*. Standard errors are clustered at the user level.

Table 2: Regression discontinuity estimates

Figure 7: Effect of unexpected internal currency wealth and supply

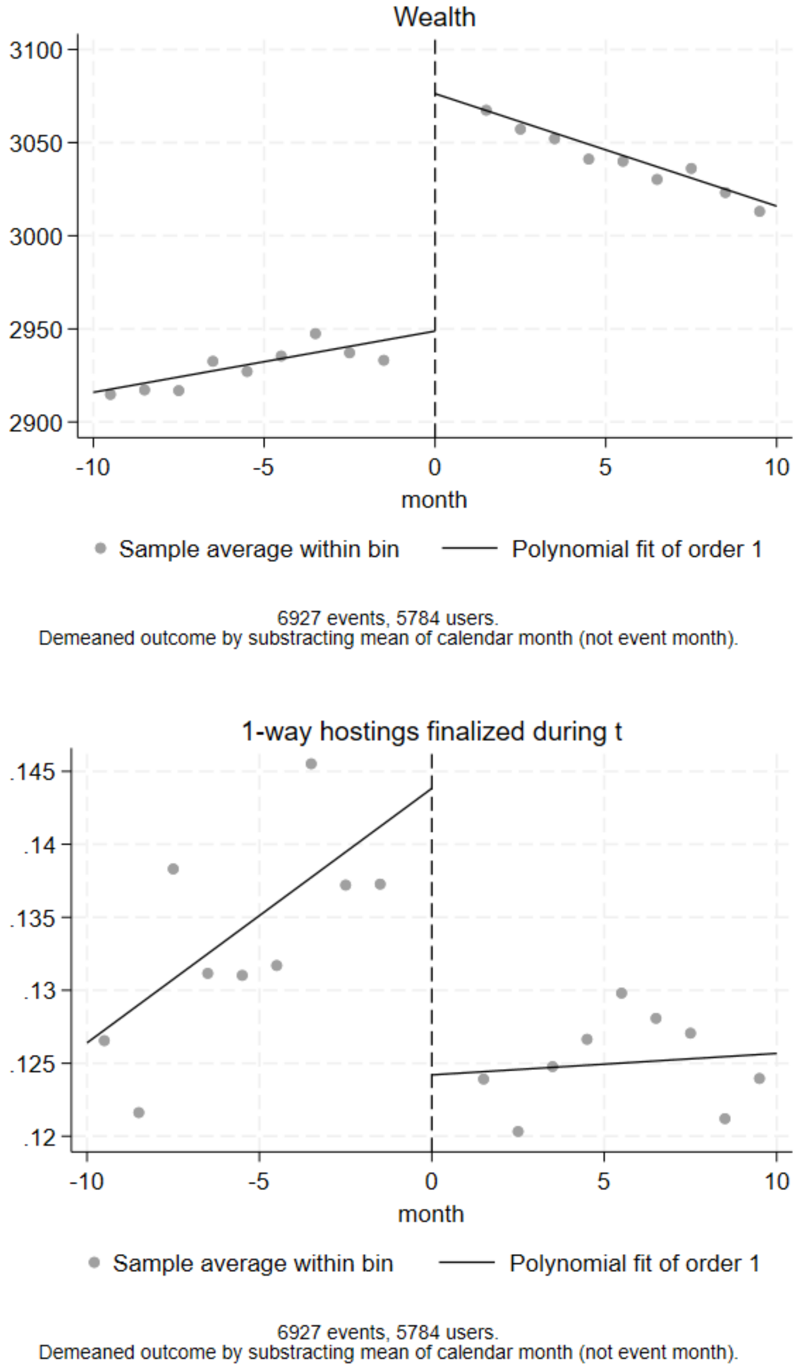

The figure shows users' wealth and supply around the time when a previously referred "godson" pays their membership fee, triggering a bonus for the referring user. Supply are 1-way hostings finalized in a given month. Bandwidth and polynomial are automatically selected using the default mserd method and triangular kernel of the *rdrobust* package introduced by Calonico et al. (2017).

Figure 8: Price-changes during reform

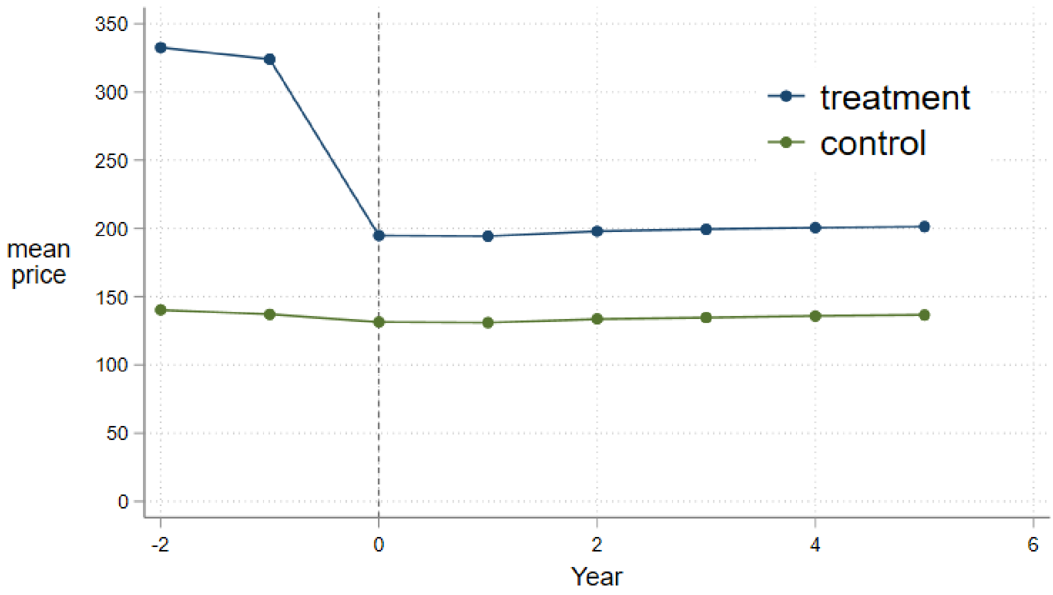

Outcome is the mean recommended price. Control and Treatment Groups: 66,183 homes with change in (-10%,+10%), 41,653 homes with change < -30%. Treatment := recommended price of own home dropped by at least 30%. Sample: all homes that are >=85% complete and have 1-15 beds .

	(1) nights	(2) nights	(3) # hosted	(4) # hosted	(5) non-recipr.	(6) reciprocal	(7) cold calls
DiD estimate	0.434***	0.415***	0.0707***	0.0676***	0.275***	0.140***	0.242***
	(0.0540)	(0.0552)	(0.00916)	(0.00950)	(0.0450)	(0.0257)	(0.0358)
Location-year fixed effects	no	yes	no	yes	yes	yes	yes
Number of hosts	97,024	96,619	97,024	96,619	96,619	96,619	96,619
Effect in % of Y-mean	22	21	20	19	18	28	28
Mean of Y	2	2	.35	.35	1.5	.49	.88

SE clustered at user-level in parentheses. All models include fixed effects for determinants of price change. In column *cold calls* the outcome is the number of hostings finalized for periods not marked as available in the calendar. Reciprocal refers to 2-way exchanges which are typically done without currency transfers. All other columns include both 1-way and 2-way hostings.

Table 3: Difference-in-difference estimates of supply effect

Figure 9: Effect of price-reform on supply (nights)

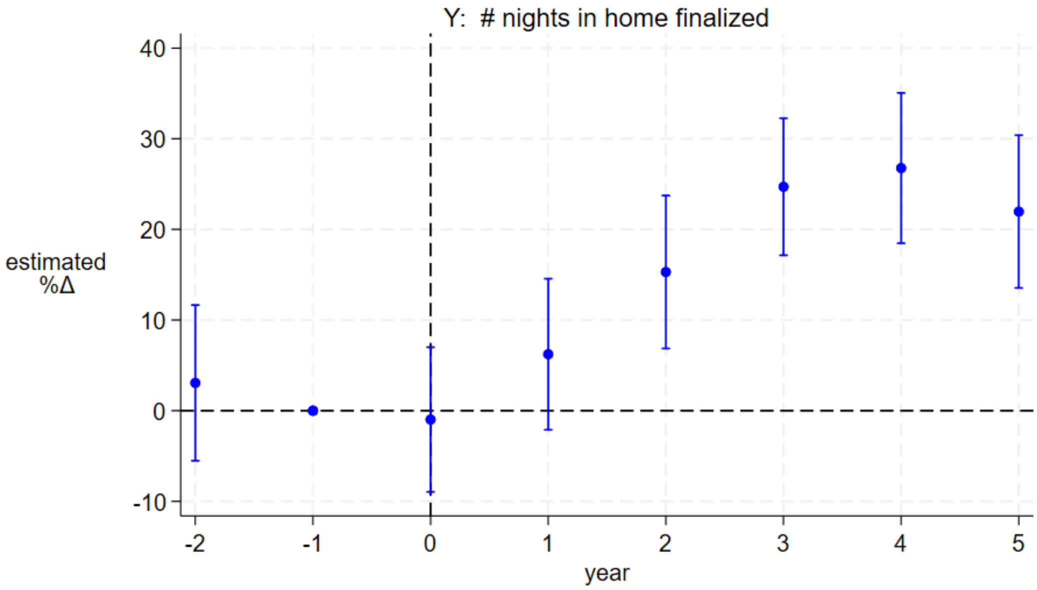

The graphs shows difference-in-difference event-study estimates (equation 38) of the effect of the price reduction on treated homes. The outcome is the number of nights in the home finalized in that year. Prices of homes in the treatment group were reduced by at least 30%. Prices of homes in the control group were reduced by $(-10\%, +10\%)$. The model includes fixed effects for time and for cells of home-types, which determine treatment status. Estimates divided by the outcome mean and multiplied by 100. Standard errors are clustered at the user level.

	(1) active	(2) hosting after	(3) # referred (sign up)	(4) # referred (home)
DiD estimate	0.0325***	0.0511***	-0.00459	0.000479
	(0.00427)	(0.00660)	(0.00744)	(0.00265)
Location-year fixed effects	yes	yes	yes	yes
Number of hosts	96,619	96,619	96,619	96,619
Effect in % of Y-mean	4.6	15	-4.8	1.1
Mean of Y	.7	.33	.096	.042

SE clustered at user-level in parentheses. All models include fixed effects for home-type cells that determine price change in reform. A user is considered *active* in a year t if they send any request or message in or after that year. *Hosting after* is defined analogously based on hostings finalized in or after that year. Column (3) shows for each existing user and year the number of new users, referred by the existing user, who signed up that year. Column (4) shows the number of referrals who completed a home listing in that year.

Table 4: Participation difference-in-difference estimates

	(1) rating received	(2) cleanliness rating	(3) rating given	(4) party size
DiD estimate	-0.0153	-0.0233	0.0106	0.215***
	(0.0101)	(0.0144)	(0.00766)	(0.0219)
Location-year fixed effects	yes	yes	yes	yes
Number of hosts	31,721	31,721	35,594	86,038
Effect in % of Y-mean	-.31	-.48	.22	7
Mean of Y	4.9	4.9	4.9	3.1

SE clustered at user-level in parentheses. All models include fixed effects for home-type cells that determine price change in reform.

Table 5: Potential side-effects - difference-in-difference estimates

Figure 10: Comparison of IC prices and imputed Airbnb prices

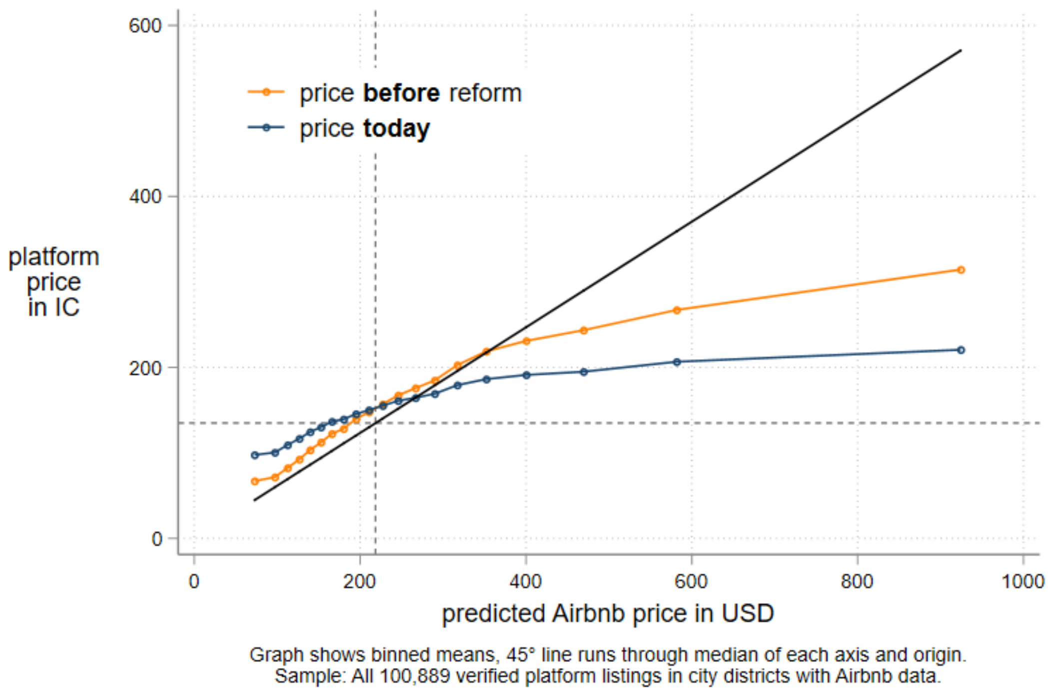

Figure 11: Share of personal residences on each platform

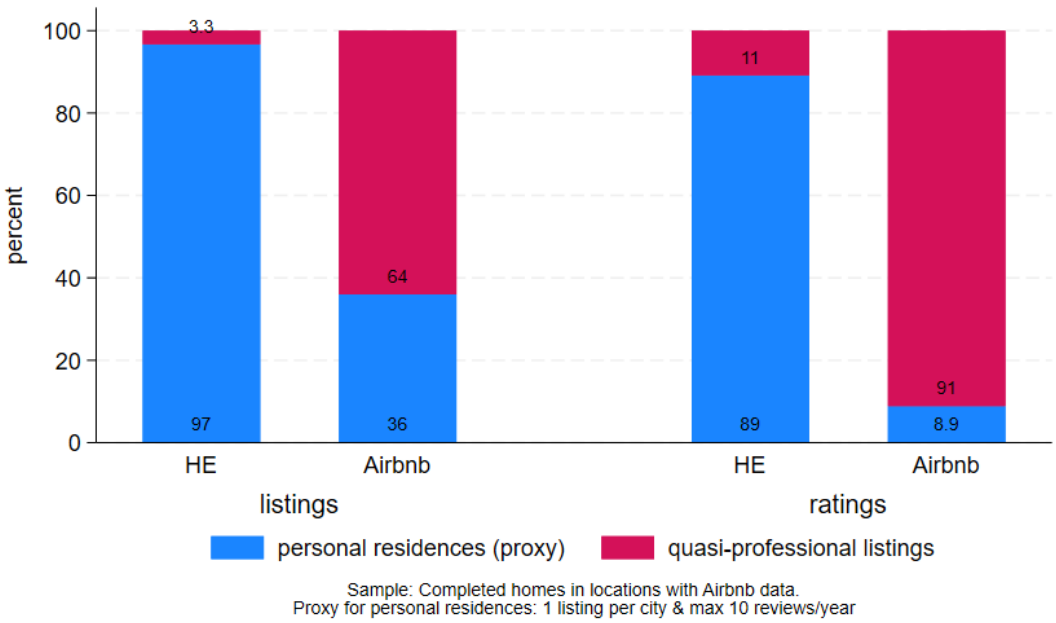

Figure 12: Number of personal residences on each platform

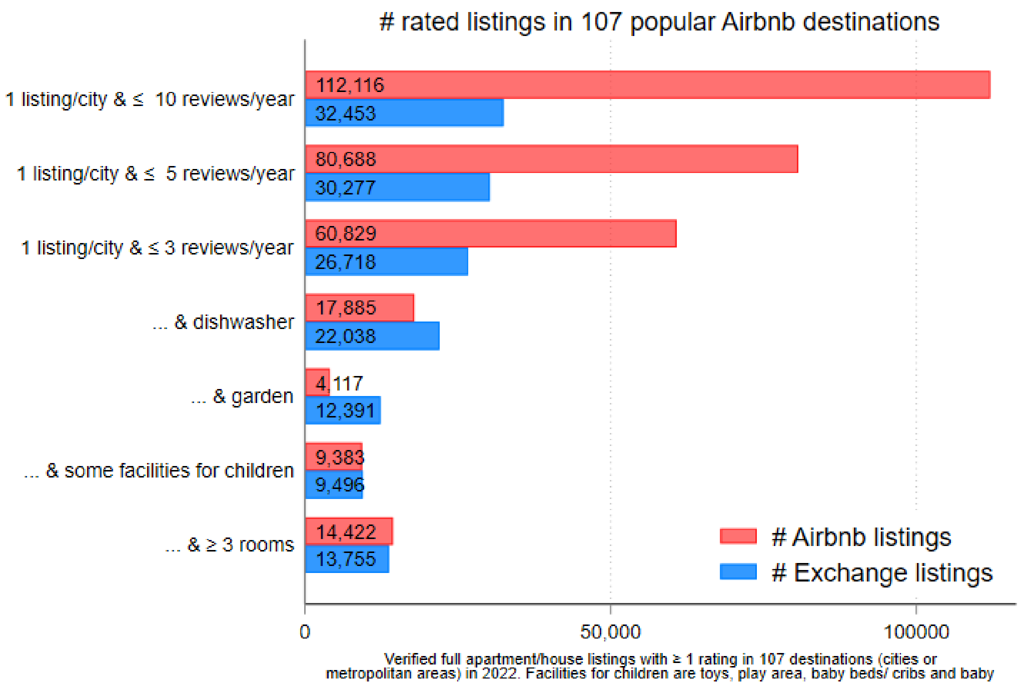

Transactions	**Foregone net Airbnb income ≥**					**# Users**
	100	200	500	1,000	1,500	
Benchmark assumptions						
all homes	0.52	0.48	0.38	0.28	0.21	48,483
homes worth $\geq$ 300 USD/night	0.81	0.77	0.66	0.53	0.44	16,580
Transaction pairs with net IC loss						
all homes	0.30	0.26	0.19	0.12	0.08	17,850
homes worth $\geq$ 300 USD/night	0.54	0.51	0.40	0.27	0.19	6,857

Table 6: Lower bounds of Willingness-to-pay

The table shows the share of users who have forgone net Airbnb income of at least 100 USD, at least 200 USD etc. up to 1,500 USD. Net Airbnb income is net of taxes, fees and imputed expenditure on destination accommodation. The benchmark assumptions include all hostings. In addition to taxes and fees Airbnb income is net of the replacement cost of earned IC (if positive IC were earned), the USD price at which users can buy IC trying to finalize a booking and having insufficient IC balance. Rows 3-4 restrict to transaction pairs where the user hosted and visited in the same week and her net IC revenue were $\leq$ 0 IC . Lost IC are never priced in USD. All scenarios exclude trips where the user visited a home with unknown Airbnb price.

