The Tragedy of Your Upstairs Neighbors: The Negative Externalities of Home-Sharing Platforms*

Apostolos Filippas John J. Horton

November 2, 2023

Abstract

Home-sharing platforms enable hosts to impose costs on their neighbors, possibly creating a market failure. To explore potential public policy responses, we develop a model of the markets for home-sharing and long-term rentals, and explore the market outcomes under different policy regimes. We show that if the decision to home-share is left to individuals there is too much home-sharing, whereas if the decision is left to a city that maximizes the surplus of long-term tenants there is too little home-sharing. However, when building owners decide on the home-sharing policy of their buildings, the externalities of home-sharing are internalized, and the level of home-sharing activity is socially optimal. Our model predicts that when building owners decide, they will be indifferent between allowing and banning home-sharing in equilibrium. We test this prediction empirically using a dataset of NYC rental apartment listings, and find that this "no policy arbitrage" property holds for similar building-level policies.

JEL Classification Codes: JEL H25, R5, J3
Keywords: home-sharing; short-term rentals; externalities; regulation; public policy

*We gratefully acknowledge generous financial support from support from the Center for Global Economy and Business. Thanks to seminar participants at the NYU Stern Technology Department and Microsoft Research New England. Thanks to Yannis Bakos, Peter Coles, Bo Cowgill, Michael Egesdal, Andrey Fradkin, Michael Luca, Igor Popov, Justin Rao, and Richard Zeckhauser for their helpful comments.

1 Introduction

The benefits of home-sharing platforms, such as Airbnb, VRBO, and Couchurfing are clear: underutilized capacity is put to use, supply flexibility increases, and consumer choice is expanded (Einav, Farronato and Levin, 2016; Filippas, Horton and Zeckhauser, 2020; Farronato and Fradkin, 2022). On the other hand, home-sharing platforms may enable "hosts" to impose costs on their neighbors: if hosts bring in loud or disreputable "guests" but, critically, still collect payment, then home-sharing platforms would seem to help create a case of uninternalized externalities that existing illegal hotel laws are intended to prevent.[1] This potential for "regulatory arbitrage" is a recurrent critique of "sharing economy" platforms more generally, and has been used in support of legislation restricting home-sharing activity, as well as in lawsuits against home-sharing platforms.[2]

The negative externalities of home-sharing are similar to those found in other economic activities: the host gets the money, and the neighbors get the noise. However, typical public policy solutions may fail to address this problem effectively, due to the nebulous property rights in large buildings, the difficulty of identifying offending parties, and the heterogeneity of residents' preferences. Motivated by this public policy question, we develop a model of the markets for home-sharing and long-term rentals, and examine the equilibrium outcomes under policy regimes that differ only in which party has the decision right to home-share. In particular, we examine four regimes where hosting decision rights are allocated to (i) individual tenants, (ii) building owners, (iii) cities, and (iv) a utilitarian social planner. For each regime, we derive the market equilibrium, and characterize the surplus of both hosting and non-hosting tenants, building owners, and guests.

Tenants are they key actors in our model. They make two choices: whether to be home-sharing hosts, and which building to live in. In deciding whether to host, tenants consider only their financial pay-off from hosting: they host if they are allowed, and if the income they receive from home-sharing exceeds their individual hosting costs. Home-sharing income is endogenous in our model, and depends on how many other tenants living in the same city choose to host. In choosing which building to live in, tenants consider the rent they will face, whether they are allowed to home-share, and the negative externality costs they bear from the home-sharing activity of other residents of the building.

We first examine the "tenants decide" (TD) regime, in which tenants are free to decide whether to become home-sharing hosts. Next, we consider the "building owners decide" (BD) regime, in which building owners set blanket policies for their buildings. In choosing whether

[1] For consistency with the literature, we use the term "hosts" to refer to users renting out their properties on home-sharing platforms, and "guests" to refer to users renting properties on home-sharing platforms.

[2] For example, the negative externality argument has been cited in recent legislative action that increased fines dramatically for hosts found to be violating local housing regulations. Furthermore, building management company AIMCO cited the negative externality argument as the main reason for a lawsuit against Airbnb.

to allow or to ban home-sharing, building owners take into account only the effect their policy choice has on the rental rates they can command from long-term tenants. Finally, we consider regimes where a central decision-maker may determine the market quantity of hosting—in practice, this quantity would be set through mechanisms such as taxation, rationed permits, and bureaucratic ordeals (Nichols and Zeckhauser, 1982). We examine two such regimes: the "city decides" (CD) regime, in which the decision maker considers only the surplus of the tenants, i.e., the residents of the city, and the "social planner decides (SD) regime, in which the decision maker takes into account the surpluses of both tenants and guests.

Our analysis shows that when individual tenants decide whether or not to host, there is too much hosting in equilibrium, in that the costs created by the marginal host exceed the benefits. Consequently, the equilibrium after the introduction of home-sharing might offer less surplus than an equilibrium before the introduction. Setting aside for a moment the case where building owners decide, we find that when the city sets the quantity, there is too little hosting. Essentially, the city behaves as a monopolist, reducing supply to raise prices, thereby transferring surplus from guests to hosts. In practice, if cities are "already" picking the profit-maximizing quantity through their regulation and taxation of the hotel industry, the city might find it optimal to ban home-sharing altogether, as the increase in supply is unwanted.[3]

The efficient quantity of hosting is obtained when the home-sharing decision is left to building owners. The driver of this efficiency result is that in equilibrium, the marginal tenant is indifferent between buildings that allow and buildings that prohibit home-sharing, and hence building owners are also indifferent between allowing or prohibiting home-sharing. The reason building owners are indifferent is that rents in a competitive long-term rental market must be the same regardless of the home-sharing policy of the respective building: rents are equal because the building's home-sharing policy imposes no direct cost on the building owner, and if a premium could be charged for one policy or the other, profit-maximizing building owners would choose whatever policy offered the premium. This building-owner self-interest equalizes long-term rental rates, and so the marginal long-term tenant—the one who is indifferent between buildings that allow home-sharing and those that do not—has a private benefit of hosting that is equal to the full costs of living in such a building. The full cost includes not only the tenants' private cost of hosting, but also the costs imposed from home-sharing hosts in the same building. Note that in this analysis, we do not have to model the surplus of the guests explicitly, as the marginal guest surplus at the market-clearing home-sharing price is the same as the private benefit to the host.

Although the model is parsimonious, the core result—the attractive efficiency properties of

[3]The high tax rates on the hospitality industry indicate that cities benefit from reducing hosting supply. For example, see http://www.wsj.com/articles/SB10000872396390443749204578048421344521076. For a list of state lodging taxes see http://www.ncsl.org/research/fiscal-policy/state-lodging-taxes.aspx.

allocating decision rights to building owners—is robust to various model extensions, including adding home-sharing supply that does not generate externalities, modeling externality costs as non-linear, allowing building owners to convert an entire property to home-sharing, giving tenants heterogeneous preferences over buildings, amenities, neighborhoods, and so on. At a high level, the reason for the invariance of our conclusions to these model extensions is that what matters for efficiency is the marginal tenant, and more complex model extensions mostly affect inframarginal market participants.

Despite the robustness of our results to several model extensions, an assumption that is critical to our results is that externalities are contained within a building. There are two strong justifications for this assumption. First, physical nuisances such as noise and smells dissipate with the cube of the distance from the source, making it hard for these kinds of costs to travel very far and remain large. Second, nuisances such as wear-and-tear, misuse of common areas, and reduced physical security, are inherently within-building problems. Despite our view that real externalities are largely contained within buildings, we do show how our model can be adapted to other cost structures.

A key prediction of our model is that, in a competitive equilibrium, building owners cannot command higher long-term rental rates through their home-sharing policy choices. It is worth emphasizing that this prediction is not equivalent to claiming that home-sharing does not increase rents. Rather, the model prediction is that rents will be equal in buildings with different home-sharing policies—albeit possibly higher than in a city without the option to home-share. This "no policy arbitrage" prediction is difficult to assess directly. The reason is that home-sharing is still a nascent phenomenon, and is subject to a constantly-shifting regulatory frameworks. Consequently, observing market equilibrium outcomes is hard, and hence data from existing markets are unlikely to offer a compelling empirical test.

To circumvent this problem, we turn to other policies routinely chosen by building owners that are conceptually similar to the decision to allow home-sharing. We examine two proxy policies. The first policy is the decision of building owners to allow subletting. Subletting is of longer duration than home-sharing, but also has slight administrative costs for the building owner, as well as a potentially large impact on would-be renters and current tenants alike. The second policy is the decision of building owners to allow dogs. This policy has slight administrative costs for building owners, but some tenants value the option to have dogs. Importantly, dogs have the potential to impose substantial negative externalities on neighbors, in the form of barking, biting, allergens, and smells.

Using a large dataset of rental listings in New York City, we find that there is no arbitrage opportunity in choosing a subletting policy. Although allowing subletting is strongly, positively correlated with rental rates, this relationship disappears when including controls. The effect of allowing subletting on rental rates is a precisely estimated near-zero when using the double-debiased machine learning (DML) approach (Chernozhukov, Chetverikov, Demirer, Duflo, Hansen, Newey and Robins, 2017). Similarly, while allowing dogs is highly correlated with higher rental rates, the effect disappears in the DML estimate. Furthermore, to build confidence in our empirical approach, we also show that a premium can be charged for "policies" that are not costless to the building-owner but valued by tenants, such as the inclusion of an in-apartment washer and dryer. We interpret these results as a case study that supports lends some empirical support to the "no policy arbitrage" prediction.

In our model, the role of home-sharing platforms is critical—their emergence is the technological shock that makes home-sharing wide-spread—but also passive with regards to the negative externality problem. Although this passivity is a useful simplification for our analysis, platforms can take an active role in addressing problems created by home-sharing. We identify measures that home-sharing platforms are already taking, and which are in agreement to the predictions of our model, including Airbnb's "resident hosting" initiative.[4] We also suggest measures that platforms can take. For example, platforms managers create tools that allow building owners to centrally impose tenant-specific hosting caps—upper bounds on individual home-sharing activity—which can be particularly important if externalities increase convexly in home-sharing activity (a possibility we discuss).