A Proofs and additional theory

A.1 Proofs

Here, we formally proof proposition 1, the characterization of the competitive equilibrium. To start, we derive the individual optimal consumption and supply choices, which in turn yield the aggregate demand and supply functions. These also feed into the proof of proposition 2 and theorem 1.

A.1.1 Optimal individual choices

Without rationing ($q_j^D = q_j^S = 1, a = H, L$) agents' optimal demand and supply choices are

$$(s^L, d_H^L, d_L^L) = \begin{cases} (1,2,0) \text{ if } p_L/2 \geq p_H \\ (2,2,0) \text{ if } p_L/2 < p_H \leq p_L \\ (2,1,0) \text{ if } p_L < p_H \leq 2p_L \text{ and } v_H^L > 2v_L^L \\ (2,0,2) \text{ if } p_H > 2p_L \text{ or } (p_L < p_H \text{ and } v_H^L \leq 2v_L^L) \end{cases} \tag{21}$$

$$(s^H, d_H^L, d_L^L) = \begin{cases} (2,2,0) \text{ if } p_H < 2p_L \text{ or } v_H^H - v_L^H > c/2 \\ (1,0,2) \text{ if } p_H \geq 2p_L \text{ and } v_H^H - v_L^H \leq c/2 \end{cases} \tag{22}$$

To see why this is optimal consider agents' possible choices. Since $v_j^i > c, \forall(i,j)$, $s^i = 0$ cannot be optimal. For the same reason, hosting only once and consuming only one unit is dominated by hosting twice and doubling consumption. Thus, agents host once only if they can nonetheless afford to consume 2 units. This requires $p_i \geq 2p_j$. Type-L agents will demand 2 units of H if they can afford them, i.e. when $p_H \leq p_L$. They prefer $(d_H^L, d_L^L) = (1,0)$ over $(d_H^L, d_L^L) = (0,2)$ iff $v_H^L > 2v_L^L$ and can afford the former iff $p_H \leq 2p_L$. Otherwise, they demand $(d_H^L, d_L^L) = (0,2)$. Type-H agents can always afford $(s^H, d_H^L, d_L^L) = (2,2,0)$. Hence $d_H^L + d_L^L = 1$ is never optimal. They prefer $(s^H, d_H^L, d_L^L) = (1,0,2)$ over $(s^H, d_H^L, d_L^L) = (2,2,0)$ iff $2v_H^H - 2c \leq 2v_L^H - c \Leftrightarrow v_H^H - v_L^H \leq c/2$ and can afford the former iff $p_H \geq 2p_L$.

Rationing does not affect what agents' most preferred bundles are. Thus, agents will try to obtain the bundles above and iff they are not served and the other good is affordable, they then demand the other good. In the following exposition we will not write out the "second-step" demand formally.

A.1.2 Aggregate demand and supply

The individual choices above imply the following aggregate demand and supply functions. Since $v^i_H > v^i_L$ for all i, we restrict attention to cases where $p_H \geq p_L$.[42]

$$S_L(p_H, p_L) = 2 \tag{23}$$

$$S_H(p_H, p_L) = \begin{cases} 2 \text{ if } p_H < 2p_L \text{ or } v^H_H - v^H_L > c/2 \\ 1 \text{ if } p_H \geq 2p_L \text{ and } v^H_H - v^H_L \leq c/2 \end{cases} \tag{24}$$

$$(D_H(p_H, p_L), D_L(p_H, p_L)) =$$

$$\begin{cases} (4,0) \text{ if } d^L_H = 2 \text{ and } d^H_H = 2 \Leftrightarrow p_H = p_L \\ (3,0) \text{ if } d^L_H = 1 \text{ and } d^H_H = 2 \Leftrightarrow \begin{array}{c} (p_L < p_H \leq 2p_L \text{ and } v^L_H > 2v^L_L) \text{ and} \\ (p_H < 2p_L \text{ or } v^H_H - v^H_L > c/2) \end{array} \\ (2,2) \text{ if } d^L_H = 0 \text{ and } d^H_H = 2 \Leftrightarrow \begin{array}{c} (p_L < p_H < 2p_L \text{ and } v^L_H \leq 2v^L_L) \text{ or} \\ (p_H > 2p_L \text{ and } v^H_H - v^H_L > c/2) \text{ or} \\ (p_L < p_H \text{ and } v^L_H \leq 2v^L_L \text{ and } v^H_H - v^H_L > c/2) \end{array} \\ (1,2) \text{ if } d^L_H = 1 \text{ and } d^H_H = 0 \Leftrightarrow p_H = 2p_L \text{ and } v^H_H - v^H_L \leq c/2 \text{ and } v^L_H > 2v^L_L \\ (0,4) \text{ if } d^L_H = 0 \text{ and } d^H_H = 0 \Leftrightarrow \begin{array}{l} (p_H > 2p_L \text{ and } v^H_H - v^H_L \leq c/2) \text{ or} \\ (p_H \geq 2p_L \text{ and } v^H_H - v^H_L \leq c/2 \text{ and } v^L_H \leq 2v^L_L) \end{array} \end{cases} \tag{25}$$

Note that $d^L_H = 2$ and $d^H_H = 0$ would require $2p_L \leq p_H \leq p_L$ and is thus never the case.

A.1.3 Proof of proposition 1

Recall proposition 1. In competitive equilibrium utilitarian welfare and aggregate supply of H are uniquely given by

$$S_H(CE(v, c)) = \begin{cases} 1 & \text{if } SSB \\ 2 & \text{else} \end{cases}$$

[42] $p_L \geq 2p_H$ would actually reduce supply of L homes, which cannot be welfare optimal.

and

$$W(CE(c,v)) = \sum_{i=H,L} \Big(\sum_{j=H,L} x^i_j v^i_j - cy^i \Big)$$

$$= \begin{cases} 2v^H_L + \ v^L_H - 3c & \text{if } SSB \\ 2v^H_H + 2v^L_L - 4c & \text{else} \end{cases}$$

Proof. First, let us determine the set of CE as a function of the preference parameters. Since $S_H(.) \in \{1,2\}$ and market-clearing requires that $D_j \leq S_j, \forall j$, there cannot be CE with $D_H(.) > 2$. Furthermore, market-clearing requires $p_H = 0$ when $D_H < S_H$. This violates the conditions for the case $(D_H(.), D_L(.)) = (0,4)$. Hence, the only two possible CE require $(D_H(.), D_L(.)) \in \{(2,2),(1,2)\}$. The conditions for $(D_H(.), D_L(.)) = (1,2)$ are $p_H = 2p_L$ and $v^H_H - v^H_L \leq c/2$ and $v^L_H > 2v^L_L$. Note that $v^L_H > 2v^L_L \Leftrightarrow v^L_H - v^L_L > v^L_L$. The assumption $v^i_j > c, \forall i,j$ implies in particular that $v^L_L > c$. Therefore, we can rewrite the condition for $(D_H(.), D_L(.)) = (1,2)$ as

$$v^H_H - v^H_L \leq c/2 < c < v^L_L < v^L_H - v^L_L$$

If this condition (SSB) is satisfied, then the unique competitive equilibrium is $p_H = 2p_L$ and $S_H(.) = 1 = D_H(.)$ and $S_L(.) = 2 = D_L(.)$. Equations 24 and 25 imply that the individual choices inducing these aggregate quantities at $p_H = 2p_L$ are optimal and markets clear. In this CE $d^L_H = 1, d^L_L = 0, d^H_H = 0, d^H_L = 2$. As all markets clear $x^i_j = d^i_j$ and $y^i = s^i$. Thus, the matching is negative assortative and does not feature full trade. The resulting welfare is $2v^H_L + \ v^L_H - 3c$. Furthermore, $(D_H(.), D_L(.)) = (2,2)$ is not a CE as all three conditions on preferences in equation 25 for this case are violated. Thus, the CE is unique.

Conversely, if SSB is not satisfied, all competitive equilibria entail $(D_H(.), D_L(.)) = (2,2)$. We know from above that SNSM is a necessary condition for a CE with $D_H(.) = 1$. When SSB is violated, either $v^L_H \leq 2v^L_L$ or $v^H_H - v^H_L > c/2$. If both $v^L_H \leq 2v^L_L$ and $v^H_H - v^H_L > c/2$, then any prices $p_H > p_L$ form a CE, because Hs always demand H and $d^L_H = 0$ unless $p_H = p_L$. Thus, $p_H = p_L$ would induce excess demand for H. If $v^L_H \leq 2v^L_L$ and $v^H_H - v^H_L \leq c/2$, CE requires $p_L < p_H < 2p_L$. Optimal individual choices then induce $D_H(.) = 2 = S_H(.)$ and $D_L(.) = 2 = S_L(.)$. If $v^L_H > 2v^L_L$ and $v^H_H - v^H_L > c/2$, CE requires $p_H > 2p_L$ because $d^L_H = 2$ at any price $p_H \leq 2p_L$ and Hs always demand H. In all three cases, $d_L^L = 2, d_H^L = 0, d_H^H = 2, d_L^H = 0$. Thus the matching is assortative and involves full trade. The resulting welfare is $2v_H^H + 2v_L^L - 4c$.

□

A.1.4 Proof that CE is a special case of URE

Remark 1. *If (p, d, s) is a competitive equilibrium, then (p, d, s) is a uniform rationing equilibrium.*

Proof. Suppose (p, d, s) is a competitive equilibrium. We will show that when applying the rationing rule while imposing the market-clearing conditions of CE, agents' maximization problem in the URE definition coincides with the CE maximization problem. Therefore, (d, s) with the associated (x, y) is also a solution to the URE problem.

Recall agents' problem in the CE

$$\begin{aligned}
&\text{For all agents } i, \\
(d^i, s^i) \in \arg\max_{d^i, s^i} \quad & \sum_t \sum_j d_{jt}^i v_j^i - c s_t^i && (26) \\
\text{s.t.} \quad & \sum_t \sum_j d_{jt}^i p_{jt} \le \sum_t p_{h(i)t} s_t^i && (27) \\
& \sum_j d_{jt}^i \le 1 \text{ for all } t && (28)
\end{aligned}$$

and in the URE

$$\begin{aligned}
&\text{For any agent i and any } TB^i, \\
(d^i, s^i) \in \arg\max_{d^i, s^i} \quad & \sum_t \sum_j d_{jt}^i v_j^i - c s_t^i && (29) \\
\text{s.t.} \quad & \sum_t \sum_j x_{jt}^i p_{jt} \le \sum_t p_{h(i)t} y_t^i && (30) \\
& d_{jt}^i p_{jt} + \sum_j x_{js}^i p_{js} \le \sum_r p_{h(i)r} y_r^i \text{ for all } t \ne s && (31) \\
& \sum_j x_{jt}^i \le 1 \text{ for all } t && (32)
\end{aligned}$$

First note that the objective functions of both maximization problems are identical. Hence, we only need to show that the constraints become equivalent when imposing market-clearing together with the rationing rule.

By definition of CE, the market-clearing conditions hold. Thus, $D_{jt} \leq S_{jt}$ for all t and j and $p_{jt} = 0$ for all t and j such that $D_{jt} < S_{jt}$. By the rationing rule $x^i_{jt} = d^i_{jt}$ both if $D_{jt} < S_{jt}$ and if $D_{jt} = S_{jt}$. Hence, $x^i_{jt} = d^i_{jt}$ for all t and j and for any tie-breaking realization.

Furthermore, $p_{h(i)t} y^i_t = p_{h(i)t} s^i_t$ for all i and t and for any tie-breaking realization. If $D_{h(i)t} = S_{h(i)t}$ then $y^i_t = s^i_t$. If $D_{h(i)t} < S_{h(i)t}$ then $p_{h(i)t} = 0$ by the market-clearing condition and the equality also holds for any i regardless of whether $y^i_t = 1$ or $y^i_t = 0$ and thus independent of the tie-breaking realization.

Since $x^i_{jt} = d^i_{jt}$, the constraint (31) is equivalent to (30), which in turn is equivalent to the CE budget constraint (27) because $p_{h(i)t} y^i_t = p_{h(i)t} s^i_t$. Similarly, the unit consumption constraint (32) reduces to the unit demand constraint (28) because $x^i_{jt} = d^i_{jt}$. As the objective functions of both maximization problems are identical, the two problems coincide.

By definition of (p, d, s) being a CE, (d, s) solves the CE problem for all i and is thus a solution to the URE problem given the rationing rule. Hence, (p, d, s) is a URE. $\square$

A.2 For-money platform as outside option

This section describes a simple extension with a competing for-money platform. Each period, agents can supply on the Exchange platform E or on the for-money platform A but not both. Similarly, they can choose between visiting on one of the platforms and not visiting at all. For simplicity, we assume that agents consider the same home-types H, L on both platforms and that agents receive the same utility from visiting a type j home on either platform. By contrast, the disutility of hosting differs across platforms. As before, hosting on platform E creates disutility c. Hosting on A creates disutility γc with $\gamma > 0$. Section 6.3 discusses this assumption. The upshot is that guests might behave better on platform E than A since all guests host themselves and are expecting a personal residence rather than a professional listing. In this case we expect $\gamma > 1$. Formally, hosting is denoted by $y^i_{tk} \in \{0, 1\}$ for $k \in \{E, A\}$ and visiting by $x^i_{jtk} \in \{0, 1\}$ for $k \in \{E, A\}$ and $j \in \{H, L\}$. x^i and y^i now additionally collect hostings and guestings on platform A. Hostings and guestings on the exchange platform need to satisfy the same token-money budget constraint as before. Hostings and guestings on platform A are paid for with real money at prices p_{iE}. Agents' utility is linear in money with slope α_i for type i agents. Hence, agents' individual optimization problem becomes the following.

$$\max \quad u(x^i, y^i) = \sum_{t=1,2} \Bigg(\sum_{j=H,L} v_j^i(x_{jtE}^i + x_{jtA}^i) - c(y_{tE}^i + \gamma y_{tA}^i)$$
$$+ \alpha_i(p_{iA} y_{tA}^i - \sum_{j=H,L} p_{jA} x_{jtA}^i) \Bigg) \tag{33}$$
$$\text{s.t.} \quad \sum_{t=1,2} \sum_{j=H,L} x_{jt}^i p_{jE} \leq \sum_{t=1,2} y_t^i p_{iE} \tag{34}$$
$$x_{jtE}^i + x_{jtA}^i \leq 1, \quad \forall t \tag{35}$$
$$y_{tE}^i + y_{tA}^i \leq 1, \quad \forall t \tag{36}$$

It is easy to show that for any (finite) values of α_i, p_{HA} and p_{LA} there exists a threshold $\bar{\gamma}$ such that neither type hosts on platform A if $\gamma \geq \bar{\gamma}$. In this case, all previous results go through. Section 5.3 shows empirically that increased price compression does not materially reduce participation on the exchange platform and the willingness-to-pay estimates in section 6.2 suggest that for many platform users the costs of hosting on Airbnb are hundreds of dollars higher than the costs of hosting on the exchange platform. Hence, high values of γ seem plausible in practice.

B Additional evidence on mechanism

B.1 Evidence that wealth reduces approval of requests

Studying hosts' decision to accept or reject requests provides further evidence that the negative effect of wealth is driven by the host. I find that a given host is much less likely to approve a given request (controlling for almost all observable characteristics) the wealthier she is at the time of the request. Formally, let $A_{ijts} = 1$ if host i approves guest j's request to visit (home h) in week s, sent at time t. My benchmark specifications resembles equation (18) and takes the following form:

$$A_{ijts}^* = \mu_i + \sum_{d=2}^{10} \beta_d \mathbb{1}(D_{d-1} < Wealth_{it} \leq D_d) + X_{ijts}\theta + \varepsilon_{ijst} \tag{37}$$

X_{ijts} are observable exchange characteristics, which include characteristics of the trip (e.g. dates, number of guests), the request (time asked in advanced, length of the message), the home (location, # beds), the guest and the host (e.g. ratings, ID verification, time on the platform) as well as the pair (e.g. indicators of shared languages and differences in age and home price). Figure D.21 shows the full list of variables and how predictive they are of acceptance in a flexible machine learning model. In this model, host's wealth at the time of a request as one of the strongest predictors of approval. As the machine learning model is hard to combine with (high-dimensional) host fixed-effects, I estimate equation 37 as a linear regression including flexible splines for the continuous variables. Figure D.20 shows that the probability that a given host accepts a request with given characteristics decreases monotonically from ca. 22% at almost 0 wealth to only 5% at 8000 IC- a reduction by more than four times.

B.2 Evidence that wealth reduces calendar availability

An alternative measure of supply is how many nights a host marks her home as available in the calendar. This measure has two drawbacks, however. First, almost half of all hostings result from users agreeing to host in periods not marked in their calendar. Second, data on historical calendar postings are limited and we do not observe the posting date of calendar entries that were subsequently booked. Therefore, we can only study the effect of wealth on nights offered in the calendar for *trips* in quarter q rather than nights *posted* in q for trips at arbitrary future dates. Figure D.19 shows the results of estimating a specification like equation 17 with four lags of wealth. Wealth in quarter q appears to increase the supply of nights for trips in q, but this very likely reflects reverse causality: many trips in q are booked before q and thus increase wealth in q. By contrast, the lagged coefficients behave as expected. Higher wealth in any of the previous four quarters significantly decreases calendar supply.[43] This confirms that the reduced number of realized nights indeed reflects a reduced willingness to host and not a demand response. There was little reason to suspect a demand response to begin with. As guests do not observe hosts' wealth (and trying to infer their wealth from ratings is very hard and certainly much more effort than sending an informal request) it seems very unlikely that guests would send fewer requests to IC-rich hosts.

B.3 Non-assortative preferences

Theorem 1 implies that if we observe a supply effect, preferences cannot be supermodular. Hence, the observed increase in supply already provides strong evidence against supermodular preferences. To further support the coherence of the theory, I nonetheless try to obtain more direct evidence on the preference structure.

[43] It is hard to meaningfully interpret the magnitudes as even the lagged coefficients likely reflect some reverse causality and are thus upward biased. Thus the true effect of wealth is likely more negative.

A simple reduced-form approach is to compare the requests that users with different types of homes make. The model was not explicit about what constitutes the type of a home. As noted before, the most plausible violation of supermodular preferences is in the location dimension. I rank locations using the success rate measure explained in section C.1. We want to know if users with homes in more attractive locations (with higher rejection rates) request more attractive destinations.[44] An important issue is that agents in more attractive locations have a higher ability to pay because they earn more IC per night and might receive more requests. Ideally, we want to compare preferences holding ability to pay fixed. To limit the impact of ability to pay I only consider the first request that agents ever make because most of their budget consists of the initial IC credit that is common across users. This does not completely eliminate the effect of ability to pay, because agents are presumably more reluctant to book a house that costs 260IC per night when knowing that their own home will only generate 90IC per night. Therefore, we will over-estimate the degree of assortativity and obtain an upper bound on how much preferences increase in own type.

Figure B.1 shows the conditional distribution destination location attractiveness given own attractiveness for all first requests of guests who have a completed home themselves. The black dots show bin means. While there is a very mild positive relationship between the attractiveness of guests' own location and their destination, it is striking how flat this relation is. A guest with a one standard deviation more attractive location on average demand a 0.1 SD more attractive location. Furthermore, guests with any house demand every type of destination in a quite similar proportion. Even with respect to home size (number of beds) and amenities (comfort score) assortativity is quite weak. Guests who have bed-space for 2 adults themselves on average request homes that can host 4 adults while guests with homes for 10-12 adults request homes for 6 people (see Figure B.2).)

An alternative index of types is homes' price, which incorporates locations, size and quality. Figure B.3 plots the conditional distribution of destination homes' prices given the price of guests' own home in bins of 25 IC . This time we consider all completed trips rather than the first request, which likely overestimates the degree of assortativity more. Nevertheless, I find that travel patterns are all over the place. While there is some sorting, it is far from perfect. If users perfectly sorted themselves perfectly on IC prices, users whose own home costs 275 ± 12.5 IC should only visit homes with a price in that interval. Instead, the average destination price is only around 180 with much variation

[44] For the small group of users with more than one home I use the most attractive location.

Figure B.1: Location-assortativity of first requests

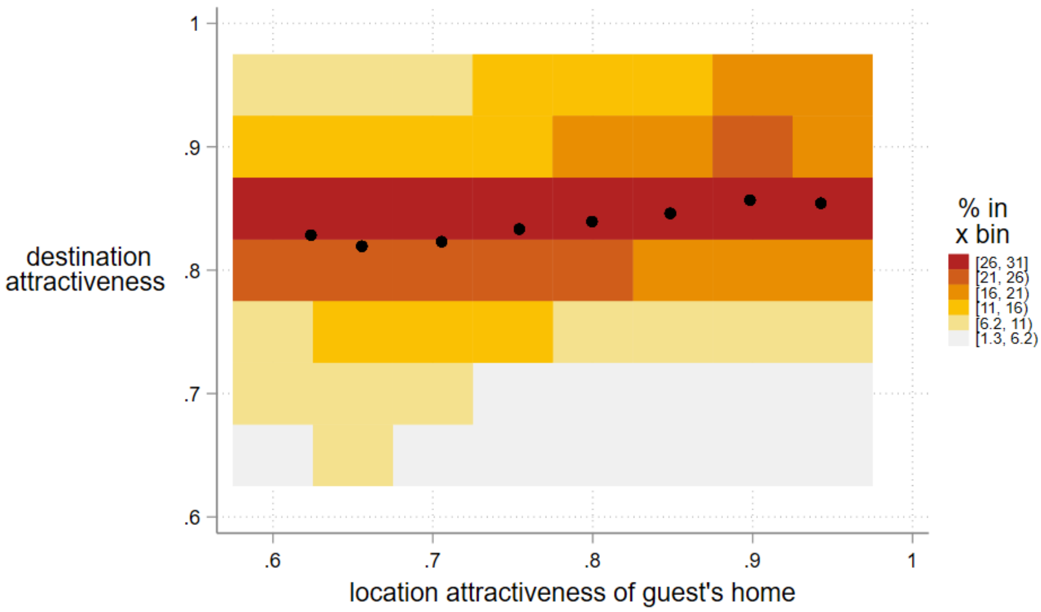

Sample: all 312395 first requests of guests with a completed home. Location attractiveness is the probability of not finding a holiday conditional on requesting the location. If the guest has multiple homes the max is used. For each value of x, shares add up to 100% vertically. Dots show bin mean.

around that. At least 10% of trips are to homes that cost less than 100 IC. Similarly, users whose own home costs around 50 visit on average places that cost about 130 - not that much lower than 180.

Overall, these results suggest that preferences cannot increase much in users' own type and could well be decreasing - consistent with the observed supply effect and the secondary model predictions.

B.4 Side-effects of compressed price

B.4.1 Lower effort

Figure D.26 shows that a given guest gives slightly worse ratings to hosts in over-demanded places. I include all 240,726 trips where the guest left a start rating and control for a range of other potential determinants of ratings. Furthermore, guests also *receive* slightly worse ratings (figure B.4). It is hard to say if guests actually behave worse or if host have higher expectations. Some indication that it is hosts comes from the fact that hosts cancel slightly more often but guests do not (figure B.5). Thus, the

Figure B.2: Assortativity of first requests on location, home size and amenities

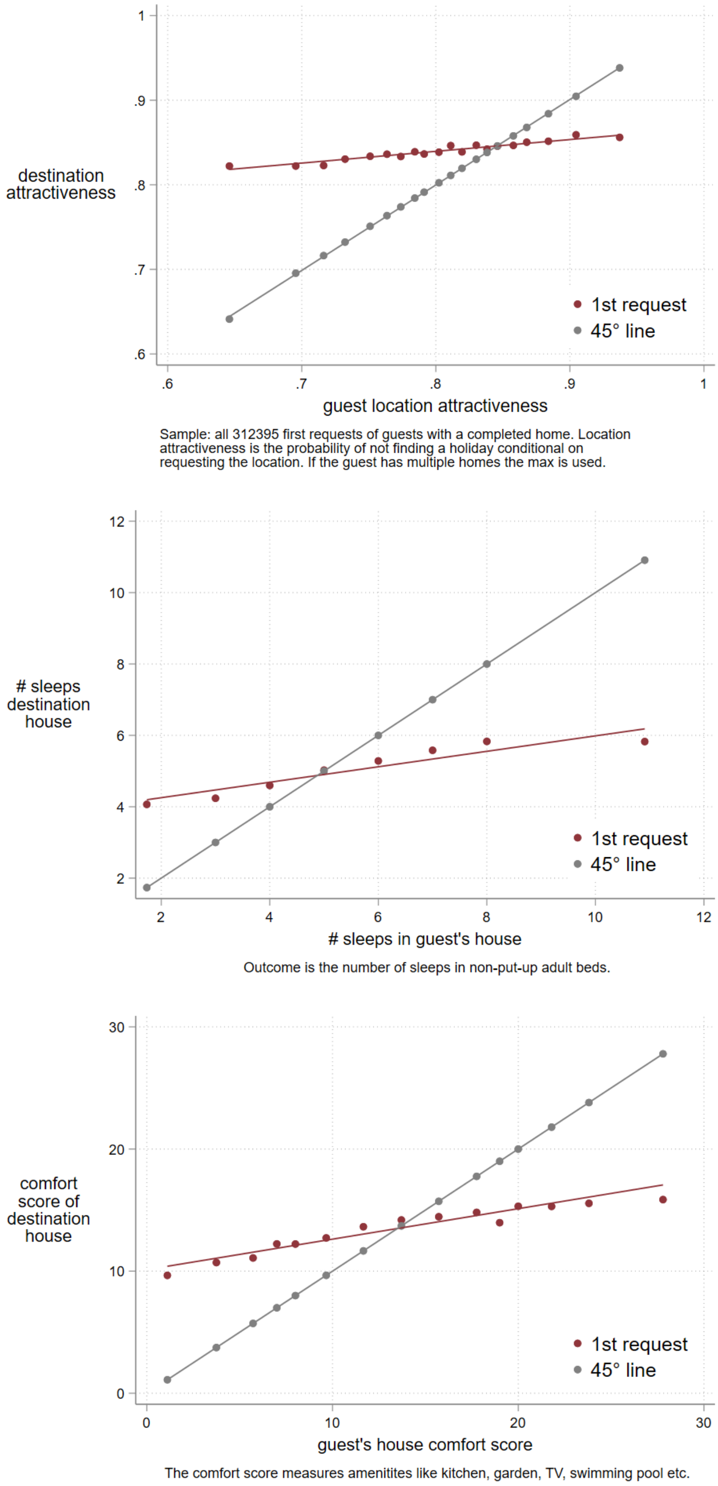

Figure B.3: Price-assortativity of finalized trips

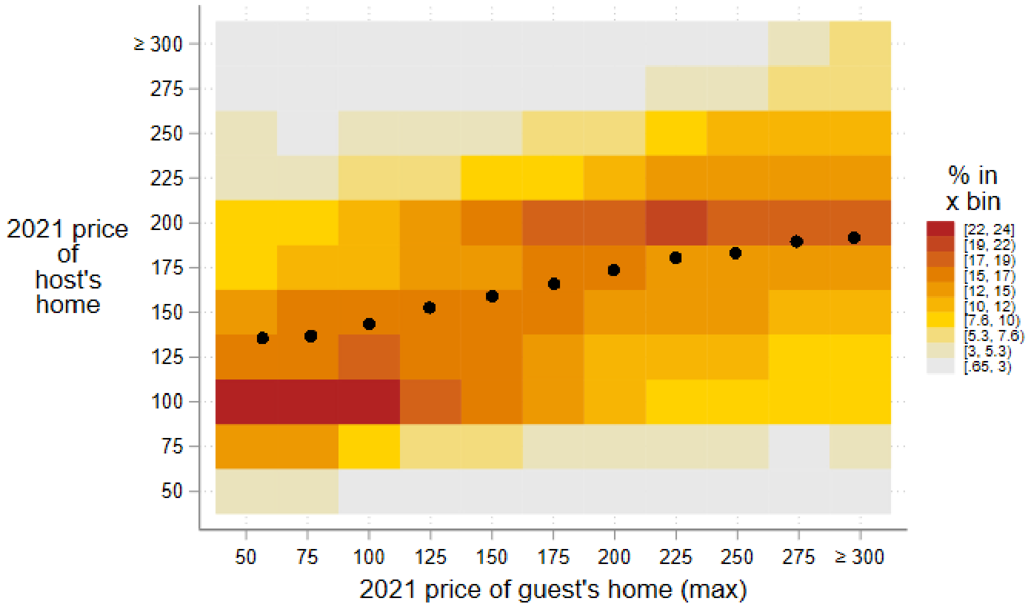

Graph compares actual 2021 price of host and guest home in 335970 completed trips. For each value of x, the shares add up to 100% vertically. Dots show bin mean. If the guest has multiple home the max is used.

negative effect on effort seems to dominate.

Yet, all of these effects are quite small. The estimated effect of moving from P10 to P90 in terms of over-demand is merely reducing the average rating from 4.92 out of 5 stars to 4.91 stars. Even in the most over-demanded places homes and guests are still rated above 4.9/5 stars on average and cancellations increase by less than one percentage point from 6.3% to 7.2%. It is not completely obvious how guests' experience maps into ratings but it seems implausible that a 0.01 point difference in average ratings reflects substantially different experiences.

B.4.2 Misallocation

In a cross-section the utilization of beds is not systematically related to how demanded a location is. Figure B.6 shows that the ratio of the ratio of guests (indicated in the request) to adult beds in the requested house does not increase with the location's success rate when flexibly controlling for the number of beds and rooms in the house. If anything, requested bed utilization seems slightly higher in more popular locations. The guest-to-bed ratio in completed trips is not related to success rates either.

Moreover, in an event-study in a narrow window around the reform I do not find that big apartments get requested by smaller parties when their price drops due to the change in pricing algorithm (Figure D.27). The event-study uses the same specification is explained in section C.1.3 but uses the number of guests as outcome. This mitigates concerns about misallocation with respect to group sizes at the current prices.

C Details on participation and Airbnb comparison

C.1 Evidence of excess demand at current prices

C.1.1 Computing a location-specific success rate

In a second approach to find out if the current prices are close to clearing the market I look for signs that demand exceeds supply in some locations. Locations are likely the key dimension of homes' type as most users presumably envisage a holiday in a certain country or city and users first need to enter a location to launch a search before possibly applying additional filters.