The main contribution of this paper lies in conceptualizing home-sharing as having the potential to create a market failure, and in developing a tractable model to examine various public policy responses. The empirical analysis supports the key prediction of the model, which in turn builds confidence in our modeling approach. Our approach is distinctive from the growing literature examining offline spillovers of online developments, in that we take spillovers as a given, and then work through their prescriptive implications. Although our analysis focuses on home-sharing, our results also have implications for platform operators who must increasingly navigate the policy landscape while pursuing new business models.[5]

The remainder of this paper is organized as follows. Section 2 reviews previous work. Section 3 develops the model and presents the main results, and Section 4 explores extensions. Section 5 empirically assesses the "no policy arbitrage" prediction. Section 6 discusses the policy prescriptions of the model. Section 7 concludes.

2 Background

Short-term rentals of personal spaces have long been possible (Jefferson-Jones, 2014; Kaplan and Nadler, 2015). Recently, a series of technological and entrepreneurial developments

[4] https://www.airbnb.com/help/article/1195.

[5] Recent controversies around for-hire vehicle caps in NYC (arguably intended to reduce congestion), and electric scooter bans (arguably intended to reduce sidewalk blockages), suggest that our "negative externality of the online platform business" focus is far from a one-off issue for would-be platform managers and entrepreneurs. For example, see https://www.kxan.com/news/local/austin/woman-s-post-about-scooters-blocking-her-path-leads-to-new-program/1386887743.

have dramatically increased the scale of home-sharing, subsequently sparking an ongoing policy debate between platforms and regulators. Home-sharing is one example of "sharing economy" platforms that span a wide range of industries, including car- and ride-sharing, micro-loans, and startup funding, and generate billions in revenue annually (Brynjolfsson, Hu and Simester, 2011; Malhotra and Alstyne, 2014; Sundararajan, 2016; Dinerstein, Einav, Levin and Sundaresan, 2018; Filippas et al., 2020; Filippas, Horton and Golden, 2022).

2.1 Offline effects of online platforms

Much of the previous work on the offline effects of online platforms has examined the effects of the entry of online platforms on offline competitors, including changes in market shares and prices (Seamans and Zhu, 2013; Kroft and Pope, 2014; Zervas, Proserpio and Byers, 2017; Farronato and Fradkin, 2022). Insofar that these effects are solely on prices, the waxing and waning of various industries does not constitute a market failure: every transaction has a buyer and a seller, and changes in price have offsetting changes in utility for the demand and the supply sides of the market. Pecuniary externalities—the effects that changes have on prices (Scitovsky, 1954; Laffont, 1989, 2008)—are distinguished in the literature from non-pecuniary externalities (also "technological," or "real"). Non-pecuniary externalities are unpriced costs and benefits, and have the potential to lead to market failure, in the sense that the decentralized market equilibrium may be characterized by inefficiently small quantities if externalities are positive, or inefficiently large quantities if they are negative.[6]

There are numerous examples of online platforms creating offline non-pecuniary externalities. In the public health sphere, Chan and Ghose (2014) present evidence that by reducing the search costs for casual sex partners, the entry of Craigslist likely caused about a 16% increase in HIV cases—at enormous social cost. As an example of a positive externality, Greenwood and Wattal (2017) exploit differences in the timing of Uber's introduction into cities in the state of California to investigate its effect on DUI arrests. They find that the effect was significant, resulting in about a 4% decrease in the rate of motor vehicle homicides. On the other hand, a negative externality of car-sharing platforms is that they exacerbate traffic congestion in urban centers (Clewlow and Mishra, 2017; Molnar and Mangrum, 2018).

[6]In contrast to non-pecuniary externalities, pecuniary externalities are often the result of positive change. For example, Sheppard and Udell (2018) provide evidence that increases in Airbnb availability are associated with increased house values—implying that home-sharing has pecuniary externalities—but also note that "Public policies that reduce house prices in pursuit of housing affordability by diminishing the efficiency with which an owner can make use of his or her property may fail to be welfare-improving, in the same way as a city that creates "affordable" housing by encouraging more crime hardly seems desirable."

2.2 Regulatory responses to home-sharing

The pecuniary externalities of home-sharing platforms are the changes in price and value brought about by the entry of the home-sharing option in a city, such as to hotels, property values, and long-term rental rates (Cusumano, 2015; Guttentag, 2015; Zervas et al., 2017; Sheppard and Udell, 2018; Farronato and Fradkin, 2022). The main policy import of these externalities is that they may have distributional consequences for different groups—owners versus renters, residents versus hosts, and so on. Several studies have established a positive effect of home-sharing on rents and house-prices. Barron, Kung and Proserpio (2017) employ an IV framework and find that a 100% increase of Airbnb activity is associated with a 1.8% increase in rents and a 2.6% increase in house values; Horn and Merante (2017) obtain similar estimates, whereas Sheppard and Udell (2018) place the house value effect estimate between 6% and 11%. These estimates are hard to interpret given the rapid growth of home-sharing, especially with regards to their temporal stability—recent studies report smaller estimates.

The non-pecuniary externalities of home-sharing describe the costs that hosts' neighbors incur due to guests. Frequently cited by critics of home-sharing, these externalities include noise, increased use of common resources, unruly behavior, threat posed by strangers, and changing the "character" of neighborhoods. To the extent that these externalities exist, public policy solutions, such as side payments, Coasian bargaining, and Pigouvian taxes, could address them in theory (Coase, 1960; Polinsky and Shavell, 1982). However, implementing these solutions can be hard in practice, given the requirement for transfers between all affected parties after every transaction, the nebulous property rights in large buildings, the difficulty in identifying offending parties, the heterogeneity in residents' preferences, and the potential for opportunistic behavior arising from side payments.[7] For that reason, we focus on potential market-based policy responses.

3 A model of the non-pecuniary externalities of home-sharing

3.1 A city without home-sharing

Consider a city comprising a fixed set of tenants $\mathcal{I}$ and a set of apartment buildings $\mathcal{A}$. Each tenant $i \in \mathcal{I}$ consumes one unit of housing capacity, and obtains utility u_0 from living in the city.[8] Each building $\alpha \in \mathcal{A}$ houses a mass of n tenants, and belongs to a building owner also indexed by α and who sets a rental rate $r_\alpha \geq 0$. We assume that there is free entry in housing development which implies that building owners make zero profit, that $\mathcal{A} = [0, A]$,

[7] "That's a nice quiet apartment you have there—it would be a shame if someone started playing loud music late at night."

[8] All quantities in our model should be thought of as flows per time period, such as yearly costs and benefits.

and that the city has fixed population level $N = nA = |\mathcal{I}|$.[9]

We begin by characterizing the equilibrium in the city without home-sharing. Our solution concept requires tenants to be indifferent between living in different buildings. Spatial indifference implies that $r_\alpha = r_0$ for all $\alpha \in \mathcal{A}$, and free developer entry yields $r_0 = 0$. Then every tenant obtains utility equal to u_0, and total tenant surplus equals $U_0 = Nu_0$. This surplus will be the welfare baseline throughout our analysis.

3.2 A city with home-sharing

When home-sharing becomes possible tenants and building owners decide whether to participate in the home-sharing economy. Let p be the going home-sharing price, which is also the benefit to a tenant who hosts. On the supply side, each tenant i can offer one unit of home-sharing supply and has an individual-specific hosting cost c_i, meaning that she is willing to host if $p \geq c_i$. The distribution of hosting costs is $F : [0, \overline{c}] \rightarrow [0, 1]$, and is differentiable with density f that is positive everywhere in its support. The home-sharing market supply $S(p)$ is the number of tenants $NF(p)$ that would host at price p, and $\hat{c}(q)$ denotes the hosting cost of the marginal host when q units of home-sharing are supplied.

On the demand side, there exists a set of consumers $\mathcal{J}$ of mass N. Each consumer $j \in \mathcal{J}$ has valuation v_j, and is willing to participate in home-sharing and become a guest if $v_j \geq p_j$. The distribution of consumer valuations is $G : [0, \overline{v}] \rightarrow [0, 1]$, and is differentiable with density g that is positive everywhere in its support. The home-sharing market demand $D(p)$ is the number of consumers $N(1 - G(p))$ that would become guests at price p. The valuation of the marginal guest when q units of home-sharing are supplied is $\hat{v}(q)$. We assume that the market clears at some price and quantity where neither a glut nor a shortage occurs.

Home-sharing activity imposes costs on a host's neighbors. Each host generates externality cost c_E to every other tenant living in the same building. We say that the hosting activity of tenant i is *socially efficient* if

$$p \geq c_i + nc_E. \tag{1}$$

Equation 1 states an intuitive criterion for assessing the impact of a host's activity: if the negative externalities generated are less than the host's private benefit—home-sharing price minus the hosting cost—then the hosting activity of that tenant is socially efficient.

Under some policy regimes, building owners may choose to allow or prohibit home-sharing in their building. We assume that allowing home-sharing is costless for building owners. The decision of building owner α is denoted by the indicator variable $h_\alpha \in \{0, 1\}$, and the fraction of buildings that allow home-sharing is $h = \frac{1}{A} \int_{a \in \mathcal{A}} h_a$.

[9] Each building α can be though of as lying at a distance α from the center of a monocentric city. This interpretation is not necessary for the analysis of this section, but will be useful in latter sections.

An equilibrium is a market configuration where (i) tenants do not want to move to a different building, (ii) owners cannot profitably change their home-sharing policies or long-term rental rates and induce tenants to move, and (iii) the home-sharing market clears. To simplify our analyses, we focus on equilibria that are symmetric in the sense that building owners who set the same home-sharing policies set equal rental rates. We examine the equilibria under four different policy regimes corresponding to potential regulatory responses to home-sharing. We then study extensions in Section 4.

3.3 The "tenants decide" (TD) regime

We first consider the policy regime where owners are not allowed to prohibit home-sharing in their buildings, and hence all tenants are allowed to individually decide whether to host. In the language of our model, $h_a = 1$ for every owner $a \in A$, and $h = 1$. The TD policy regime has undesirable properties and can create a market failure.