Therefore, I compute a location-specific success rate that captures how likely a candidate

Figure B.4: Rating guests receive vs. over-demand

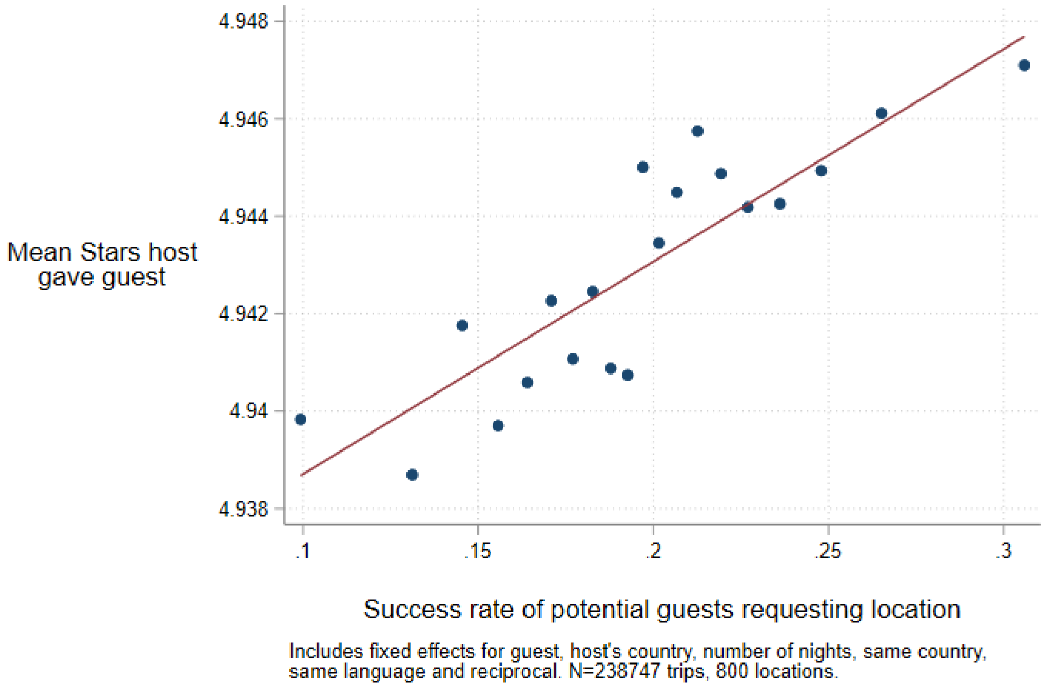

Figure B.5: Cancellations vs. over-demand

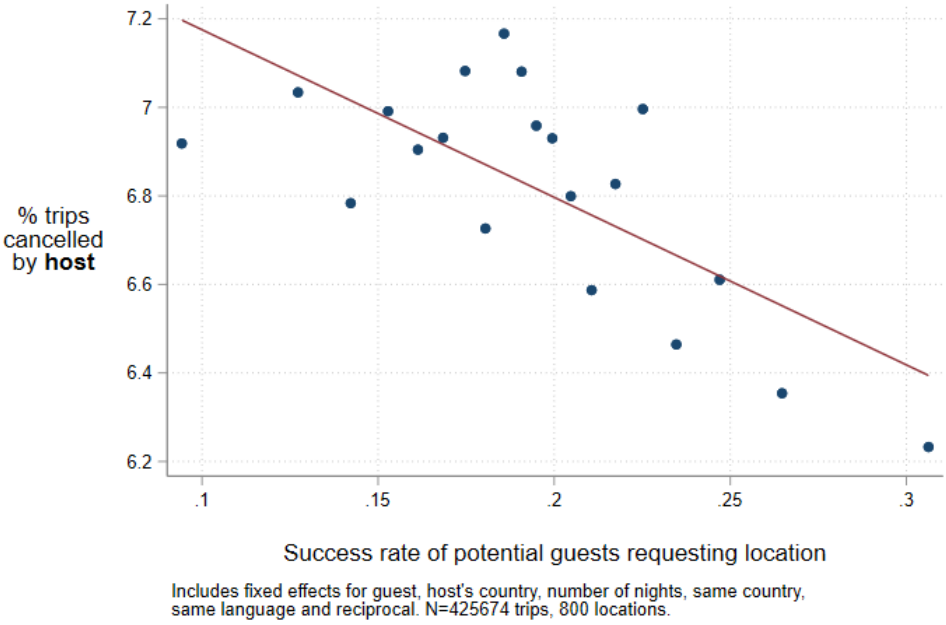

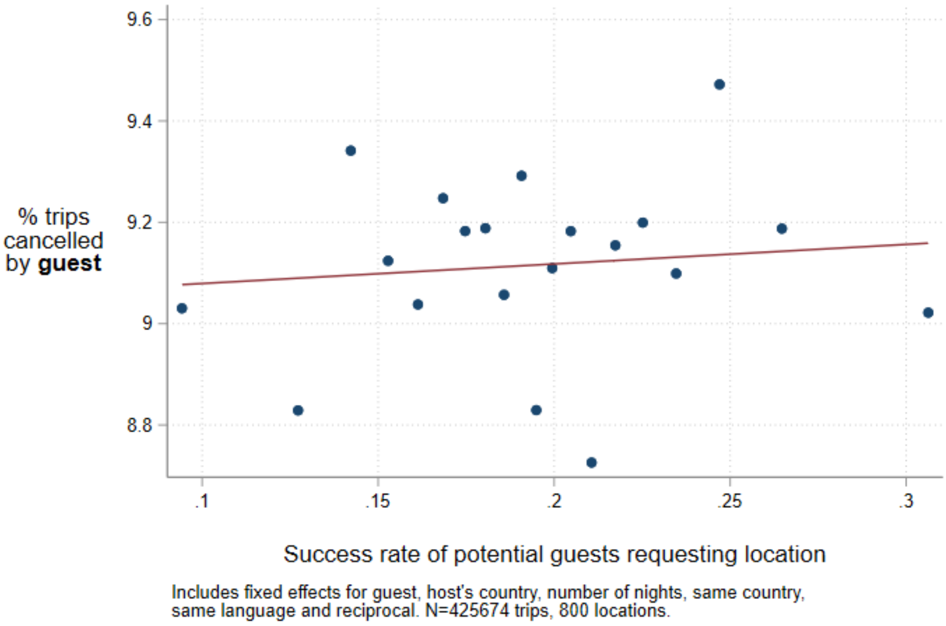

Figure B.6: Bed utilization vs. over-demand

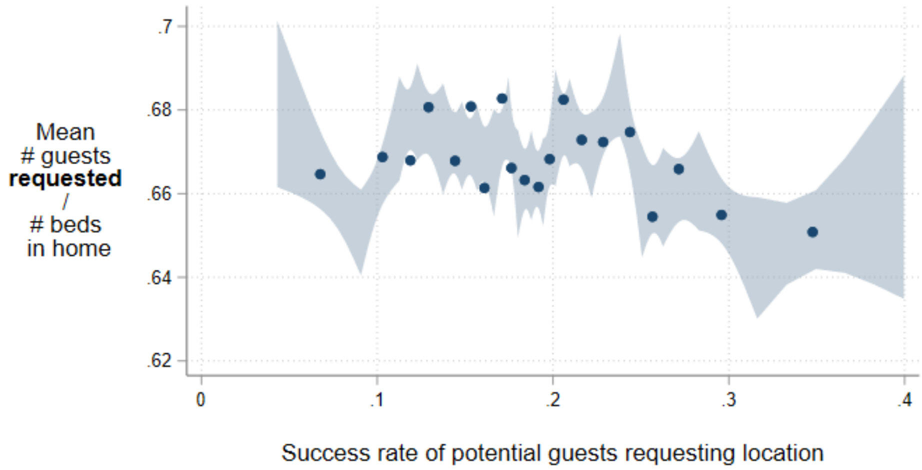

The outcome is the average number of guests (group-size) that users indicate in requests to a house divided by the number of non-put-up adult beds in the home. Includes fixed effects for # non-put-up adult beds, # total beds, # double beds and # rooms using the procedure in Cattaneo et al. (2022). 95%-CI are constructed using a degree 1 polynomial which was automatically selected as optimal by the procedure in Cattaneo et al. (2022). Sample: all 84617 homes that hosted at least one trip with known number of guests.

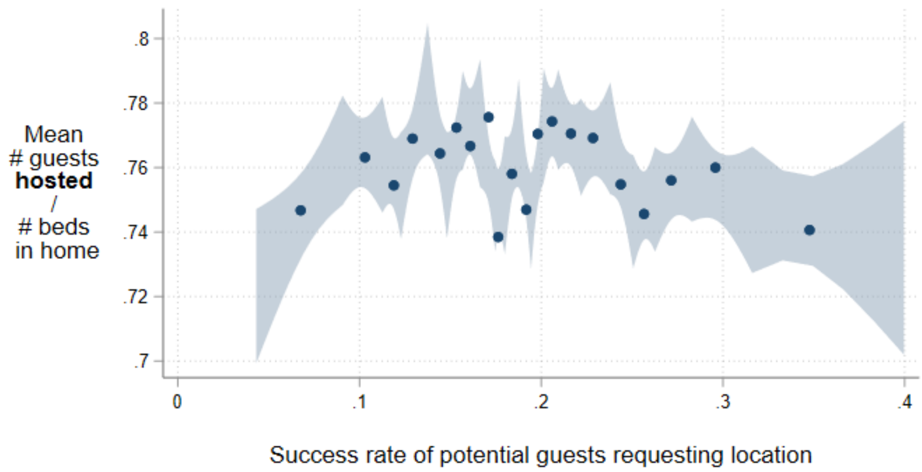

The outcome is the average number of guests (group-size) hosted divided by the number of beds non-put-up adult beds in the home. Includes fixed effects for # non-put-up adult beds, # total beds, # double beds and # rooms using the procedure in Cattaneo et al. (2022). 95%-CI are constructed using a degree 1 polynomial which was automatically selected as optimal by the procedure in Cattaneo et al. (2022). Sample: all 84617 homes that hosted at least one trip with known number of guests.

Figure C.7: Locations available on Inside Airbnb vs. studied platform

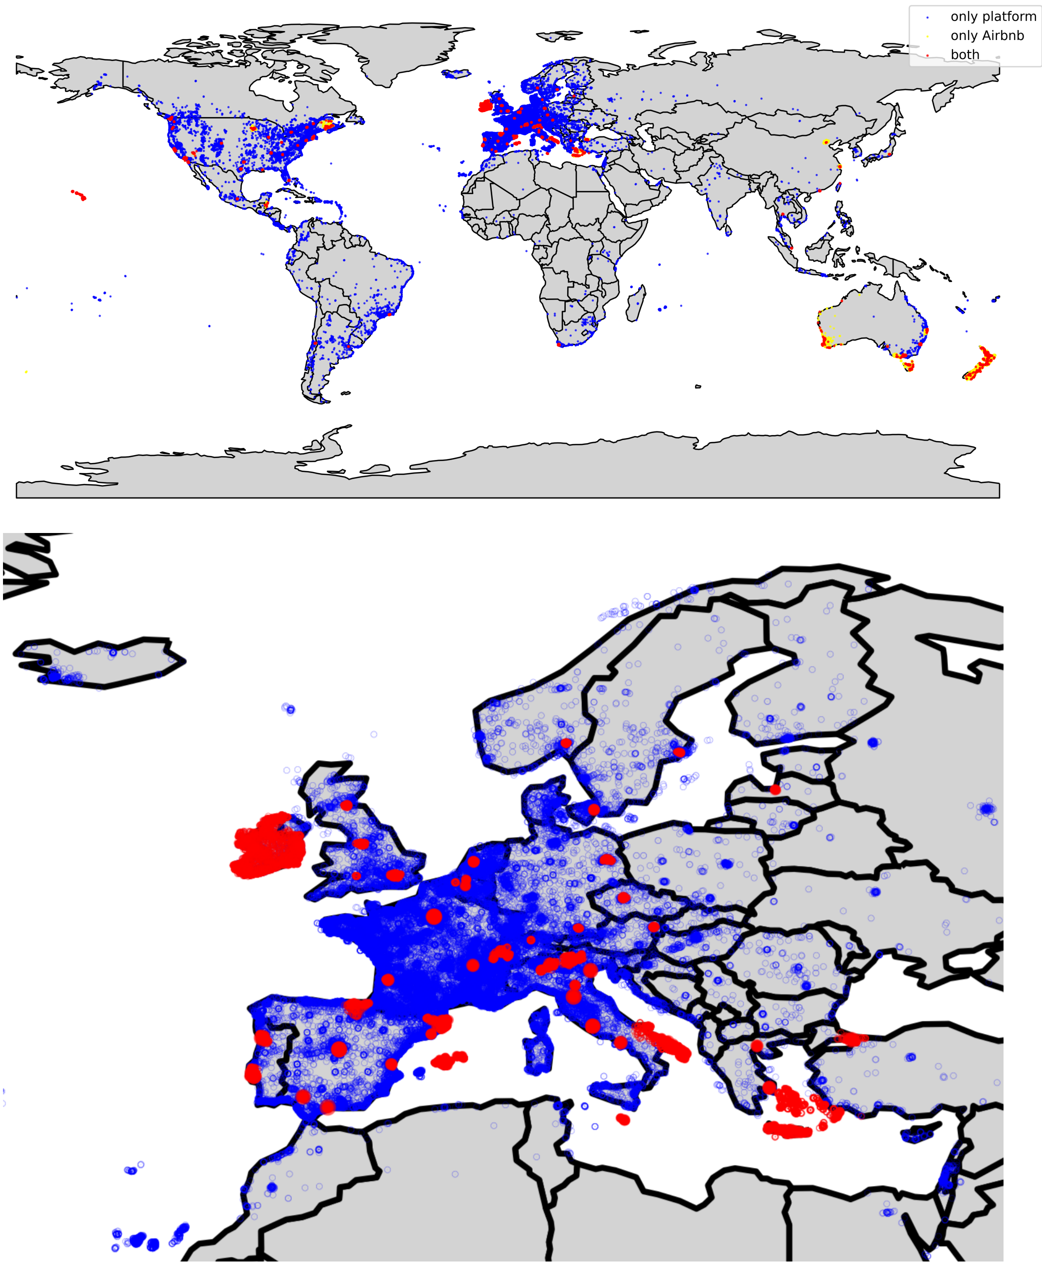

Each dot represents a 1km grid cell that contains a least one house. Red dots are cells with homes on both websites, yellow dots cells with only Airbnb homes, blue dots cells with platform homes but no Airbnb data.

Figure C.8: Out-of-sample fit of hedonic Airbnb model

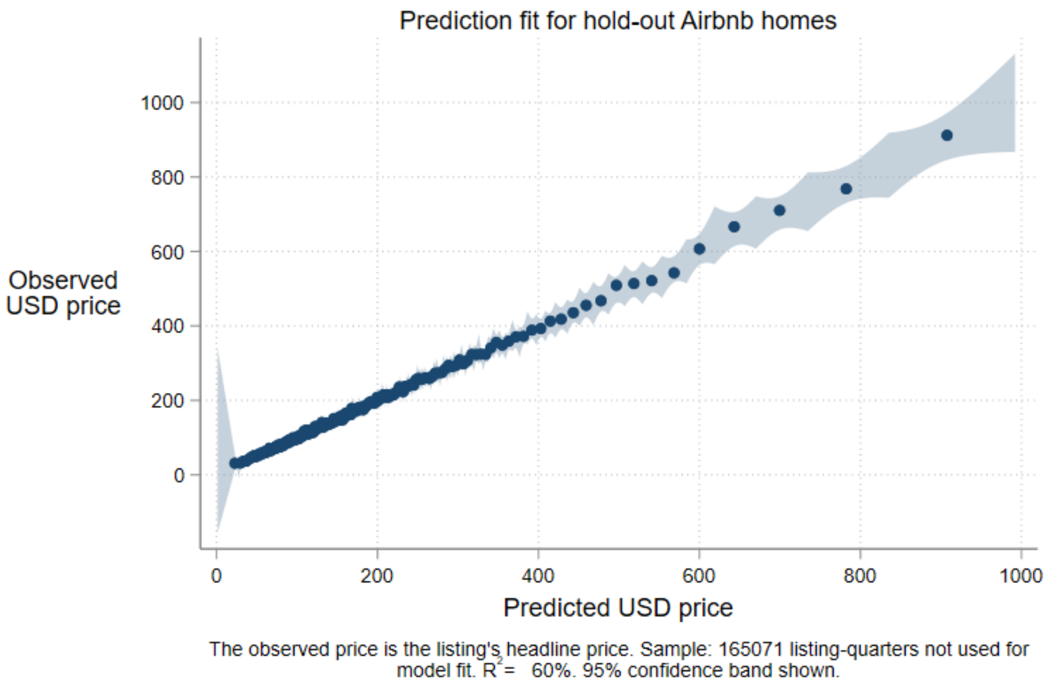

	Locations with Airbnb data				All locations
	Personal residences		All listings		
Listing characteristic	Airbnb	Exchange	Airbnb	Exchange	Exchange
Bathroom Count	1.4	1.7	1.4	1.7	1.7
Bedroom Count	1.9	2.5	1.9	2.4	2.7
Capacity	4.3	5.6	4.5	5.6	6.4
Garden	.072	.44	.078	.43	.6
Balcony Terrace	.28	.64	.38	.64	.7
Elevator	.25	.34	.26	.35	.2
Swimming Pool	.13	.13	.16	.12	.21
Aircondition	.43	.4	.55	.4	.36
Microwave	.36	.79	.55	.79	.8
Dishwasher	.28	.8	.35	.8	.79
Washing Machine	.76	.94	.72	.93	.91
Dryer	.33	.56	.36	.56	.52
Wifi	.9	.95	.93	.95	.91
Bathtub	.18	.69	.2	.68	.65
Number Of Reviews	4.1	1.9	33	2.5	2.3
Reviews Per Year	1.4	.74	10	1.1	
Personal Residence Proxy	1	1	.29	.93	
N	125040	61316	434948	65665	110196

The table shows the mean characteristics of listings on Airbnb and exchange platform. Columns 1-2 restrict to listings that meet the personal residence definition. Column 5 shows all completed and verified listings on exchange platform by 2022, including locations not covered by the InsideAirbnb data. The number of observations is larger than in figure 12 because this table also includes completed and verified but unrated listings.