Proposition 1. *The equilibrium of the TD regime is characterized by a unique tuple (p_T, q_T) and the following properties: (i) there exist tenants whose hosting activity is socially inefficient, and (ii) home-sharing activity is equally dispersed across buildings.*

Proof. Every tenant who wants to be a participate in home-sharing is allowed, and hence tenant i hosts if $c_i \leq p$. Due to the monotonicity and continuity properties of the supply and the demand, a unique market-clearing tuple (p_T, q_T) exists, where $p_T = \hat{c}(q_T)$. Note that $q_T < N$ and $p_T > 0$ by our assumption that the home-sharing market clears. Because $nc_E > 0$, it follows that $p_T < \hat{c}(q_T) + nc_E$, and Equation 1 does not hold for all hosts in equilibrium, which implies that there exist tenants whose hosting is socially inefficient. For tenants to be indifferent between buildings, they should not be able to incur a lower externality cost by moving to another building; as such, an equal share of home-sharing activity takes place in each building. □

The TD policy regime has several drawbacks. First, there are inefficiently many hosts: hosts with hosting costs in the interval $(p_T - nc_E, p_T]$ generate externalities that outweigh their individual benefits. Second, externalities are not internalized: hosts occupy apartments in every building, and hence tenants who do not host incur externalities and see their utilities decrease to $u_0 - \frac{S(p_T)}{A} c_E$.

A market failure may occur in the TD equilibrium. The total tenant surplus is

$$U_T = U_0 + \Big(\int_0^{q_T} p_T - \hat{c}(q)\, dq \Big) - q_T n c_E. \tag{2}$$

The first term in Equation 2 is the surplus due to tenants occupying apartments, the second term is the net surplus generated from hosting (market price minus hosting costs), and the last term is the sum of the home-sharing externalities. Tenant surplus decreases if the externalities generated outweigh the sum of the hosts' benefits.

Figure 1 illustrates this situation. Point EQ_1 indicates the TD equilibrium. The total cost curve $\hat{c}_t(q) = \hat{c}(q) + nc_E$ captures the social cost of the marginal listing when q units of home-sharing are supplied. It is the difference between the individual hosting cost and the total cost that creates the potential for market failure. The light gray area depicts the positive contribution to the aggregate tenant surplus, which is due to those tenants with low enough hosting cost that their profit from hosting outweighs their individual hosting cost plus the negative externalities of their hosting. The dark gray area depicts the negative contribution to the total tenant surplus, coming from those tenants with a low enough hosting cost to still want to host at the equilibrium price, but not low enough to outweigh the sum of their individual costs and the externality costs.

Figure 1: Tenant surplus in the "tenant decides" market equilibrium.

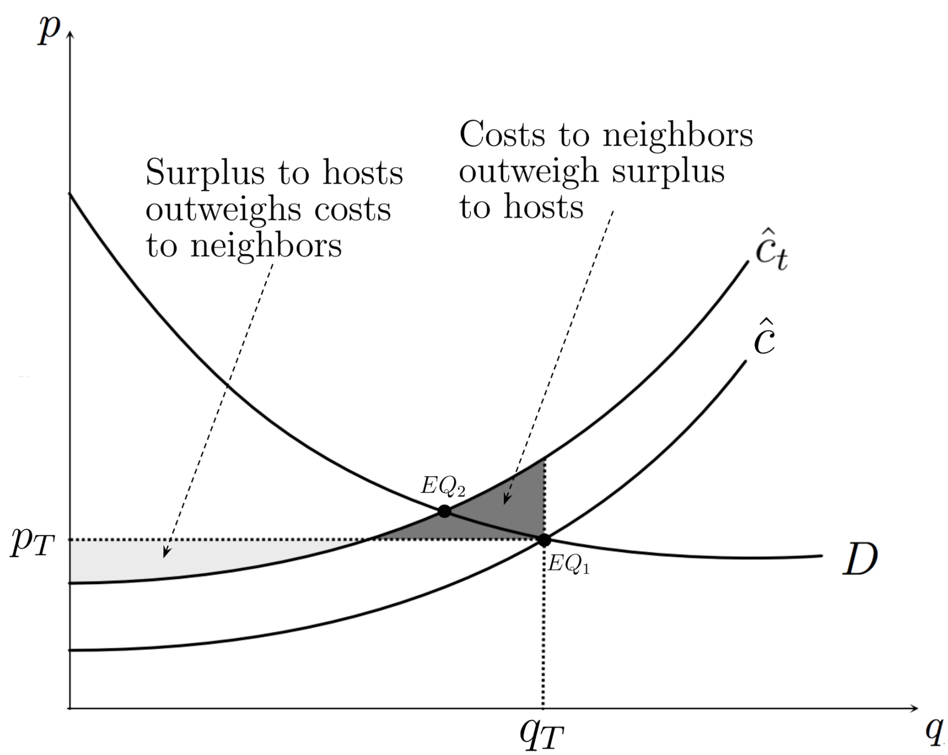

The potential for market failure relates directly to the elasticity of the supply curve. Consider the scenario where supply is perfectly elastic, that is, $c_i = c_H$ for every tenant i. In this case, $p_T = c_H$, and consequently $\int_0^{q_T} p_T - \hat{c}(q)\, dq = 0$ and $U_T < U_0$. As the supply elasticity decreases and holding all other factors fixed, it becomes more likely that the aggregate supply-side welfare will not decrease.[10]

[10]If we assume that hosting costs are drawn from some distribution, the same intuition holds for the variance

3.4 The "building owners decide" (BD) regime

We next examine the policy regime where building owners may choose to either prohibit home-sharing in their buildings, or to allow all tenants to host. Note that a tenant living in a home-sharing-friendly building will not host if her hosting cost exceeds the home-sharing price. We show that market efficiency is recovered in the BD equilibrium.

Proposition 2. *The equilibrium of the BD regime is characterized by a unique tuple* (p_B, q_B)*, and the following properties: (i) rental rates are equal across buildings with different home-sharing policies, (ii) all tenants' hosting activity is socially efficient, and (iii) externalities are internalized.*

Proof. Let h_B be the fraction of buildings that allow home-sharing, $q_B = h_B An$ be the corresponding maximum home-sharing supply, r_1 be the long-term rental rate in buildings that allow home-sharing, and r_0 be the long-term rental rate in building that do not. Free developer entry implies that $r_0 = 0$. Non-hosting tenants are spatially indifferent only if they live in buildings that prohibit home-sharing, as they would otherwise be better off moving to buildings without home-sharing and incur no externalities. This implies that all tenants in home-sharing-friendly buildings host, and that q_B is the actual home-sharing supply. Furthermore, if a tenant i with hosting cost c_i hosts in equilibrium, then a tenant j with hosting cost $c_j < c_i$ must also host; otherwise, either tenant i or tenant j are not indifferent between buildings with different policies. Consider the home-sharing price p_B and the marginal host who has hosting cost $\hat{c}(q_B)$. If $p_B < \hat{c}(q_B) - nc_E$, then the marginal host would be better of moving to a building that does not allow home-sharing. If $p_B > \hat{c}(q_B) - nc_E$, then the marginal host obtains a positive premium from living in a home-sharing building, and there either exist non-hosting tenants who are not spatially indifferent, or the no profit condition does not hold. Therefore $p_B = \hat{c}(q_B) + nc_E$ in equilibrium, with the market-clearing quantities being unique due to the monotonicity and continuity properties of the supply and the demand. Because the marginal host makes zero profit, then $r_1 = r_0 = 0$. Equation 1 holds for every host in the BD equilibrium, and only tenants who host are subject to home-sharing externalities. □

The BD policy corrects the shortcomings of the TD policy. First, home-sharing externalities are internalized: because tenants are sorted according to their preferences for home-sharing, only hosts incur negative externalities from home-sharing. Second, all hosting

of that distribution. As the variance of the distribution of hosting costs tends to zero, market failure becomes more likely. As the variance of the distribution increases, U_T increases as well, making it more likely that $U_T > U_0$.

activity is socially efficient. The surplus of tenants in the BD equilibrium is

$$U_B = U_0 + \left(\int_0^{q_B} p_B - \hat{c}(q)\, dq \right) - q_B n c_E. \tag{3}$$

Tenant surplus never drops below that of a market without home-sharing.

Proposition 3. *In the BD equilibrium, tenant welfare (weakly) increases compared to the market without home-sharing; that is,* $U_B \geq U_0$.

Proof. We have

$$U_B - U_0 = \left(\int_0^{q_B} p_B - \hat{c}(q)\, dq \right) - q_B n c_E = \int_0^{q_B} p_B - \hat{c}(q) - n c_E\, dq \geq \int_0^{q_B} p_B - \hat{c}(q_B) - n c_E\, dq,$$

where the inequality is due to $\hat{c}$ being increasing in q. By definition, $p_B = \hat{c}(q_B) + n c_E$ so the last expression is equal to zero, proving our result. □

As in the TD equilibrium, the increase in tenant surplus in the BD equilibrium is rooted in the heterogeneity of hosting costs. Consider again the case where $c_i = c_H$ for all tenants i. We get that $p_B = c_H + n c_E$, and as a result, $U_B = U_0$. This showcases the robustness inherent in the BD equilibrium: tenant surplus does not decrease even under the worst-case distribution of hosting costs, where supply is perfectly elastic.

Tenant surplus under the BD regime always compares favorably to that of the TD regime.

Proposition 4. *In the BD equilibrium, tenant welfare (weakly) increases compared to the TD equilibrium; that is,* $U_B \geq U_T$.

Proof. Subtracting the two quantities gives us

$$U_B - U_T = (p_B - p_T) q_B + \int_{q_B}^{q_T} \hat{c}(q) + n c_E - p_T\, dq \geq (p_B - p_T) q_B + \int_{q_B}^{q_T} \hat{c}(q) + n c_E - p_B\, dq.$$

The first term is nonnegative as $p_B \geq p_T$. Since $\hat{c}$ is increasing and $p_B = \hat{c}(q_B) + n c_E$ by definition, the integrand is nonnegative on the $[q_B, q_T]$ interval. This proves the result. □

To see why building-specific policies improve upon the supply-side surplus of the TD regime, we need to observe that there are two channels through which an additional home-sharing listing may decrease tenants' surplus: (i) the listing imposes externalities greater than the corresponding benefits, and (ii) the corresponding benefits are less than the utility lost among all previous hosts due to the decrease in price that the higher supply results in. We showed that Equation 1 holds for all hosts in the BD equilibrium, and hence no tenant surplus is lost due to excessive hosting externalities. While this condition is sufficient to guarantee that tenant surplus always increases compared to either the TD market or the market with no home-sharing option, it does not imply that tenant surplus is maximized.

3.5 The "city planner decides" (CD) regime

In the rest of this section we turn our attention to centralized decision-makers who may control home-sharing supply, through means such as issuing individual- or building-level permits, taxing, and imposing transaction costs to home-sharing. We assume that centralized planners can allocate the right to home-share to the tenants that have the lowest-hosting costs.

We first examine city-level regulatory bodies, which we refer to as the "city planner decides" (CD) policy regime. We model the city planner as having the incentive to maximize tenant surplus. Our choice is motivate by the fact that city planners collect taxes from accommodation-related activities (see Footnote 3), and that city residents—not guests—shape voting outcomes on the city level. The city planner's intervention in the home-sharing market lowers the supply relative to that of the BD regime.

Proposition 5. *Home-sharing supply q_C in the CD regime is (weakly) lower than the home-sharing supply q_B in the BD regime.*

Proof. The quantity that maximizes tenant surplus is found by solving

$$\underset{q\in[0,q_B]}{\arg\max} \quad \left(\int_0^q p(q) - \hat{c}(x)\,dx\right) - qnc_e. \tag{4}$$

Note we can impose the upper bound q_B on the feasible region without loss of generality, as tenant surplus strictly decreases for quantities greater than q_B. The optimal solution q_C satisfies the optimality condition

$$\frac{\partial p}{\partial q}q + p(q) = \hat{c}(q) + nc_E, \tag{5}$$

which states that the quantity q_C is that where the marginal revenue (left-hand side) equals the marginal cost (right-hand side). Since $\frac{\partial p}{\partial q} \leq 0$, the city planner potentially restricts the number of home-sharing buildings, and we get $q_C \leq q_B$ and $p_C \geq p_B$. □

In the CD regime, there exist tenants who are prohibited from hosting, but whose value from hosting would be greater than the corresponding marginal social cost. As a result, while tenant surplus in the CD regime is maximized, this surplus is distributed to fewer tenants—those tenants with the lowest hosting costs. Lower supply implies higher home-sharing prices, and hence guest surplus decreases as well, creating a welfare transfer from guests to tenants. For $q \in [0, q_C)$ both market sides incur losses, although surplus never becomes negative for neither side, and hence social welfare does not drop below that of a market without the home-sharing option.

It is worth making two observations. First, the marginal host under the CD policy regime obtains a positive premium from home-sharing. This premium can be captured building owners through higher rents, or by the centralized planner through costly home-sharing permits. Second, although the optimal quantity is invariant to how the city planner controls the market supply, other properties of the market equilibrium depend on the exact mechanism through which supply is restricted. We elaborate on this point in Section 6.

3.6 The "social planner decides" (SD) regime

We now consider a central planner who can control home-sharing supply, but optimizes for the surplus that home-sharing creates on both the supply and the demand sides of the market. We refer to this case as the "social planner decides" (SD) policy regime. The social planner's choice coincides with the BD equilibrium.

Proposition 6. *Home-sharing supply q_S in the SD regime is equal to the home-sharing supply q_B in the BD regime. As such, the optimal social welfare is obtained in the BD equilibrium.*

Proof. The social welfare maximization problem is

$$SW = \max_{q \in [0, q_B]} \quad \int_0^q p(q) - \hat{c}(x) - nc_E dx + \int_0^q \hat{v}(x) - p(q) dx. \tag{6}$$

It is straightforward to show the maximizer of Equation 6 satisfies $\hat{c}(q) + nc_E = \hat{v}(q)$. This condition also holds for $q = q_B$, and hence the BD equilibrium quantity maximizes social welfare, that is $SW = U_B$. The monotonicity of the supply and demand curves guarantee that this is also the unique optimal solution. □

Proposition 6 reveals an important advantage of the BD policy regime: in equilibrium, not only are hosting externalities internalized, but also hosting quantity is optimal with respect to social welfare. Figure 1 illustrates this situation, where point EQ_2 indicates the BD (and SW) equilibrium. From a social welfare perspective, too much home-sharing is allowed in the TD regime, and too little home-sharing is allowed in the CD regime.

4 Extensions

As in any model, we make assumptions and leave out real-world complexities. We next examine how alternative assumptions and extensions to the base model change its predictions.

4.1 Alternative forms of within-building externalities

Our base model assumes that the marginal externality cost of a host to neighboring tenants is fixed. Another possibility is that externalities increase superlinearly in the hosting activity that takes place within a building. We next consider the case where the externality costs tenants incur are convex in the number of hosts living in the same building.

Toward that end, suppose that each tenant incurs externality cost $c_E(x)$, where x is the number of tenants that host within the same building, with $c_E(0) = 0$, $c'_E > 0$, $c''_E > 0$, and $c_E(x) > xc_E$ for all $x \in (0, n]$. The convex cost assumption does not change the positive sorting and efficient hosting properties of the BD equilibrium, but the BD equilibrium welfare is no longer socially optimal.

Proposition 7. *The equilibrium of the BD regime with convex, superlinear externalities is characterized by a unique tuple (p_B^c, q_B^c), with $q_B^c < q_B$ and the following properties (i) rental rates are equal across buildings with different home-sharing policies, (ii) all tenants' hosting activity is socially efficient, (iii) externalities are internalized, and (iv) social welfare is optimal amongst building-specific policies, but not socially optimal.*

Proof. All arguments invoked in the proof of Proposition 2 carry through to the case of convex externalities. The equilibrium home-sharing supply and demand is the tuple (p_B^c, q_B^c) that satisfies $p_B^c = \hat{c}(q_B^c) + c_E(n)$, which is the spatial indifference condition for the marginal tenant. The same condition implies that changing the number of home-sharing friendly buildings will decrease social welfare, and hence the BD equilibrium results in the highest social welfare among allocations utilizing building-specific policies. Because $c_E(n) > nc_E$, we get $q_B^c < q_B$. To show that the BD regime is no longer socially optimal, let h_B^c be the fraction of home-sharing-friendly buildings in the BD equilibrium, and note that the total externality cost generated in the BD equilibrium equal $h_B^c nAc_E(n)$. In the allocation where supply is kept fixed but hosts are equally distributed among buildings the total externality cost is reduced to $nAc_E(h_B^c n)$, while all other quantities remain equal. $\square$

With convex externality costs, Proposition 6 no longer holds: convex costs create a centralized planner incentive to minimize the number of hosts within the same building. The optimal allocation spreads hosts equally across all buildings, which is, by definition, impossible through building-based home-sharing policies. As such, with convex costs there exists a tradeoff between maximizing surplus and internalizing externalities. Furthermore, the surplus-optimal allocations suffer from implementability issues, which we discuss in Section 6.

4.2 Neighborhood- and city-level externalities

One consideration potentially relevant to policy decisions at the city level, and that is not captured in our framework, is the impact of guests on the local economy. The positive impact from every additional guest is not only generated through lodging payments, but also through activities such as dining, shopping, and sightseeing. Incorporating this additional benefit to our model would push the tenant-optimal fraction of home-sharing supply to be higher than θ_C. However, it is important to note that these effects are pecuniary externalities, meaning there is unlikely to be a market failure rationale for considering these effects.

The optimal home-sharing quantity under the presence of these positive, system-wide externalities directly depends on additional assumptions on guest behavior. We consider the special case where each guest has an individual-specific budget b_i for their trip, spends an amount p for accommodation, and the remaining budget, $b_i - p$, on city activities.[11] Following through with the analysis of Section 3.6, we can then show that the BD equilibrium is optimal for the local economy. Furthermore, as listings on home-sharing platforms are more geographically dispersed than hotels, these benefits are also likely to be more geographically spread out (Coles, Egesdal, Ellen, Li and Sundararajan, 2019). However, it is worth noting that this increase in consumption is possibly offset by a decrease in consumption of other activities.

Outside-building externalities of guests may also be negative. For example, extraordinarily noisy guests may impose negative externalities to tenants residing in neighboring buildings. However, our view is that between-building externalities are likely to be small in magnitude. First, physical nuisances such as noise and smells dissipate with the cube of the distance from the source, making it hard for these kinds of costs to travel very far, and certainly not to neighboring buildings. Second, nuisances such as wear-and-tear, misuse of common areas, and reduced physical security, are inherently within-building externalities.

If we do assume that between-building externalities exist and they are negative, the optimal amount of home-sharing would change. Consider a configuration where all buildings exist in a line and that spillovers occur to the buildings left and right of the focal building. In the simplest case, we may assume that the marginal guest in building i does not only impose a negative externality c_E on every tenant of building i, but also a fraction $\alpha < 1$ of this externality on each tenant of buildings $i-1$ and $i+1$. An immediate implication of such externalities is that the gap between the cost curve $\hat{c}$ and the total cost curve $\hat{c}_t$ grows by a factor of 2α (see Figure 1). As a result, the socially optimal home-sharing quantity would decrease. At the same time, the TD equilibrium q_T would remain unaffected, as individual decision-makers only care about their own profit, and market failure would be more likely to occur.

In the BD regime, the presence of negative outside-building externalities implies that, in all non-trivial cases, there will be some tenants who do not participate in the sharing economy but who incur externalities. Equilibrium rents are now not equalized, but rather depend on the number of buildings that allow for home-sharing. In the example of linearly ordered buildings, there now are three equilibrium rents, reflecting the three potential states a building can be in relative to its "neighbor" buildings: a building can have one, two, or zero

[11] Internal Airbnb studies have shown that the average Airbnb guest stays two days longer and spends an additional \$200 on local businesses, compared to tourists staying in hotels (see http://www.airbnb.com/press/news/new-study-airbnb-generated-632-million-in-economic-activity-in-new-york). Furthermore, Alyakoob and Rahman (2022) show that increased home-sharing activities has a positive and salient impact on restaurant employment in New York City.

adjacent buildings that allow for home-sharing, with average rents declining in the number of adjacent home-sharing buildings.

While the exact characterization of the new equilibria hinges upon additional assumptions, an interesting case is that of would-be hosts who experience lower externalities, i.e., c_i is correlated with $c_{E,i}$. We can show that there now exists a unique equilibrium where buildings that allow home-sharing cluster: non-hosts incur higher between-building externalities than would-be hosts, and are willing to pay more to live away from home-sharing-friendly buildings. Therefore, the equilibrium of the BD regime remains socially efficient. Furthermore, this result straightforwardly extends to general topologies, such as grids.