Table C1: Home characteristics on Airbnb and Exchange platform guest is to find a home in a given location for a given period conditional on sending at least one request to that location for that period.[45] I define locations as countries unless a sub-national geographic unit has more than 500 homes, in which case this becomes a location of its own. For instance, Spain might be divided into "Barcelona", "Madrid", "[the rest of] Catalonia", "Andalucía", "Balearic islands" and "[the rest of] Spain". This flexible definition has the advantage that each location has a sufficient number of homes to estimate the success rate without too much noise while allowing a more fine-grained analysis for popular destinations. This definition yields 801 distinct locations with at least 500 homes. Most results are qualitatively similar when use 1000, 5000 and 10000 homes as threshold. For each user, I define a holiday search as a group of requests with travel dates within 30 days of some other requests. For each holiday search I consider all requests sent within 24 hours of the first request. For each location I then compute the share of holiday searches that were approved by at least one host.[46]

C.1.2 Results

In a competitive equilibrium we would expect a success rate of 100% everywhere. Market-clearing requires that the number of users willing to pay the posted price for a given type of home equals the supply of such homes. However, even on Airbnb there are substantial search frictions and more than 40% of requests get rejected (Fradkin, 2017). Here, the largest approval rate in any location is around 45% and most locations are substantially below. This is not too surprising as we know that users are not professionals who frequently update their availability and hosts presumably require stronger trust as they offer their primary residences. More importantly, Figure C.10 shows that success rates vary wildly across locations. The location with the lowest success rates turn out

[45] If we observed search data we could alternatively define demand as the number of users searching for a given location at given dates and compare this to the number of homes marked as available for those dates. However, this definition has its own weaknesses. First, many hosts do not use the calendar actively and are willing to host even if they did not mark their home as available. About half of realized trips were not marked as available in the calendar. Thus we would underestimate supply. Second, some users may browse a range of locations without a serious intention to visit. Thus we might overestimate demand.

[46] At least one is used because formal approval typically only occurs after guest and host have verbally agreed on the trip and I want to ignore the unusual cases where more than one host formally approves. As users often send requests to more than one location as defined above, approvals are weighted by the share of requests sent to that location. For example: A user sends 6 requests with travel dates in July 2019, 4 to Montpellier and 2 to Marseille. If any host in Montpellier approved this is weighted as 6/4. Again the reason is that typically only one request will be formally approved. Suppose all 6 hosts verbally signal willingness to host and the guest randomly picks one of them to continue the booking process, there is only a 4/6 chance that we observe a formal approval in Montpellier although willingness to host is 1.

to be New York. In Queens, the success rate is only 2%, in the rest of NYC 4% and in the rest of New York State 10%. Other very demanded locations include London with a rate of 6% (again, specific districts like St. James, Kings and Kent are below), Tel Aviv (6%), the center of Paris (1st arrondissement, 7%) and some expensive seaside tourist destination like Miami (7%), Biarritz and Marbella (8%). By contrast, less touristy districts of Paris have much higher success rates around 27% (e.g. Paris 19 and 20). The highest success rates are observed in Hungary and some parts of France that have many platform users but are not extremely touristy. This makes sense as the platform is very popular in France so that hosts may struggle to find guests. The current results presumably mask some further heterogeneity by season. Some of the beach locations are likely much more demanded in summer and some mountain locations more in winter.

The massive differences in success rates across locations are also in line with the results of a machine learning exercise presented in figure D.21. I explore which features best predict whether an individual request will be accepted by the host. I consider a large set of variables describing characteristics of the trip (e.g. dates, number of guests, time asked in advanced), the home (location, # beds), the guest and the host (e.g. ratings, ID verification, time on the platform) as well as the pair (e.g. age difference, same main language). In a flexible XGBoost model the location of the requested house turns out to be the most important predictor of approval (measured as total gain across splits), followed by the host's IC wealth at the time of the request and the month of the trip. An obvious interpretation is that hosts in over-demanded places are less likely to accept a given request, because they are only able or willing to host a certain number of nights per year.

If the platform wanted to reduce demand for places like NYC to average levels it would need to raise recommended (and thus actual) prices in over-demanded locations substantially. An event-study in section C.1.3 suggests that prices in NYC might have to increase 2-4 times. Overall, the results in this section suggest that current prices are strongly compressed. Both approaches suggest that the price of the most demanded homes would at least have to double to clear the market.

C.1.3 Gauging the price-elasticity of demand

I estimate with an event-study how demand responds to the sudden price changes induced by the reform. This gives an idea how much prices in over-demanded locations would need to rise to substantially reduce excess demand. It is also reassuring to check

Figure C.9: keyword-based rejections vs. price disconnect

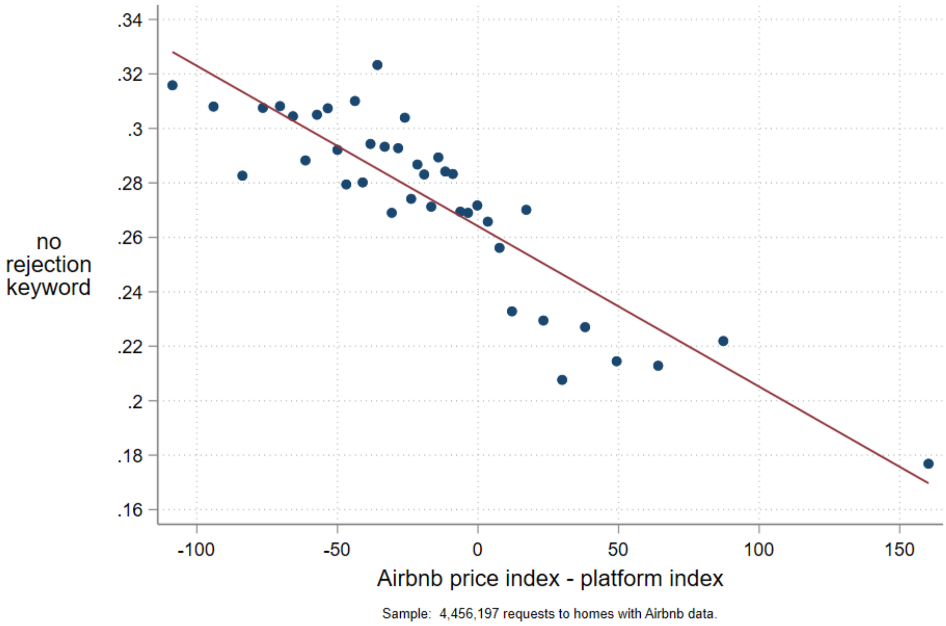

that users perceive of IC as something valuable that should not be squandered and rationally respond to that.

In addition to changing the algorithm that determines homes' default prices the platform limited users' ability to adjust prices by imposing the cap of 30IC above the recommended nightly price. Furthermore, the prices of all homes were reset to the new recommended price on the day of the reform, which creates large short-term variation in actual prices that I will exploit in the event-study. For the difference-in-difference design we have to use recommended (default) prices because the actual prices are endogenously chosen by hosts. While the recommended price did not increase substantially for any homes, some homes had chosen a price substantially below the old recommended price and experienced large increases in actual price on the day of the reform. This allows us to create a price increase treatment group. Since users the owners of those homes could subsequently reduce their prices again, we cannot use this treatment for the difference-in-difference design in section 20.

The reform happened "over night" and was not expected by users. This is supported by the fairly flat pre-trends of both treatment groups. It seems unlikely that hosts' behavior changed much in the week of the reform. We have seen before that members use the

Figure C.10: Success rate across locations

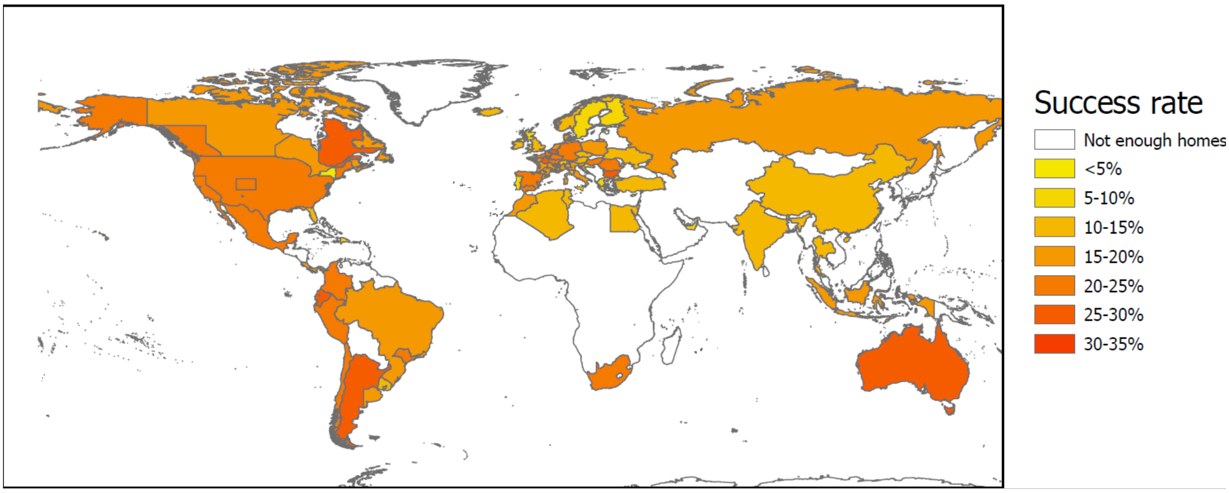

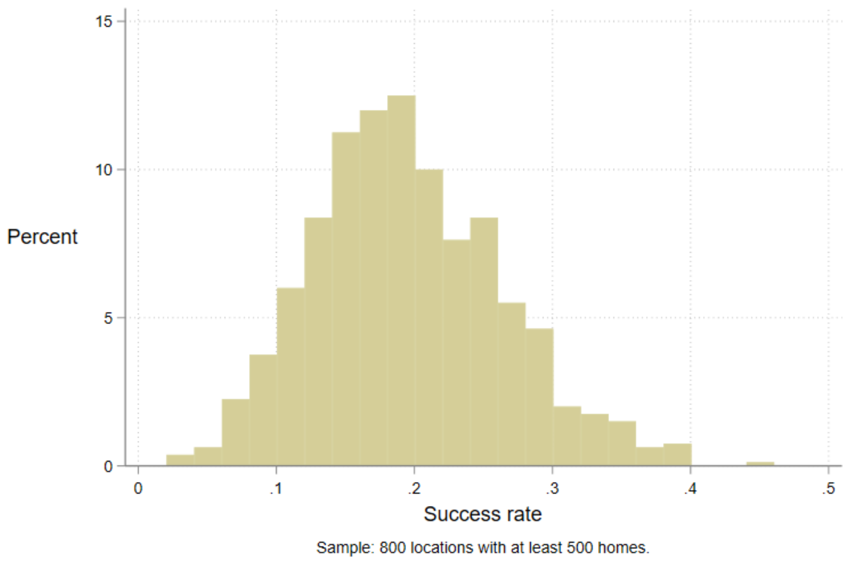

The success rate is the share of "holiday searches" for that location that were approved by at least one host. Locations are defined as countries or sub-national regions with at least 500 homes in the histogram and at least 10,000 homes in the map for readability.

platform very casually and do not update their availability very often. Presumably, most hosts noticed weeks or months later that the price of their home changed when checking out their profile. This is supported by lagged supply response in section 20. By contrast, Users searching for a home a day before the reform would observe a different price than users looking at the same home a day later. Thus, demand should respond immediately.

I somewhat arbitrarily define two treatment groups with quite extreme changes in the (actual) price, one with a large price drop (at least 50%) and one with a large surge (at least 50%), and a control group whose prices changed less than 10%. I discard homes with a price change in $\pm(10\%, 50\%)$. [47] The main specification is

$$\begin{aligned} Y_{it} = {} & \alpha_1 P_i^{up} + \alpha_2 P_i^{down} + \mu_{l(i)m(t)} \\ & + \sum_{s \neq -1} \delta_s \mathbb{1}(s = t) + \sum_{s \neq -1} \beta_s \mathbb{1}(s = t) P_i^{up} + \sum_{s \neq -1} \gamma_s \mathbb{1}(s = t) P_i^{down} + \varepsilon_{ijst} \end{aligned} \tag{38}$$

where Y_{it} is the number of requests home i receives in week t, and the treatment indicators $P_i^{up} \equiv \mathbb{1}(p_{i0}/p_{i-1} \leq 0.5)$ and $P_i^{down} \equiv \mathbb{1}(p_{i0}/p_{i-1} \geq 1.5)$. I then estimate changes in Y_{it} for both treatment groups during the 10 weeks before and after the reform using the specification below. It includes location-month fixed effects $\mu_{l(i)m(t)}$ to reduce noise stemming from the different timing of school holidays in different locations.

Figure C.11 shows the resulting estimates for β_s and γ_s. The fairly flat pre-trends of both treatment groups support the idea that users did not anticipate the reform. Exactly at the week during which the reform occurred demand for homes that suddenly became cheaper jumps up by ca. 60% while demand for homes that became more expensive sharply drops by about 30%. Thus, demand responds to prices as expected. Since the treatments are price changes of at least 50%, the estimates above suggest that the average price-elasticity of demand (for homes with a large price change) cannot be much above one.

We can use these estimates for a quick back-of-the-envelope calculation how much prices in the most over-demanded locations would need to rise to reduce demand enough to

[47] In the event-study, the treatment is defined based on the difference between the last *actual* price before the reform and the new price after the reform because the actual price can be treated as exogenous over the short 10-week time period since users change their prices infrequently and did not anticipate the reform. The actual (chosen) price is less plausibly exogenous over the 7-year period studied now. Therefore, the DiD treatment is defined using *recommended* prices, which are deterministically computed based on the characteristics of the house and changed permanently in the reform. As the changes in the recommended price are less extreme than the (short-term) changes in actual price I use a smaller threshold to define the treatment group.

Figure C.11: Immediate effect of reform on demand

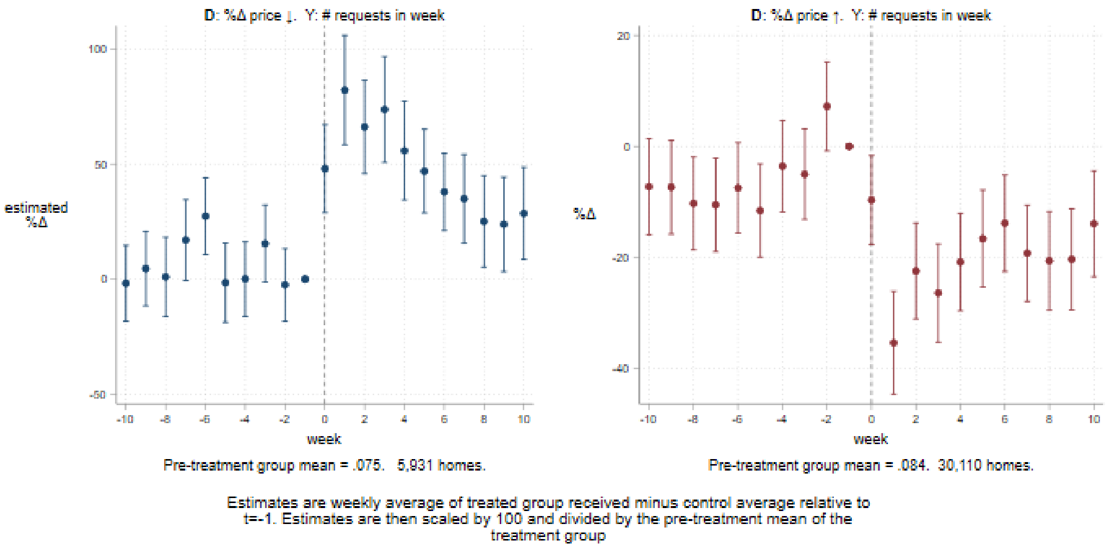

reach average success rates holding supply constant. The average success rate across locations is about 20%, the top 1% (10%) of locations have a success rate of about 5% (10%). Increasing the success rate from 10 to 20 percent requires halving demand. Assuming an elasticity of one prices of the 10% most popular locations would need to double and prices in the top 1% - such as NYC and London - would have to quadruple. This simple calculation assumes homogeneous elasticities. However, even if we believed the price-elasticity to be higher, say two, in popular locations (which is not obvious) or at higher price levels, this would still imply very substantial increases.

D Additional figures

D.1 Price compression in various internal currency systems

Figure D.12: Degree of price compression in example internal currency systems around the world

compress prices more

	goods	services	homes	institutional
Market-clearing prices (or unrestricted decentralized prices)	• Bunz	• Simbi	• Trade to Travel • ILikeYourSpace	• Feeding America • CourseMatch • UMBS Bidding
partially compressed prices			• Love Home Swap • Third home • IVHE	• NKR Center Liquidity Contribution Program
1:1 prices	• Geev • PaperBackSwap	• Baby-sitting co-op • Timebanks	• People Like Us • 48 dots	• Unpaired kidney exchange

Platforms in the "goods" column enable (local) residents to trade (used) goods like cloths, books or appliances. Services often include informal help like practicing to speak in a foreign language or help with computer problems. Homes refer to holiday homes. The institutional column pools various systems that are not peer-to-peer and often managed by public institutions. Bunz and Simbi use a decentralized pricing approach where users can choose an arbitrary price. Thus, the platform does not limit price dispersion. Systems colored gray are not yet or only partly implemented.