4.3 Reaching equilibrium

The externality problem is "fixed" in the BD equilibrium by tenants moving to buildings that match their "type:" tenants who wish to host move to buildings that allow for home-sharing, and those who do not move to buildings that prohibit it. As such, the preference elicitation and tâtonnement mechanisms are similar to the "foot voting" proposed by Tiebout (1956).[12]

A potential challenge with tenant sorting is the costs tenants incur to move to apartments with the appropriate home-sharing policy under the BD regime. To study the tâtonnement process, we develop an agent-based model of a market under the BD regime. We find that convergence to the BD equilibrium is rapid, under several initial conditions and behavioral assumptions. However, moving costs can decelerate the tâtonnement process, and decrease the efficiency of the resulting equilibrium. The efficiency decrease comes from tenants who are "locked in," and cannot move to buildings with their preferred home-sharing policy.

The agent-based model also allows us to examine other real-life factors that can affect the BD equilibrium. For example, within-building tenant "type" correlation—tenants with similar hosting costs living in the same buildings—accelerates convergence to the BD equilibrium. The details of the agent-based model and all results can be found in Appendix A.

5 Assessing the "no policy arbitrage" prediction

5.1 Empirical strategy

Profit-maximizing building owners should be indifferent in a competitive equilibrium over policies that have no cost to them, even if these policies have within-building negative exter-

[12]Our model departs from Tiebout (1956) in at least two ways. First, a building owner's home-sharing policy directly affects all other buildings owners and tenants, because additional supply reduces the home-sharing price, and hence decreases the home-sharing benefit for all hosts. Second, the number of buildings with the "right" policy is determined endogenously in our setting, and the size of each building is fixed; as such, there is no need to assume that sufficient quantities of communities and tenant "types" exist.

nalities for tenants.[13] In the case of home-sharing, this "no policy arbitrage" prediction is challenging to assess empirically for two reasons. The first reason is that home-sharing is a nascent phenomenon, and the legislative framework surrounding it is changing continually. As such, we lack data on building-level home-sharing policy decisions, and we are unlikely to observe long-run equilibrium market outcomes. The second reason is the fundamental problem of causal inference: it is not possible to observe rental rates for the same building at the same time under two different home-sharing policies.

To circumvent the first problem—that home-sharing is still a nascent phenomenon, and home-sharing policies are not observable in data from existing rental markets—we turn to proxy policies that are conceptually similar to home-sharing policies. The first proxy policy we use is the building owner's decision to allow or prohibit subletting. Although of longer duration than home-sharing, setting a subletting policy is conceptually similar to home-sharing, as it has slight administrative cost implications for the building owner, a large financial impact for tenants who sublet, and potentially large negative externalities for neighbors. The second proxy policy we use is the building owner's decision to allow or prohibit dogs. Allowing or banning dogs has slight administrative costs for building owners, but some tenants value this option; importantly, dogs have the potential to impose substantial negative externalities on neighbors, in the form of barking, biting, allergens, and smells.

To circumvent the second problem—the need to observe counter-factual rental rates under different policies—we construct a predictive model that estimates the long-term rental rates of listings. The predictive model predicts what an apartment "should" rent for based on fundamental, non-policy features. The idea is simple: if we observe the rental rate for an apartment in building A that allows subletting, the rental price in a building B across the street that prohibits subletting might provide us with a good counter-factual. However, building B might not have a roof garden or the square footage might be smaller; these factors could affect the rental rate, which would in turn reduce B's usefulness as a counter-factual. But a predictive model that accurately predicts rental rates based on building attributes can account for differences in amenities, insofar that the values of different amenities and dis-amenities are common in the market.

To fix ideas, consider a listing i for which we observe rental rate $\log r_i$, a set of non-policy features X_i, and a policy of interest $\textsc{Policy}_i$. We can decompose the rental rate as

$$\log r_i = \beta \textsc{Policy}_i + \log \rho_i + \epsilon, \tag{7}$$

where $\log \rho_i$ is what the apartment "should" rent for based only on the non-policy features

[13] It is worth stressing that the "no policy arbitrage" prediction *is not equivalent* to stating that home-sharing will not result in higher long-term rental rates. Rather, it implies that buildings with different home-sharing policies will have equal long-term rental rates— albeit potentially higher than in a city without home-sharing.

X_i, and ϵ is an idiosyncratic error, with $\mathbf{E}[\epsilon] = 0$. In our context, the parameter β is the premium of allowing home-sharing, and the "no policy arbitrage" prediction is that $\beta = 0$.

A "naive" regression of $\log r_i$ on SUBLET alone—without including $\log \rho_i$ on the right hand side—would yield a biased estimate for $\hat{\beta}_{\text{naive}}$ due to omitted variable bias. We can reduce this bias by approximating $\log \rho_i$ using a predictive model trained on the non-policy features X_i, such that $\log \rho_i = \widehat{\log \rho_i} + \eta$, with $\mathbf{E}[\eta] = 0$, residualizing the rental rates, and performing the regression

$$\log r_i - \widehat{\log \rho_i} = \beta \text{POLICY}_i + \eta + \epsilon. \tag{8}$$

The resulting estimate $\hat{\beta}_{\text{hedonic}}$ is an unbiased estimate of β only if there are no selection effects on the "treatment" level based on the non-policy features X_i, but otherwise it is potentially biased. The double-debiased machine learning (DML) approach attempts to address this problem, by using the covariate vector X_i to learn a predictive model of the policy feature, and residualizing not only the dependent variable, $\log r_i$, but also the policy variable, POLICY_i (Chernozhukov et al., 2017). The resulting estimate, $\hat{\beta}_{\text{DML}}$, is our preferred estimate. Appendix B.2 provides an overview of DML.

5.2 Data

Our data consists of 21,262 New York City apartment listings across 13,243 buildings. We collected the data in February 2017 from StreetEasy, one of the leading online rental advertising platforms.[14] StreetEasy receives listings data directly from large and small brokers and rentals brokerage firms, individual owners, co-ops, and homeowner associations. As the success of the website's business model depends on providing renters with accurate information on all attributes relevant to their decision-making, StreetEasy ensures the accuracy of the listings' information by monitoring for and removing fraudulent listings, verifying the identity of brokers and brokerage/management agencies, and keeping the listings information up to date by frequently contacting agencies and owners.

Table 1 provides descriptions and statistics of key variables. We note that subletting-friendly policies are somewhat rare in our data set, with only around 1.1% of the listings explicitly allowing subletting.[15] Additional details on our data are provided in Appendix B.

[14] For more details, see http://www.streeteasy.com/.

[15] Though NYC law mandates that subletting cannot be unreasonably refused, tenants must obtain approval from the property owner or landlord before subleasing their apartments. In practice, landlords have several ways to make this process more or less costly for tenants, including affecting the speed and alacrity with which a security deposit is returned and repair requests are answered, whether a renewal offer is extended, and so on. Therefore, we interpret allowing sublets as the landlord signaling that she will not obstruct the process.

Table 1: Definitions and summary statistics of key variables

Variable	Definition	Mean	Median	Variance	Min	Max
price	Monthly rental price	3,760	2,800	4,161.3	750	10,000
sqft	Square footage	1,024	961	542.6	100	12,173
bd	Number of bedrooms	1.63	2	1.04	0	7
age	Building age (years)	76.39	72	38.87	0	120
a_sublets	Sublets allowed (1=Yes)	0.01	0	0.01	0	1
a_dogs	Dogs allowed (1=Yes)	0.18	0	0.14	0	1
a_washrdryr	Washer/Dryer in unit (1=Yes)	0.23	0	0.18	0	1

5.3 Effects of policies on long-term rental rates

We estimate the effects of two costly, building-level policies: whether the listing allows subletting, and whether the listing allows dogs. These policies are conceptually similar to home-sharing, and hence are likely to be useful proxies for our empirical exercise. For each policy, we report three estimates: (i) a "naive" estimate that simply regresses long-term rental rates on the policy treatment indicator, (ii) the hedonic pricing estimate described in Equation 7, and (iii) and the DML estimate. Figure 2 reports the results as percentage changes in the dependent variable, along with a 95% confidence interval around each point estimate.

Starting with subletting in Panel (a), we see that buildings allowing subletting have on average about 10% higher rental rates. A naive interpretation would be that building owners could increase profits by 10% simply by allowing subletting—assuming there are no real additional costs to this policy. However, as we discussed, insofar the subletting policy is correlated with building attributes that affect rents, this estimate is likely to be (severely) biased. Indeed, when we residualize the rental rates using the predictions of the hedonic model, we see that the "subletting premium" is likely entirely due to omitted variables bias: the effect of allowing subletting is slightly negative, and statistically indistinguishable from zero. The DML estimate is a precisely estimated zero.

With respect to allowing dogs, the naive estimate of Panel (b) suggests that building owners can may command a large rental premium by setting a dog-friendly policy. Similarly to subletting, this premium vanishes in both the hedonic pricing and the DML estimates. Together, the estimates from the two proxy policies support the prediction that owners cannot command higher rents by simply changing a costless policy. We interpret these results as evidence supporting the "no policy arbitrage" prediction of our model.

Our approach should predict that a premium is possible when "policies" are not costless—say, adding some amenity to a building. In Panel (c), we estimate the effects of whether the apartment has an in-apartment washer and dryer. This "policy" is costly for owners and clearly valued by at least some renters. Building confidence in our approach, the magnitude of the naive estimate is unreasonably large, and diminishes but remains positive for both the

Figure 2: Estimates of the effects policy decisions on long-term rental rates

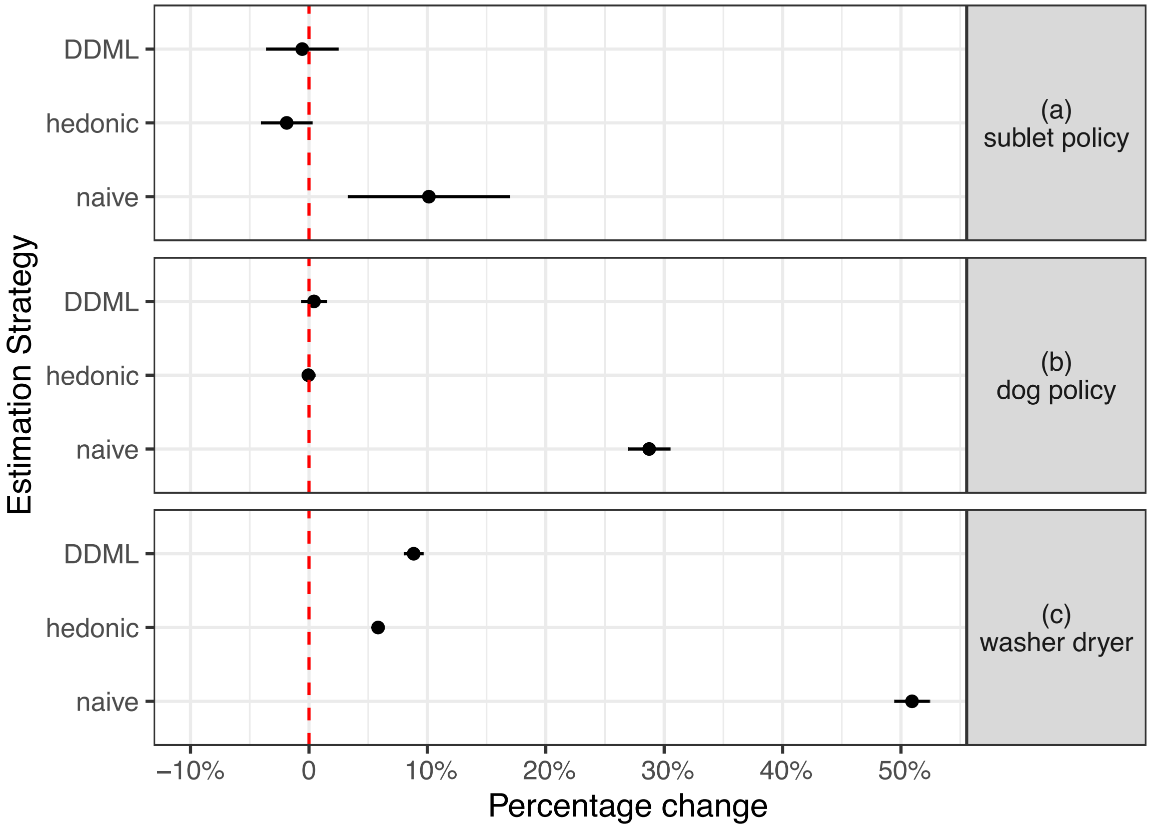

Notes: This figure reports estimates of the effects of policy decisions on long-term rental rates. The policies are (a) whether the listing allows subletting, (b) whether the listing allows dogs, (c) whether the listing has a washer and dryer. The "naive" estimate is calculated by regressing the long-term rental rates on the policy indicator variable. The hedonic estimates first residualize the long-term rental rates using a hedonic pricing model learnt from non-policy listing attributes. The DML estimates are obtained by first residualizing both the long-term rental rates, and the policy indicator variable. Each panel plots the estimated effect as the percentage change in the dependent variable. A 95% confidence interval is plotted around each estimate. Section 5.1 provides more details on the estimation strategy. Regression tables are presented in Appendix B.

hedonic and the DML estimates.

5.4 Limitations

Our empirical strategy relies on observational data, and we do not have access to a source of plausibly exogenous variation. In our context, however, we believe that this is less of a shortcoming that it might first appear. First, hedonic pricing models have been widely used in real estate markets, as these markets are characterized by high vertical differentiation, and the attributes consumers care about most—e.g., geographic location and size—are typically measured without error (Sirmans, Macpherson and Zietz, 2005). Second, our data comes from a thick and highly competitive rental market. Third, landlords and would-be tenants have strong incentives to reduce search costs by sharing as much match-relevant information about apartments as possible (e.g., number of bedrooms, location, building amenities). The same incentive is shared by the data broker through which we obtained data, as their business model solely relies on data accuracy.

6 Policy implications

We developed a model of a market for home-sharing and long-term rentals, derived the equilibria under policy regimes that differ only in which party the right to host is allocated to, and studied several extensions of the basic framework.

Our analysis of the TD regime reveals that a market where all tenants are allowed to home-share suffers from two fundamental problems. First, the amount of hosting is inefficient, in that there will be tenants whose hosting activity generates externality costs that outweigh the home-sharing benefits. Second, externalities are not internalized, and tenants not willing to participate in the home-sharing economy are always worse off compared to a market without the home-sharing option.

The BD regime fixes the problems of the TD regime: hosting activity is socially optimal, and hosts and non-hosts are sorted across buildings. The BD regime is also an "information-light" market-based policy response to home-sharing. As such, it does not require a central regulatory body with complete information about externality and hosting costs. Instead, these quantities are taken into consideration through the choices of the market participants who know them—this is a considerable advantage in this context, because these costs are highly idiosyncratic and hard to quantify. Market-based policies are also self-adjusting, and hence robust to structural shifts in market quantities; instead, centralized policy needs to be subject to periodical reevaluation to remain efficient.[16]

Today, city and state regulators take a wide array of approaches to home-sharing policies, while home-sharing platforms are lobbying for or against some of them, and propose their own policies. In the rest of this section, we examine some of these policies, and use our framework to study their efficiency and distributional consequences.

[16] Tirole (2015) also makes this case, providing historical and contemporary examples of this problem. In addition to home-sharing, the NYC taxicab medallion supply problem is a conceptually similar case in point, where a market inefficiency was created by supply failing to meet the growth in market demand due to regulatory restrictions (Tullock, 1975).

6.1 Centrally restricting supply

A city planner who wants to maximize tenant-side surplus has incentives to decrease supply below the social optimum (see Proposition 5). This distortion is in congruence with the findings of previous work on centrally restricting housing supply through regulation.[17] In the case of home-sharing, the exact supply restriction mechanism has important implications.

One way to restrict supply is by imposing direct costs on tenants, such as requiring them to obtain individual permits and licenses.[18] Higher hosting costs shift the supply curve $\hat{c}$ upwards, excluding some would-be hosts from home-sharing; as such, home-sharing profits accrue to fewer tenants. The resulting equilibrium is identical to that of the TD regime, and tenants do not sort according to their home-sharing preferences. These difficulties carry through to regulatory approaches such as increased home-sharing taxes, which reduce home-sharing activity by decreasing tenants' benefits from hosting. To bypass this problem, the city could instead allow building owners to set home-sharing policies, but make allowing home-sharing costly, e.g., by auctioning off a limited number of building-level licenses. In this case, tenant sorting would take place in equilibrium, but building owners in possession of home-sharing licenses would increase their long-term rental rates (see Section 6.4).

It is worth noting that, as we discussed in the beginning of this section, optimally reducing supply is hard regardless of the city planner's objectives. First, the city planner would need to have perfect information of all relevant market quantities, such as hosting and externality costs, which are highly idiosyncratic and hard to measure. Second, after setting an initial supply level, the city planner would have to engage in a potentially costly re-evaluation of these policy decisions in order to respond to shifts in market quantities. Furthermore, if externality costs are not homogeneous, any one tax level would be inefficient. In contrast, market-based solutions based on sorting of "types" retain their positive properties under heterogeneity assumptions (Tirole, 2015).

6.2 Hosting caps

Hosting caps limit the number of nights an apartment can be home-shared each calendar year. The economic rationale behind this increasingly popular policy is that it can render making properties exclusively available for home-sharing unprofitable for owners.[19] Applying

[17] For example, Glaeser, Gyourko and Saks (2005) examine the gap between building costs and market prices, and find that stricter zoning laws result in a 10-30% increase in house prices.

[18] For example, see https://www.engadget.com/2018-01-19-airbnb-san-francisco-listings-cut-in-half.html.

[19] A criticism of hosting caps is that they are often set lower than their "break-even" value—the value that would make owners indifferent between long-term rentals, and making properties available exclusively for home-sharing. Coles et al. (2019) estimate that the "break-even" hosting caps exceed 180 nights across all NYC boroughs in 2016. For example, home-sharing is limited to thirty days per year in Amsterdam, which is likely lower than the break-even value—see https://techcrunch.com/2018/01/10/amsterdam-to-halve-airbnb-style-tourist-rentals-to-30-nights-a-year-per-host.

a hosting cap reduces home-sharing supply from existing hosts, thereby increasing the going price for home-sharing, and inducing tenants with higher hosting costs to become hosts. Because new hosts have higher hosting costs, hosting caps result in lower market supply. Unlike the supply restriction mechanism that we examined in Section 6.1, hosting caps expand the number of tenants to which home-sharing benefits accrue. However, market failure may still occur in equilibrium, and non-hosts incur externality costs.

To bypass these problems, hosting caps can be applied concurrently with the BD policy regime. Because more tenants want to become hosts, more owners will allow home-sharing in equilibrium (see Propostion 2). As the new hosts are characterized by higher hosting costs—otherwise they would have already been hosting—hosting caps shift the supply curve upwards, and decrease social surplus. One exception can be found with the case of convex externality costs, where hosting caps may instead increase the efficiency of the BD regime. In this case, the BD regime is too conservative, whereas the socially optimal solution spreads hosting activity equally amongst buildings, but does not allow for tenant sorting (see Section 4.1). Applying hosting caps in conjunction with the BD regime spreads out hosting activity across more buildings, pushing it closer to the socially optimal solution. Whether hosting caps increase or decrease surplus ultimately depends on the elasticities of the market demand and supply, as well as on the degree of convexity of the externality cost function.

6.3 The role of the platform

Online platforms reduce search and transaction costs by aggregating supply and demand, maintaining reputation systems, offering transaction insurances, and automating large parts of each transaction. In the context of home-sharing, this reduction in transaction costs can be thought of as a reduction in hosts' hosting costs, and has been the main contributor of the rapid proliferation of home-sharing (Filippas et al., 2020).

Platforms could increase the social surplus that home-sharing generates by reducing its negative externalities. In terms of our model's parameters, reducing externality costs implies pushing the total cost $\hat{c}_t$ closer to the hosting cost $\hat{c}$, which results in higher equilibrium hosting quantities across *all* regimes we examined, and higher surplus for both sides of the market. As home-sharing platforms typically leverage a fixed percentage fee on transactions, they have strong incentives to reduce externality costs.

Home-sharing platforms already take several steps towards reducing externality costs. Part of the effort centers on informing hosts and guests about the specifics of each building, neighborhood, and city, such as noise ordinance laws and expected behavior, and providing insurance to both hosts and building owners for misuse and damages. Another interesting measure is Airbnb's provision of a platform for neighbors of hosts to complain about cases where guests generated extensive negative externalities, such as noise issues or misuse of common spaces.[20] Home-sharing platforms also maintain reputation systems in order to enforce better behavior, and to remove bad actors from the market. However, the effectiveness of such mechanisms erodes over time (**?**).