Figure D.13: Actual prices vs. default prices

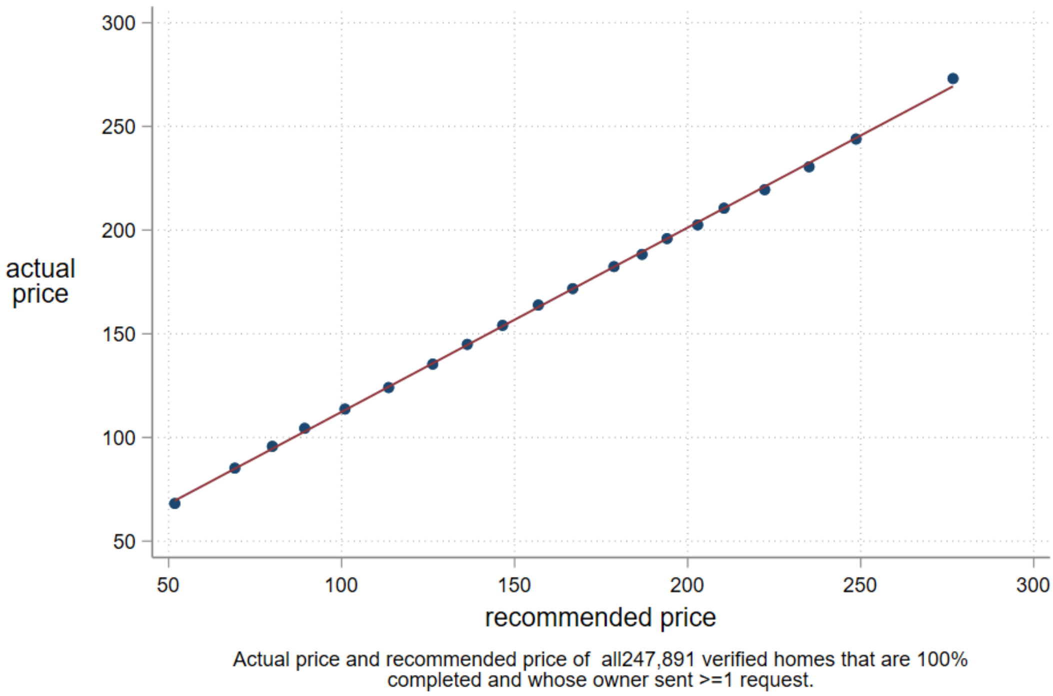

Figure D.14: Quarter-to-quarter wealth dynamics

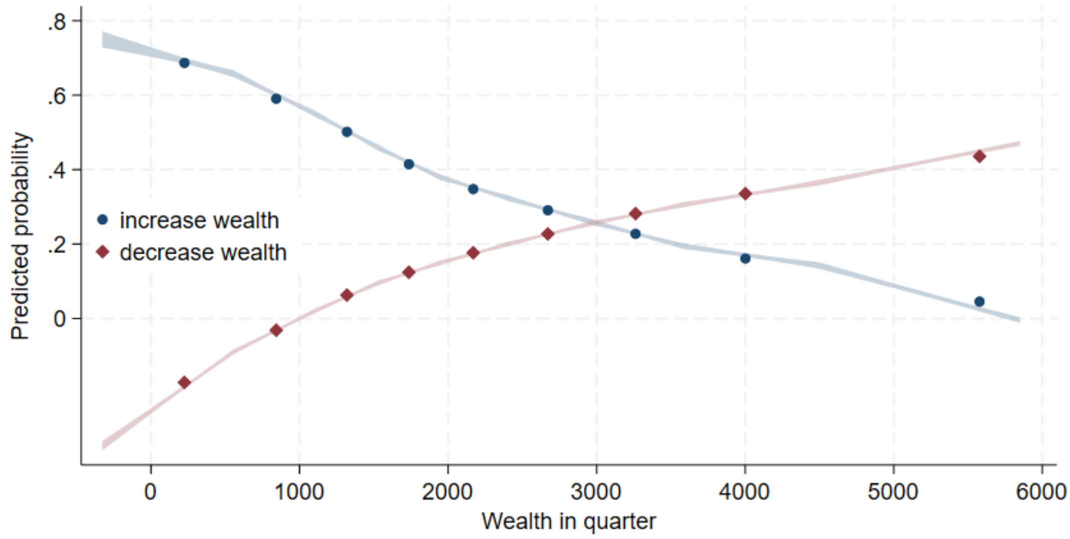

Outcomes are indicators for increasing/decreasing wealth until next quarter. Model includes indicators for deciles of wealth and fixed effects for home, quarters since sign-up and location-by-quarter. Sample: verified homes that have hosted 2+ times in past. N=958103 home-quarters, 79,950 homes, 76,150 users. SE clustered at user-level. 95% CI shown.

Figure D.15: Robustness of referral RDD

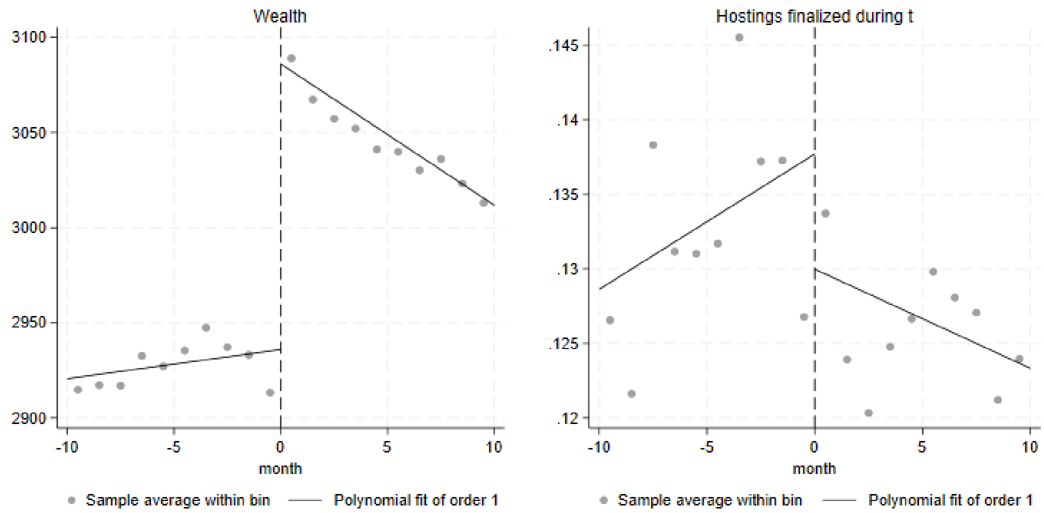

6927 events, 5784 users.
Demeaned outcome by substracting mean of calendar month (not event month).

The figure shows users' wealth and supply around the time when a previously referred "godson" pays their membership fee, triggering a bonus for the referring user. Supply are 1-way hostings finalized in a given month. The specification is the same as in figure 7 but includes event and the directly adjacent period.

Figure D.16: Effect of wealth on nights hosted in quarter

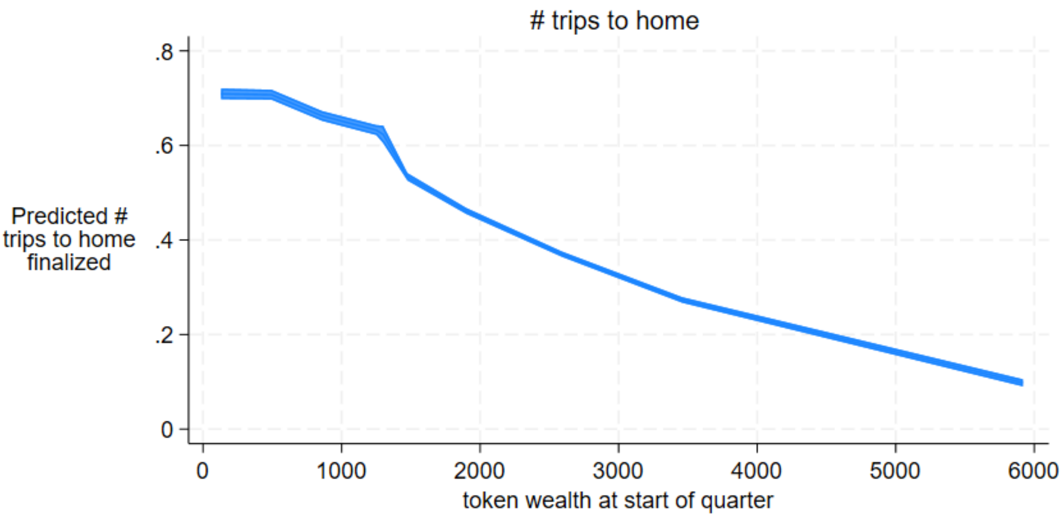

Outcome are nights finalized in that quarter. Model includes indicators for deciles of wealth and fixed effects for home, location-by-quarter, quarters since sign-up, # hosted in last quarter and in last year. Sample: verified homes that have hosted 2+ times in past. N=984593 home-quarters, 81,075 homes, 77,202 users. Y-mean = .39. SE clustered at user-level. 95% CI shown.

Figure D.17: Linear effect of wealth on various outcomes

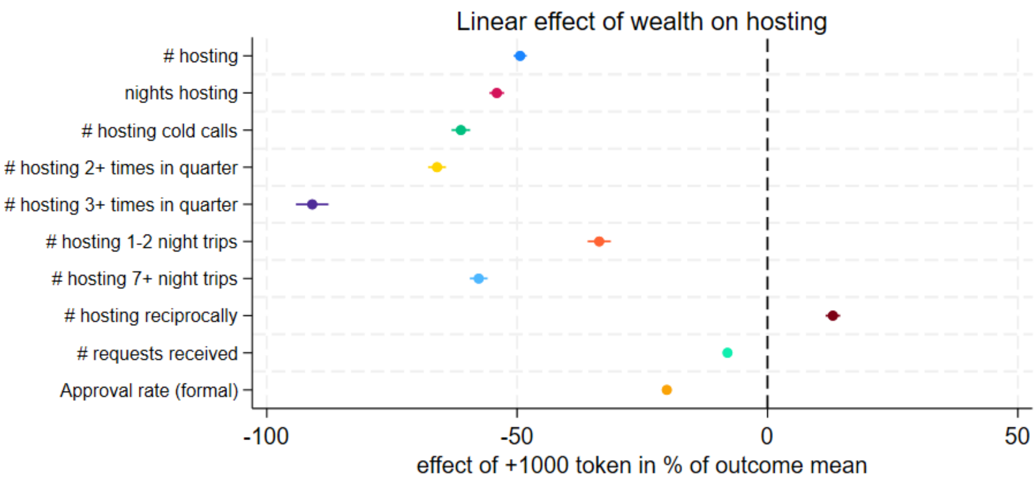

Graph shows the estimated linear effect of wealth at the beginning of a quarter on trips *finalized* and requests *received* in that quarter. Model includes indicators for deciles of wealth and fixed effects for home, location-by-quarter, quarters since sign-up, # hosted in last quarter and in last year. Sample: verified homes that have already hosted 2+ times. N= 984,575 home-quarters, 81,074 homes, 77,201 users. SE clustered at user-level. 95% CI shown.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	# hosted	nights	# hosted	nights	host 3+	cold calls	reciprocal
wealth (1000s) at start of quarter	-27.25***	-29.61***	-27.12***	-28.33***	-45.71***	-33.90***	7.205***
	(0.427)	(0.465)	(0.423)	(0.463)	(1.004)	(0.584)	(0.472)
Benchmark controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
FE lagged hostings	No	No	Yes	Yes	Yes	Yes	Yes
Mean of Y	.19	1.1	.17	1	.013	.067	.04
R^2	0.29	0.24	0.29	0.24	0.18	0.22	0.18
N	2,849,197	2,849,197	2,490,167	2,490,167	2,490,167	2,490,167	2,490,167
Number of hosts	162,666	162,666	151,386	151,386	151,386	151,386	151,386

SE clustered at user-level in parentheses. Effects are normalized as percent of the outcome mean. All models include fixed effects for home, location-by-quarter and quarters since sign-up. Some model additionally include fixed effects # hosted in last quarter and in last year. In column 'cold calls' the outcome is the number of hostings finalized for periods not marked as available in the calendar. Reciprocal are 2-way exchanges which are typically done without currency transfers. All other columns include both 1-way and 2-way hostings. Full sample of verified homes.

Table D2: Effect of wealth on supply when including users with no or little hosting record

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
DiD estimate	0.417***	0.410***	0.415***	0.431***	0.476***	0.376***	0.389***
	(0.0495)	(0.0512)	(0.0552)	(0.0649)	(0.0745)	(0.0633)	(0.0525)
Location-year fixed effects	yes	yes	yes	yes	yes	yes	yes
Price change treatment group	< -20%	< -25%	< -30%	< -35%	< -40%	< -30%	< -30%
Price change control group	[-10%, 10%]	[-10%, 10%]	[-10%, 10%]	[-10%, 10%]	[-10%, 10%]	[-5%, 5%]	[-15%, 15%]
Number of hosts	118,402	110,605	96,619	82,008	72,885	74,724	109,929
Effect in % of Y-mean	21	21	21	22	25	19	19
Mean of Y	2	2	2	2	1.9	2	2

Outcome is all columns is nights hosted. SE clustered at user-level in parentheses. All models include fixed effects for determinants of price change. in (under-demanded) locations with above median success rates.

Table D3: Difference-in-difference estimates for alternative treatment definitions

	calendar nights		calendar nights NR	
	(1)	(2)	(3)	(4)
DiD estimate	1.946**	1.902**	-0.247	-0.398
	(0.609)	(0.640)	(0.237)	(0.258)
Location-year fixed effects	no	yes	no	yes
Number of hosts	97,024	96,619	97,024	96,619
Effect in % of Y-mean	8.1	7.9	-4.2	-6.7
Mean of Y	24	24	5.9	5.9

SE clustered at user-level in parentheses. All models include fixed effects for determinants of price change. Calendar nights are the number of nights that a home is marked as available for in the calendar in the year. NR refers to calendar entries specifically for non-reciprocal exchanges (which are paid with tokens). in (under-demanded) locations with above median success rates.

Table D4: Effect of calendar postings

	(1)	(2)	(3)	(4)
	bottom 50%	top 50%	top 25%	top 10%
DiD estimate	0.309***	0.542***	0.541***	0.570^{+}
	(0.0678)	(0.0914)	(0.143)	(0.292)
Location-year fixed effects	yes	yes	yes	yes
Number of hosts	52,967	45,176	21,689	7,959
Effect in % of Y-mean	17	25	26	26
Mean of Y	1.8	2.2	2.1	2.2

SE clustered at user-level in parentheses. All models include fixed effects for home-type cells that determine price change in reform. Column "bottom 50%" shows estimates for the subsample of homes that are in (less-demanded) locations with above median request success rates. See section C.1 for details. Columns "top 50%" (top 25% and tops 10%) show estimates for the subsample of homes that are in (over-demanded) locations with below median success rates (with success rates below P25, below P10 of all locations).

Table D5: Difference-in-difference estimates in high vs low demand locations

Figure D.18: Linear effect of wealth on various outcomes when including users with no or little hosting record

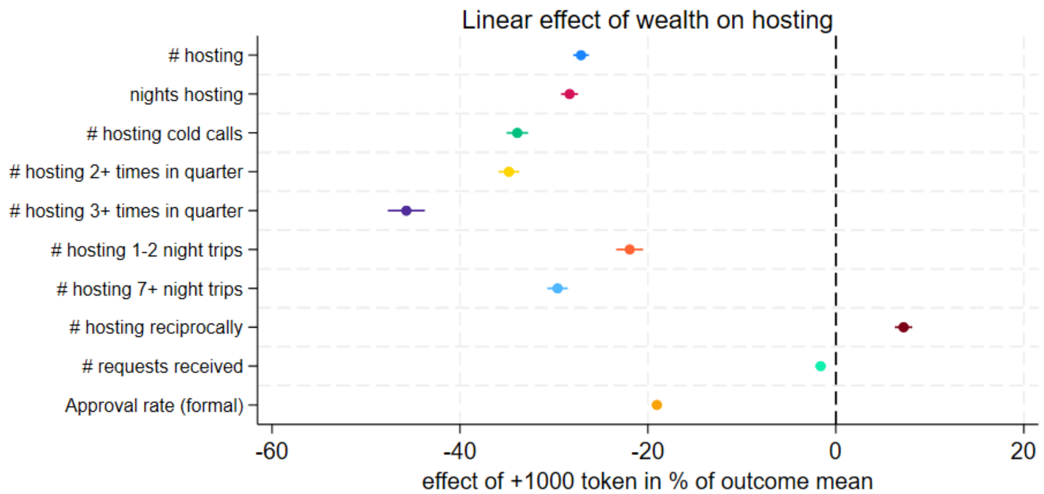

Graph shows the estimated linear effect of wealth at the beginning of a quarter on trips *finalized* and requests *received* in that quarter. Model includes indicators for deciles of wealth and fixed effects for home, location-by-quarter, quarters since sign-up, # hosted in last quarter and in last year. Full sample of verified homes. N= 2,490,167 home-quarters, 163,985 homes, 151,386 users. SE clustered at user-level. 95% CI shown.

	nights hosted		# hosted		ever hosted after year	
	(1)	(2)	(3)	(4)	(5)	(6)
	min $t \leq -1$	min $t \leq -2$	min $t \leq -1$	min $t \leq -2$	min $t \leq -1$	min $t \leq -2$
DiD estimate	0.395***	0.362***	0.0743***	0.0662***	0.0325***	0.0276***
	(0.0671)	(0.0909)	(0.0118)	(0.0162)	(0.00506)	(0.00686)
Location-year fixed effects	no	no	no	no	no	no
Number of hosts	14,377	7,203	14,377	7,203	14,377	7,203
Effect in % of Y-mean	28	26	29	28	10	9.2
Mean of Y	1.4	1.4	.26	.24	.31	.3

Outcome is all columns is nights hosted. SE clustered at user-level in parentheses. All models include fixed effects for determinants of price change. in (under-demanded) locations with above median success rates.

SE clustered at user-level in parentheses. All models include fixed effects for home-type cells that determine price change in reform. Column "min $t \leq -1$" and "min $t \leq -2$" respectively show shows estimates based on the panel of all homes which were all created until $t = -1$ and $t = -2$ respectively.

Table D6: Difference-in-difference estimates in strongly balanced panel of homes active before reform

Figure D.19

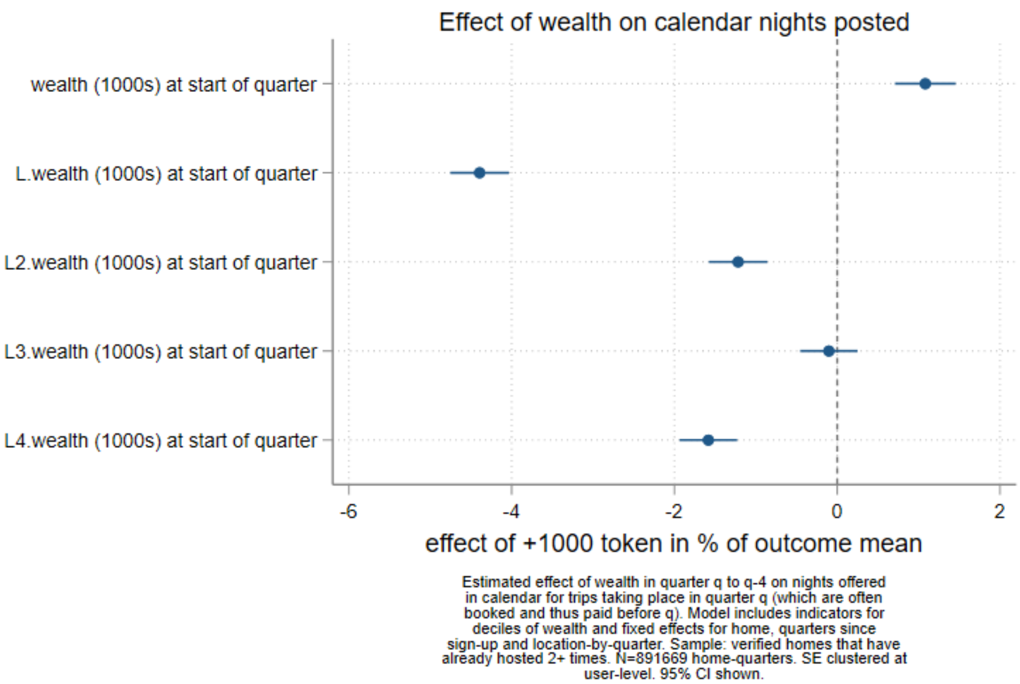

Effect of wealth on calendar nights posted
wealth (1000s) at start of quarter
L.wealth (1000s) at start of quarter
L2.wealth (1000s) at start of quarter
L3.wealth (1000s) at start of quarter
L4.wealth (1000s) at start of quarter
-6
-4
-2
0
2
effect of +1000 token in % of outcome mean
Estimated effect of wealth in quarter q to q-4 on nights offered in calendar for trips taking place in quarter q (which are often booked and thus paid before q). Model includes indicators for deciles of wealth and fixed effects for home, quarters since sign-up and location-by-quarter. Sample: verified homes that have already hosted 2+ times. N=891669 home-quarters. SE clustered at user-level. 95% CI shown.

Figure D.20: Effect of wealth on approving requests

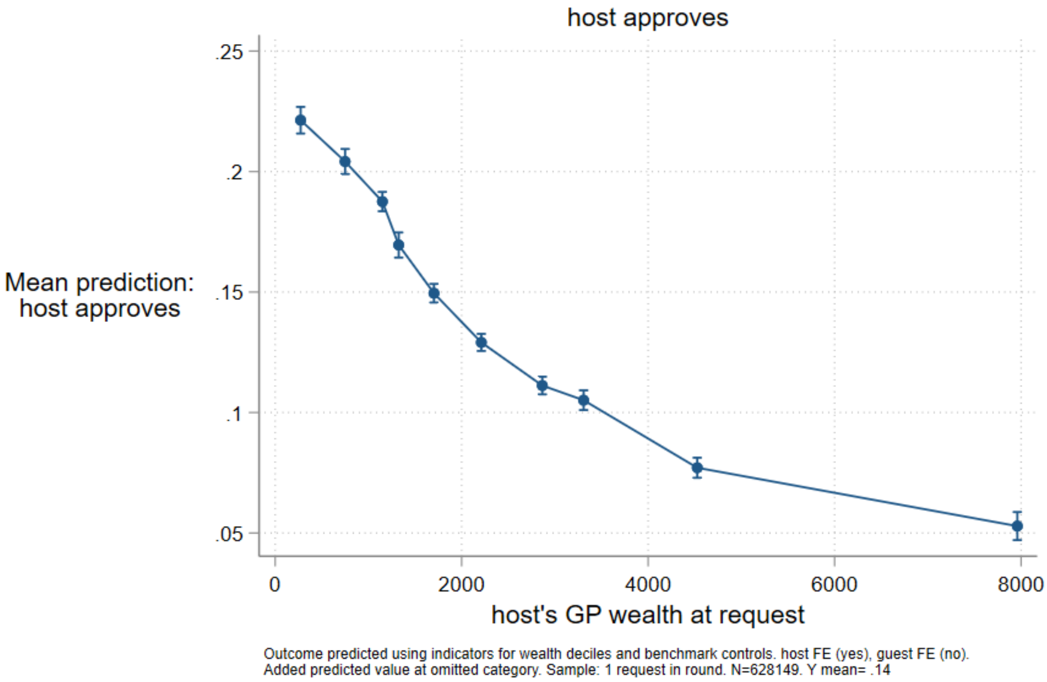

Figure D.21: Most important predictors of approval

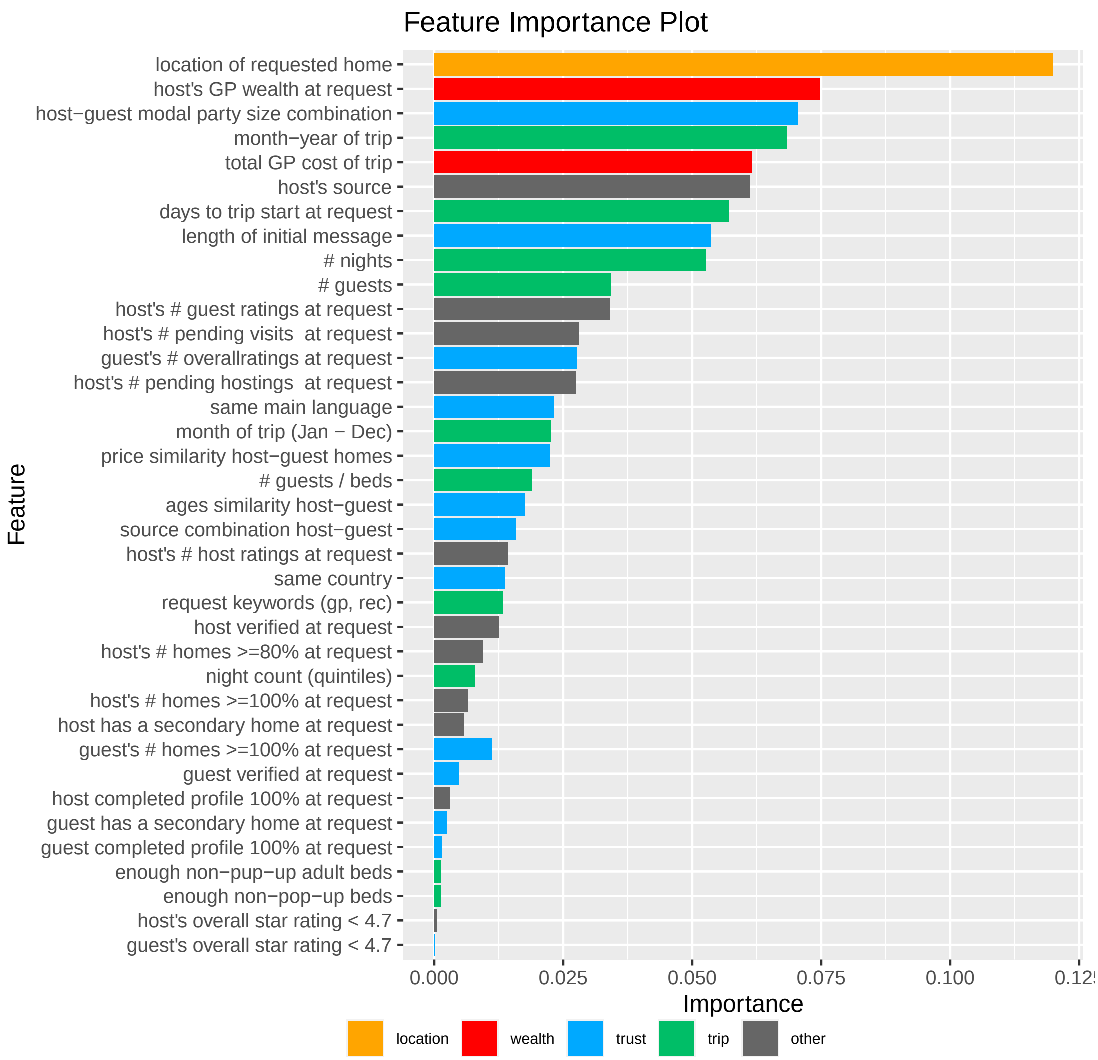

Sample: single requests where the guest can afford the house at the start of the conversation.

Importance is defined as the total gain across splits using this variable in the XGBoost model. The sample are all requests where the guest sent no other request 12h before or after and can afford trip at start of conversation. This ensures that formal approval mostly reflects the host's decision and not how many simultaneous requests the guest sent.

Figure D.22: Calendar postings in difference-in-difference event-study

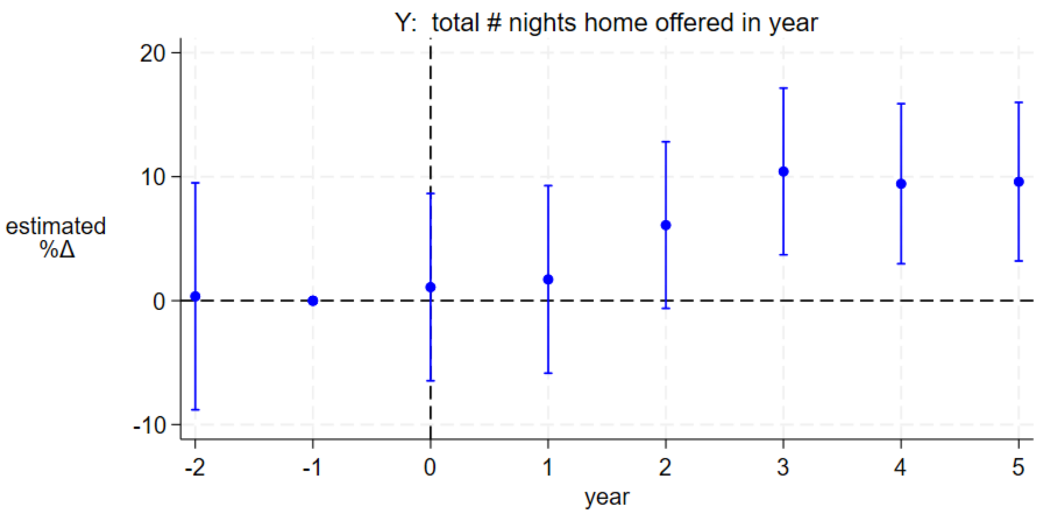

The model includes fixed effects for time and for cells of home-types, which determine treatment status. Estimates divided by the outcome mean and multiplied by 100. Treatment := recommended price of own home dropped by at least 30%. Sample: all homes that are >=85% complete and have 1-15 beds . 41,653 homes. Standard errors are clustered at the user level.

The graph shows difference-in-difference event-study estimates (equation 38) of the effect of the price reduction on treated homes. The outcome is the number of nights the home is marked as available in the calendar. Note that many users do not use the calendar feature, attenuating treatment effects. Prices of homes in the treatment group were reduced by at least 30%. Prices of homes in the control group were reduced by $(-10\%, +10\%)$. The model includes fixed effects for time and for cells of home-types, which determine treatment status. Estimates divided by the outcome mean and multiplied by 100. Standard errors are clustered at the user level.

Figure D.23: Difference-in-difference event-study without location-year fixed-effects

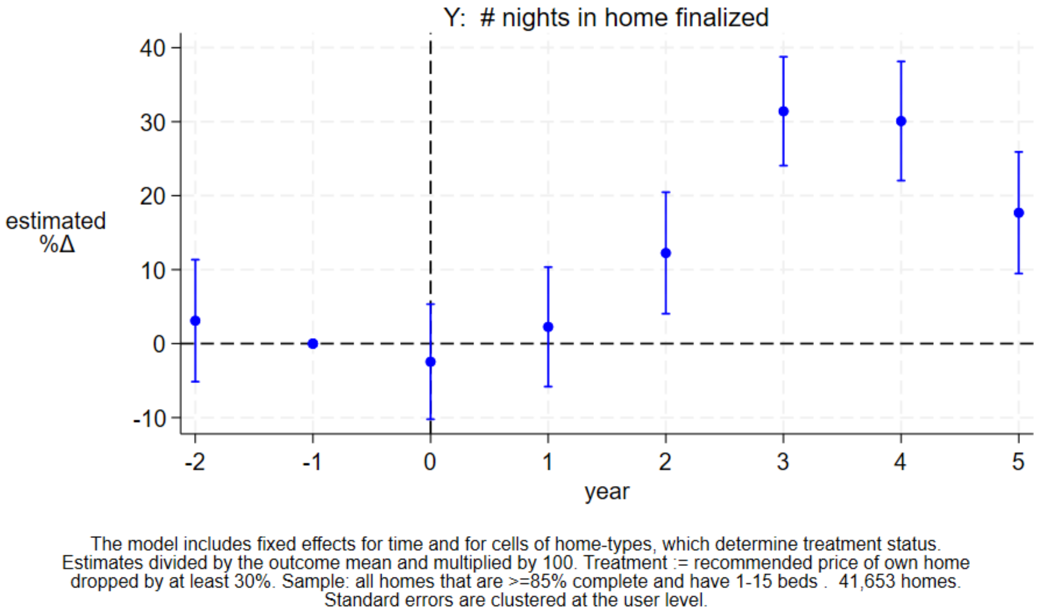

The graph shows difference-in-difference event-study estimates (equation 38) of the effect of the price reduction on treated homes. Prices of homes in the treatment group were reduced by at least 30%. Prices of homes in the control group were reduced by $(-10\%, +10\%)$. The model includes fixed effects for time and for cells of home-types, which determine treatment status. Estimates divided by the outcome mean and multiplied by 100. Standard errors are clustered at the user level.