The second important dimension of the problem is whether the externality costs of home-sharing are internalized. We have stressed throughout the paper that this property is obtained *only* in the presence of building-wide policies, as the externalities of home-sharing are internalized only if hosts and non-hosts are sorted. As such, we expect home-sharing platforms to encourage a move in this direction. Interestingly, Airbnb has initiated a "friendly buildings" program, coinciding with the prescription of our paper.[21]

Incumbents often employ lobbying in an attempt to pose regulatory barriers to the entry and growth of technology firms (Djankov, La Porta, Lopez-de Silanes and Shleifer, 2002; Cusumano, 2015). As a response, sharing economy platforms have recently intensified their lobbying efforts.[22] In the context of home-sharing, our model shows that these lobbying efforts should be directed towards state rather than city regulators, as city regulators have incentives to reduce supply, at the cost of restricting the growth of home-sharing platforms. Airbnb's recent lobbying efforts have been in accordance with this finding.[23]

6.4 Policy arbitrage

A prediction of our model is that there is "no policy arbitrage," that is, long-term rental rates are equal among buildings with different home-sharing policies in the BD equilibrium (see Section 3.4). This prediction stems from the assumption that home-sharing imposes no cost on building owners, and that buildings are identical. Relaxing these assumptions changes the predictions of our model.

If some buildings have features are more attractive to hosts than to non-hosts, equilibrium rents will not be equal. For example, suppose that a building a is equipped with keyless unlocking technology which decreases the hosting costs of each host i to $c_i' = c_i - k$. Following the steps of Proposition 2, we can show that it is optimal for owner a to allow home-sharing, as she can command a rental rate $r_1 = r_0 + k$ in the BD equilibrium.

Another case where the "no policy arbitrage" prediction may not hold is if allowing home-sharing is costly for building owners. Let c_H denote the owner's cost for allowing home-sharing, such as . For example, the cost c_H may stem from guests inflicting additional wear-and-tear on apartments and common building resources, or the administrative costs associated with home-sharing. Following the steps of Proposition 2, it is easy to show that

[20] See also https://www.airbnb.com/neighbors.

[21] See also https://www.airbnb.com/help/article/1195/what-s-the-airbnb-friendly-buildings-program.

[22] For example, see https://qz.com/630939/charted-which-tech-companies-spend-millions-in-lobbying-the-us-government.

[23] For example, see https://www.theinformation.com/articles/uber-airbnb-fight-cities-by-lobbying-states.

for owners to be indifferent between allowing and prohibiting home-sharing, $r_1 = r_0 + c_H$ in the BD equilibrium. In words, rental rates between buildings with different home-sharing policies are no longer equal, but owners do not profit from their choice of home-sharing policy. In real life, the owners costs for allowing home-sharing are likely slight: hosts have the incentive to keep their apartments in good condition to make them attractive to guests, and home-sharing platforms typically offer insurance against guest damage to the building. We empirically assess the "no policy arbitrage prediction" in Section 5.

7 Conclusion

Our model suggests allowing individual building owners discretion in setting home-sharing policies. Under this policy regime, hosts and non-hosts can sort across buildings with the preferred home-sharing policy, and the social welfare obtained coincides with that of a market regulated by a social planner. The reason is that terms between different types of buildings are equalized in a competitive long-term rental market, and the marginal host's individual benefit does not exceed the full cost of him living in such a building. The two alternatives we examined—allocating decision rights to the individual tenant or to the city—are likely to lead to too much, and too little, hosting, respectively.

Our empirical analysis of the NYC rental market strongly suggests that, as predicted by our model, building owners cannot extract a premium through policy decisions that are costless to them, but that potentially imply negative externalities for other tenants. Employing an agent-based modeling approach, we exhibit that a market under the building-specific policies regime always converges to equilibrium. Higher moving costs reduce tenant surplus, while within-building tenant type correlation decreases the amount of moving necessary for the equilibrium to be reached.

As technological innovations continue to bring forth applications previously thought not possible, policy-makers will debate about policy that addresses externality issues, and managers will strive to aid this effort proactively or face significant regulatory pushback. Our paper adds rigor to the policy debate about home-sharing, introduces a theoretical framework that can generally be applied to externalities caused by online platforms, and offers clear prescriptions for policy makers and platform managers.

A natural direction for future work would be to empirically investigate aspects of the model. For example, it might be illuminating to interview building owners making decisions, and examine how they are dealing with existing and prospective tenants. Another direction is to test whether cities with particularly inelastic travel demand—and hence the ability to extract substantial rents—are also the cities most interested in restricting home-sharing.

References

Alyakoob, Mohammed and Mohammad Saifur Rahman, "Shared Prosperity (or Lack Thereof) in the Sharing Economy," *Information Systems Research*, 2022.

Arthur, W Brian, "Complexity and the economy," *Science*, 1999, *284* (5411), 107–109.

_ , "Out-of-equilibrium economics and agent-based modeling," *Handbook of Computational Economics*, 2006, *2*, 1551–1564.

Athey, Susan, "The impact of machine learning on economics," in "The economics of artificial intelligence: An agenda," University of Chicago Press, 2018, pp. 507–547.

Barron, Kyle, Edward Kung, and Davide Proserpio, "The Sharing Economy and Housing Affordability: Evidence from Airbnb," *Working paper*, 2017.

Bonabeau, Eric, "Agent-based modeling: Methods and techniques for simulating human systems," *Proceedings of the National Academy of Sciences*, 2002, *99* (suppl 3), 7280–7287.

Brynjolfsson, Erik, Yu Hu, and Duncan Simester, "Goodbye pareto principle, hello long tail: The effect of search costs on the concentration of product sales," *Management Science*, 2011, *57* (8), 1373–1386.

Chan, Jason and Anindya Ghose, "Internet's dirty secret: Assessing the impact of online intermediaries on HIV transmission," *MIS Quarterly*, 2014, *38* (4), 955–976.

Chang, Ray M, Wonseok Oh, Alain Pinsonneault, and Dowan Kwon, "A network perspective of digital competition in online advertising industries: A simulation-based approach," *Information Systems Research*, 2010, *21* (3), 571–593.

Chen, Tianqi, Tong He, Michael Benesty et al., "Xgboost: extreme gradient boosting," *R package version 0.4-2*, 2015, pp. 1–4.

Chernozhukov, Victor, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins, "Double/debiased machine learning for treatment and causal parameters," *The Econometrics Journal*, 2017.

Clewlow, Regina R and Gouri Shankar Mishra, "Disruptive transportation: the adoption, utilization, and impacts of ride-hailing in the United States," Technical Report, Research Report–UCD-ITS-RR-17 2017.

Coase, Ronald, "The problem of social cost," *Journal of Law & Economics*, 1960, *3*, 1–44.

Coles, Peter A, Michael Egesdal, Ingrid Gould Ellen, Xiaodi Li, and Arun Sundararajan, "Airbnb Usage Across New York City Neighborhoods: Geographic Patterns and Regulatory Implications," *Cambridge Handbook on the Law of the Sharing Economy*, 2019.

Cusumano, Michael A, "How traditional firms must compete in the sharing economy," *Communications of the ACM*, 2015, *58* (1), 32–34.

Daskalakis, Constantinos, Paul W Goldberg, and Christos H Papadimitriou, "The complexity of computing a Nash equilibrium," *SIAM Journal on Computing*, 2009, *39* (1), 195–259.

Dinerstein, Michael, Liran Einav, Jonathan Levin, and Neel Sundaresan, "Consumer price search and platform design in internet commerce," *American Economic Review*, 2018, *108* (7), 1820–59.

Djankov, Simeon, Rafael La Porta, Florencio Lopez de Silanes, and Andrei Shleifer, "The regulation of entry," *The quarterly Journal of economics*, 2002, *117* (1), 1–37.

Dube, Arindrajit, Jeff Jacobs, Suresh Naidu, and Siddharth Suri, "Monopsony in online labor markets," *American Economic Review: Insights*, 2020, *2* (1), 33–46.

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

Farronato, Chiara and Andrey Fradkin, "The welfare effects of peer entry: the case of Airbnb and the accommodation industry," *American Economic Review*, 2022, *112* (6), 1782–1817.

Filippas, Apostolos, John J Horton, and Joseph Golden, "Reputation inflation," *Marketing Science*, 2022.

_ , _ , and Richard J Zeckhauser, "Owning, using and renting: Some simple economics of the "sharing economy"," *Management Science*, 2020.

Galla, Tobias and J Doyne Farmer, "Complex dynamics in learning complicated games," *Proceedings of the National Academy of Sciences*, 2013, *110* (4), 1232–1236.

Glaeser, Edward L, Joseph Gyourko, and Raven Saks, "Why is Manhattan so expensive? Regulation and the rise in housing prices," *The Journal of Law and Economics*, 2005, *48* (2), 331–369.

Greenwood, Brad N and Sunil Wattal, "Show me the way to go home: An empirical investigation of ride-sharing and alcohol-related motor vehicle fatalities," *MIS Quarterly*, 2017, *41* (1).

Guttentag, Daniel, "Airbnb: Disruptive innovation and the rise of an informal tourism accommodation sector," *Current Issues in Tourism*, 2015, *18* (12), 1192–1217.

Horn, Keren and Mark Merante, "Is home sharing driving up rents? Evidence from Airbnb in Boston," *Journal of Housing Economics*, 2017, *38*, 14–24.

Jackson, Joshua Conrad, David G Rand, Kevin Lewis, Michael I Norton, and Kurt Gray, "Agent based modeling: A guide for social psychologists," *Working Paper*, 2016.

Jefferson-Jones, Jamila, "Airbnb and the housing segment of the modern sharing economy: Are short-term rental restrictions an unconstitutional taking," *Hastings Constitutional Law Quarterly*, 2014, *42*, 557.

Kaplan, Roberta A and Michael L Nadler, "Airbnb: A case study in occupancy regulation and taxation," *University of Chicago Law Review Dialogue*, 2015, *82*, 103.

Kroft, Kory and Devin G Pope, "Does online search crowd out traditional search and improve matching efficiency? Evidence from Craigslist," *Journal of Labor Economics*, 2014, *32* (2), 259–303.

Laffont, J-J, "Externalities," in "Allocation, Information and Markets," Springer, 1989, pp. 112–116.

Laffont, Jean-Jacques, "The New Palgrave Dictionary of Economics," *Steven N. Durlauf and Lawrence E. Blume Eds*, 2008.