Figure D.24: Difference-in-difference event-study in strongly balanced panel I

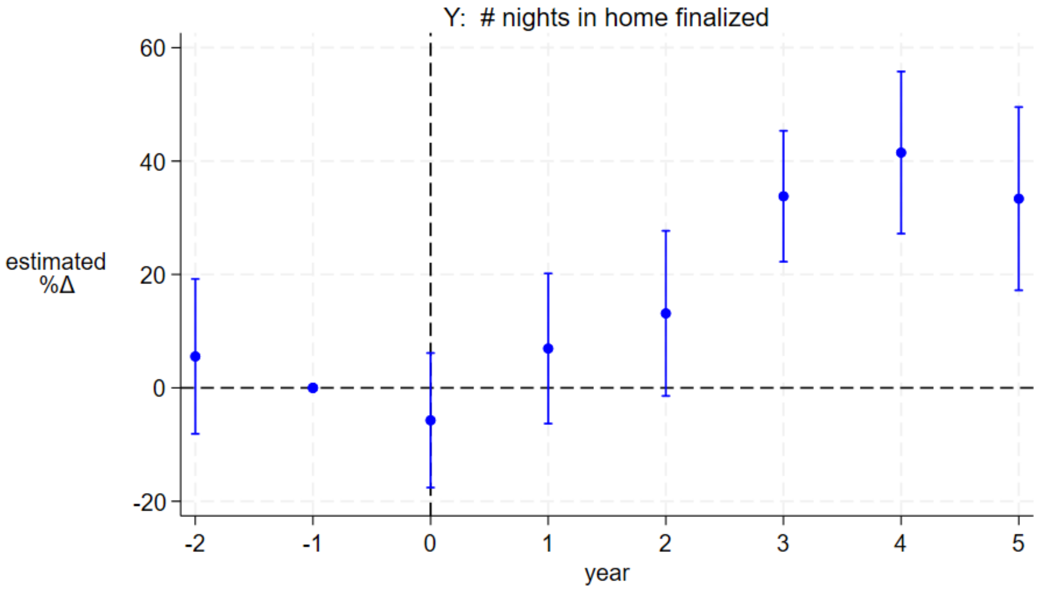

The graph shows difference-in-difference event-study estimates (equation 38) of the effect of the price reduction on treated homes. Prices of homes in the treatment group were reduced by at least 30%. Prices of homes in the control group were reduced by $(-10\%, +10\%)$. The model includes fixed effects for time and for cells of home-types, which determine treatment status. Estimates divided by the outcome mean and multiplied by 100. Standard errors are clustered at the user level.

Figure D.25: Difference-in-difference event-study in strongly balanced panel II

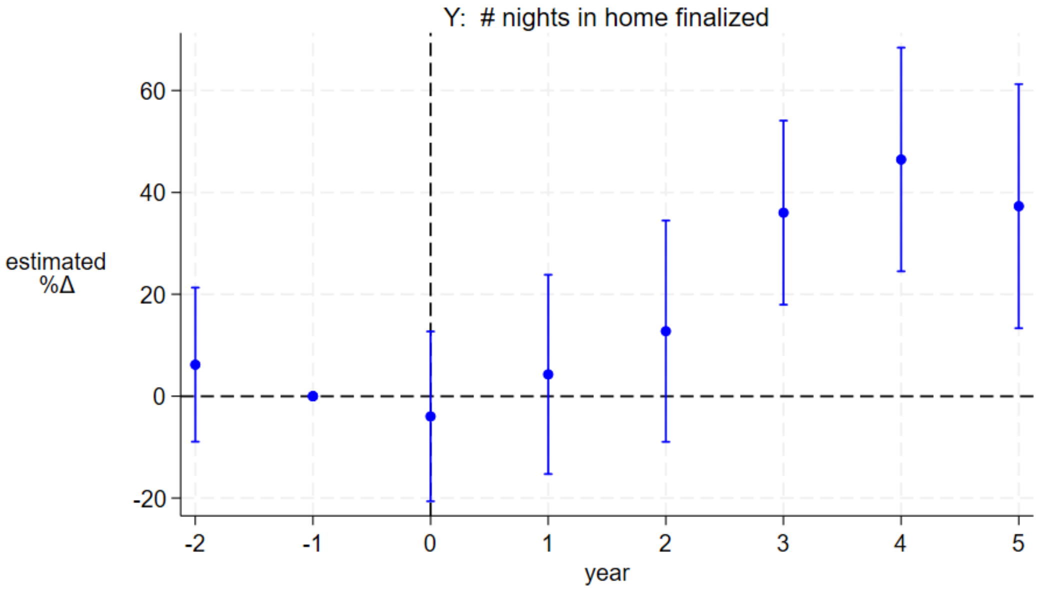

The graph shows difference-in-difference event-study estimates (equation 38) of the effect of the price reduction on treated homes. Prices of homes in the treatment group were reduced by at least 30%. Prices of homes in the control group were reduced by $(-10\%, +10\%)$. The model includes fixed effects for time and for cells of home-types, which determine treatment status. Estimates divided by the outcome mean and multiplied by 100. Standard errors are clustered at the user level.

Figure D.26: Rating hosts receive vs. over-demand

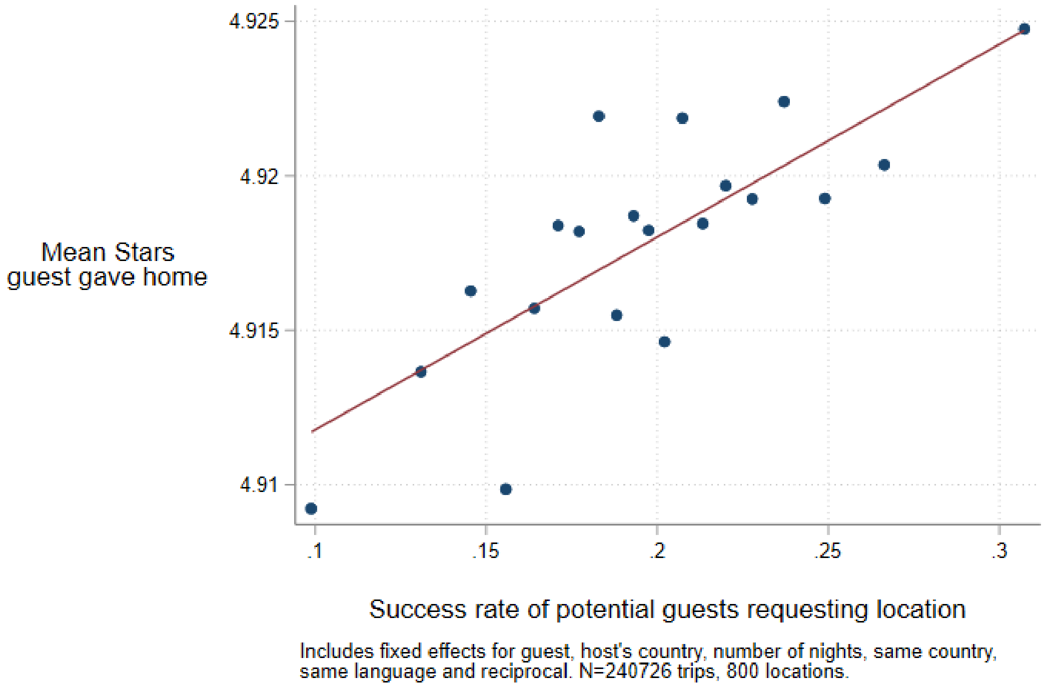

Figure D.27: Requested party size around price reform

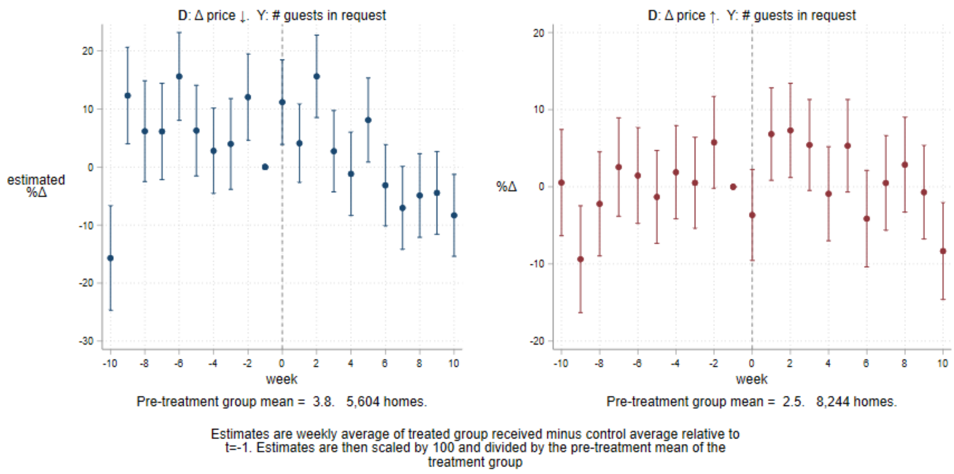

The outcome is the number of guests that the sender indicates in her request. The left panel shows requests to homes whose displayed price dropped by at least 50 % in week 0, right panel homes whose price increased by at least 50% (because all displayed prices were reset to recommended price at the day of the reform).

Figure D.28: Evolution of number of active homes around reform

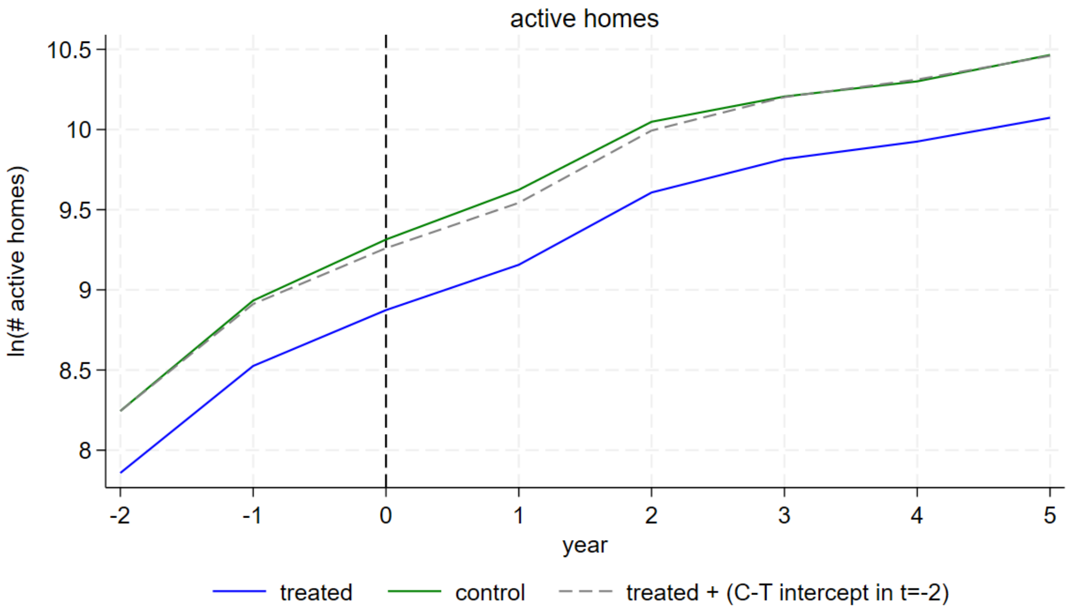

The figure shows how the total number of active homes evolved on the platform, separately for homes in the treatment group and homes in the control group. Active homes are defined as in table 4. The home prices for the treatment group were reduced by at least 30 percent in the reform year. The dashed line shows the evolution of the treatment group shifted upward by the initial difference between both groups.

Figure D.29: Effect of price reduction on participation (proxy for remaining active)

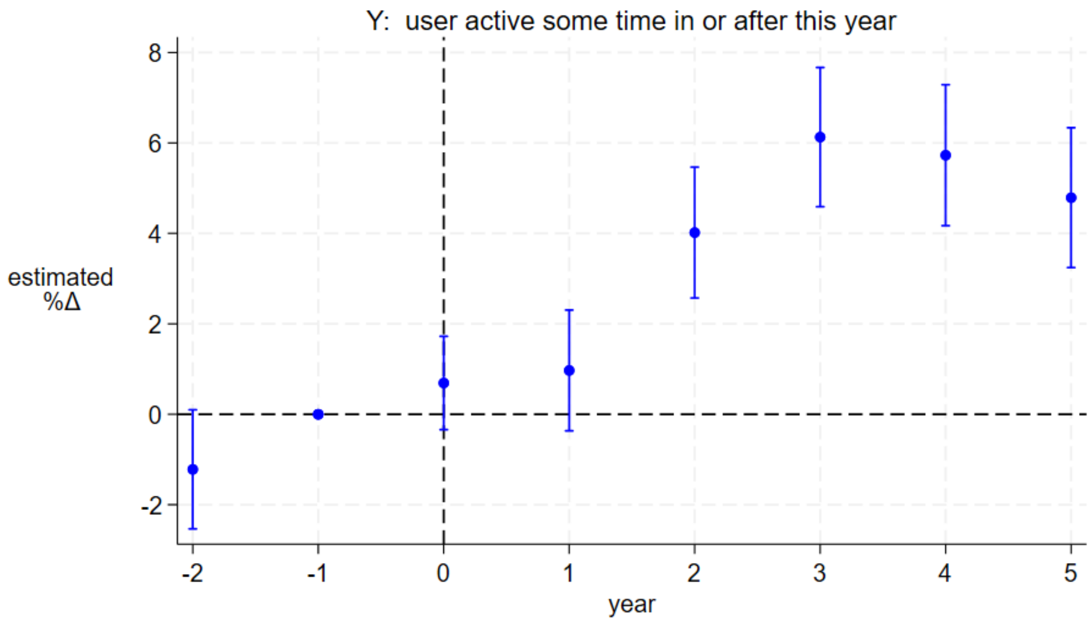

The graph shows difference-in-difference event-study estimates (equation 38) of the effect of the price reduction on treated homes. The outcome is an indicator if the the user is active in or after that year, where sending a message is the minimum required activity. Prices of homes in the treatment group were reduced by at least 30%. Prices of homes in the control group were reduced by $(-10\%, +10\%)$. The model includes fixed effects for time and for cells of home-types, which determine treatment status. Estimates divided by the outcome mean and multiplied by 100. Standard errors are clustered at the user level.

Figure D.30: Sensitivity to differential trends

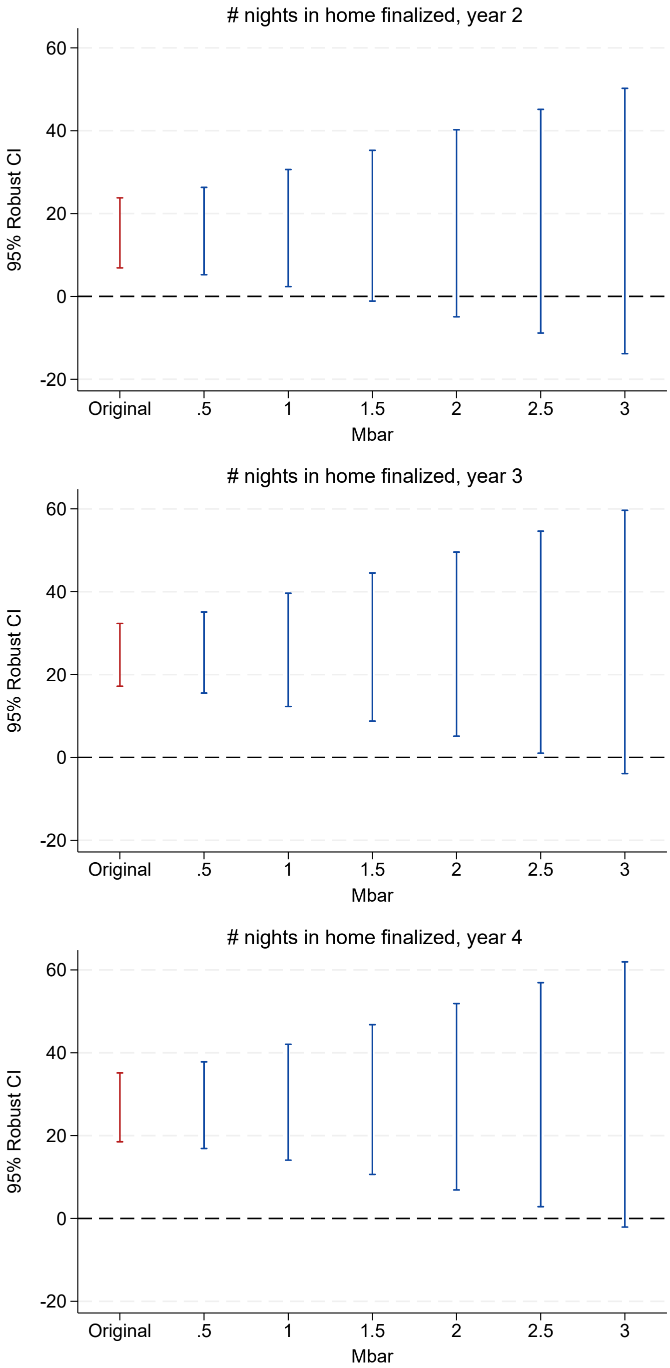

The figure shows the sensitivity of the benchmark event-study estimates in figure 9 to differential trends. Using the honestDiD package developed by (Rambachan and Roth, 2023) I calculate the confidence sets for various values of M. A value of $M = 2$ ($M = 3$) means that differential post-treatment are allowed to be two times (three times) as large as the worst violation in the pre-treatment period. The top panel shows the sensitivity of the event-study estimate for year 2, the middle one for year 3 and the bottom panel for year 4.