Malhotra, Arvind and Marshall Van Alstyne, "The dark side of the sharing economy . . . and how to lighten it," *Communications of the ACM*, 2014, *57* (11), 24–27.

Marcet, Albert and Juan P Nicolini, "Recurrent hyperinflations and learning," *The American Economic Review*, 2003, *93* (5), 1476–1498.

Molnar, Alejandro and Daniel Mangrum, "The marginal congestion of a taxi in New York City," *Working paper*, 2018.

Nichols, Albert L and Richard J Zeckhauser, "Targeting transfers through restrictions on recipients," *The American Economic Review*, 1982, *72* (2), 372–377.

Oh, Wonseok, Jae Yun Moon, Jungpul Hahn, and Kim Taekyung, "Research Note: Leader influence on sustained participation in online collaborative work communities: A simulation-based approach," *Information Systems Research*, 2016, *27* (2), 383–402.

Polinsky, A Mitchell and Steven Shavell, "Pigouvian taxation with administrative costs," *Journal of Public Economics*, 1982, *19* (3), 385–394.

Rahmandad, Hazhir and John Sterman, "Heterogeneity and network structure in the dynamics of diffusion: Comparing agent-based and differential equation models," *Management Science*, 2008, *54* (5), 998–1014.

Schelling, Thomas C, "Dynamic models of segregation," *Journal of Mathematical Sociology*, 1971, *1* (2), 143–186.

Scitovsky, Tibor, "Two concepts of external economies," *The Journal of Political Economy*, 1954, *62* (2), 143–151.

Seamans, Robert and Feng Zhu, "Responses to entry in multi-sided markets: The impact of Craigslist on local newspapers," *Management Science*, 2013, *60* (2), 476–493.

Sheppard, Stephen and Andrew Udell, "Do Airbnb properties affect house prices?," *Working Paper*, 2018.

Sirmans, Stacy, David Macpherson, and Emily Zietz, "The composition of hedonic pricing models," *Journal of real estate literature*, 2005, *13* (1), 1–44.

Sundararajan, Arun, *The sharing economy: The end of employment and the rise of crowd-based capitalism*, MIT Press, 2016.

Tebbens, Radboud J Duintjer and Kimberly M Thompson, "Priority shifting and the dynamics of managing eradicable infectious diseases," *Management Science*, 2009, *55* (4), 650–663.

Tesfatsion, Leigh and Kenneth L Judd, *Handbook of Computational Economics: Agent-based Computational Economics*, Vol. 2, Elsevier, 2006.

Tiebout, Charles M, "A pure theory of local expenditures," *The Journal of Political Economy*, 1956, *64* (5).

Tirole, Jean, "Market failures and public policy," *American Economic Review*, 2015, *105* (6), 1665–82.

Tullock, Gordon, "The transitional gains trap," *The Bell Journal of Economics*, 1975, pp. 671–678.

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

A Reaching equilibrium

Even though the regime wherein owners decide on their building's home-sharing policy is socially optimal, convergence to the market equilibrium would require tenants to "sort" into buildings of the appropriate policy, thereby creating two potential problems. First, individually rational behavior is not guaranteed to converge to a steady market state, or may require a prohibitively large amount of time to do so. The resulting fluctuations in prices as well as changes in other market quantities could require substantial tenant sorting to "fix." A long line of game-theoretic research shows that systems comprising individually rational decision makers are not guaranteed to self-stabilize. For example, agents, often modeled as best-responding to the current system state, may get trapped in cycles of suboptimal states, and the market may either fail to reach equilibrium or require a prohibitively large amount of time (Arthur, 1999; Marcet and Nicolini, 2003; Arthur, 2006; Daskalakis, Goldberg and Papadimitriou, 2009; Galla and Farmer, 2013). Furthermore, tenant "types"—those that want to host and those that do not—are initially mixed across buildings. Any policy imposed by a landlord will leave some of them happy and others unhappy. Tenants who are dissatisfied will subsequently look to move to a building with their preferred home-sharing policy. However, to do so they would have to incur costs such as time spent in searching and evaluating, realtor fees, moving expenses, and so on. The sorting mechanism is costly, and these costs could dissipate the surplus of home-sharing. As a consequence, some tenants may get "locked into" their current building, and the market may fail to reach the state that the BD equilibrium predicts.

These two issues—(1) can the equilibrium be reached and (2) what are the implications of adjustment costs—may raise questions about the applicability of the building-specific policy approach to real-life markets. To explore the tâtonnement process by which an equilibrium is obtained, we construct an agent-based model of the home-sharing rentals market. Agent-based models (ABMs) are computational simulations in which entities are programmed to interact and respond to their environment over time (Jackson, Rand, Lewis, Norton and Gray, 2016). ABMs are commonly used to study emergent and transitory macro-level phenomena created by micro-level behavior, which would otherwise be theoretically intractable (Schelling, 1971; Bonabeau, 2002; Tesfatsion and Judd, 2006; Rahmandad and Sterman, 2008; Tebbens and Thompson, 2009; Chang, Oh, Pinsonneault and Kwon, 2010; Oh, Moon, Hahn and Taekyung, 2016).

We first show that the market operating under the building-specific policy regime converges to the competitive equilibrium under a variety of initial conditions. We then incorporate moving costs to the model and find that a 1% increase in moving costs results in roughly a 1% decrease in the tenant surplus generated through home-sharing, compared to the case where moving costs are zero. While the home-sharing equilibrium supply only marginally decreases with higher moving costs, some tenants are "locked into" buildings with undesirable (for them) home-sharing policies. As a result, tenants with higher hosting costs end up becoming home-share hosts, and tenants with lower hosting costs are excluded from home-sharing, creating an inefficiency. Nevertheless, the net effect of home-sharing on tenant surplus is always positive. It is also worth noting that the moving expense is likely a one-time cost, as we find that in almost all cases tenants will select into buildings of the right "type." Finally, we show that including within-building correlation in tenant types—captured through correlated hosting costs for tenants residing in the same building—leads to faster convergence, as well as to a decrease in the number of tenant moves necessary for the market equilibrium to be reached.

A.1 An agent-based model of the BD regime

We build our ABM analogously to the model of Section 3. We begin our description by focusing on tenants. At time t, tenant $i \in I$ lives in a building $b_i(t) = j \in J$, and can home-share only if the policy of the building allows for hosting. If he is allowed, tenant i hosts if the market price for home-sharing, $p(t)$, exceeds his personal hosting cost, c_i. If $k_j(t)$ other hosts live in the same building, then tenant i incurs total externality costs $k_j(t)c_E$. Buildings that allow for home-sharing charge rent $r_1(t)$, and buildings that prohibit home-sharing charge rent $r_0(t)$.

When tenants would be better off living in another building, they enter a pool of tenants who want to move from their apartments. To move, tenants incur a cost $c_{i,m}$. There are two cases in which tenants move. First, tenants want to move if they are currently not allowed to host and hosting would increase their utility. In the language of our model, tenant i wants to move if there exists some building j' such that

$$u_0 - r_0(t) \leq u_0 - c_{i,m} - r_1(t) + p(t) - c_i - k_{j'}(t)c_E.$$

Second, tenants want to move if they are currently allowed to host, but would be better off in a building that prohibits home-sharing as they would not have to incur the externalities from other tenants' hosting activity. Formally, tenant i wants to move if there exists some building j' such that

$$u_0 - r_0(t) - c_{i,m} \geq u_0 - r_1(t) - k_j(t)c_E + \max\{0, p(t) - c_i\}.$$

We assume that tenants only consider their present utility from living in a home-sharing friendly apartment against not being able to host, i.e., they do not form expectations about others' behavior, they are "small" relative to the market. The reason for this assumption is that the agents' decision process is in practice stationary: in our simulations we find that tenants (almost) never move buildings twice, and owners (almost) never change their building's policy more than once: agents, both owners and tenants, spend the rest of their time in the state they move to.

Market clearing is brought about through both rent and home-sharing policy adjustments. Building owners adjust rents and home-sharing policies in response to the relative demand for moving to home-sharing friendly and unfriendly buildings. For example, if there are more tenants looking to move to buildings that allow for home-sharing than to those which prohibit it, then rents in the former buildings increase in the next period, while the latter may convert to a home-sharing-friendly policy. It is worth mentioning here that tâtonnement requires both rent and policy adjustments. While the theoretical model we developed predicts "no policy arbitrage" in equilibrium, i.e., that rents are equalized in across building "types," we do not disallow rent adjustments in the ABM, as we want to examine whether this property is an "organic" market outcome in our simulations. Similarly, assuming that home-sharing policy adjustments do not take place would impose a constant supply constraint on the building owner side.

As moving decisions usually take place on a yearly basis, each period in our ABM can be thought of as a year in a real-life rental market. Each instance of our computational model is carried out for 50 periods, or until the market reaches a steady state. Initial building policies are randomly selected with equal probability; other methods of initialization that we tried do not qualitatively change our results.

We describe the order in which events take place in every period below.

1. **Pool of movers is identified.** Tenants who are dissatisfied with their building's current home-sharing policy and who would be better off incurring the cost of moving to another building enter the pool of potential movers to and away from home-sharing -friendly apartments, creating market demand for the corresponding building "type."

2. **Building-specific policies are adjusted.** Building policies respond to the market demand. For example, if more tenants want to move to home-sharing-friendly buildings, then the home-sharing-unfriendly buildings probabilistically change their policies to cover, in expectation, a percentage of the excess demand. The exact percentage is a parameter of the ABM, and our results are qualitatively insensitive to whether too few or too many buildings change their policies to cover the excess demand. If there is no net difference in demand, policies remain unaffected.

3. **Rents are adjusted**. Rents also respond to the aggregate demand. Buildings with policies for which there is higher demand increase their rental prices by a constant amount, while rents in the other category remain unchanged. Similarly to policies, if the two type of demands are equal, there is no change in rents.

4. **Tenants move**. After rents and building policies are adjusted, tenants determine whether they want to change buildings. A tenant attempts to move if the difference in utility obtained by changing apartments is higher than his moving cost. If the sets of tenants that want to move to buildings with different policies are both non-empty, we randomly select pairs of tenants and switch the building in which they reside. In the case where the demand to move to one type of building exceeds the other, some tenants will not be able to move.

5. **Market quantities are updated.** The tenants update their hosting decisions. The price of home-sharing rentals, modeled as a decreasing linear function of supply, responds to the new market state.

These five steps constitute a period in our model, and are repeated until the system converges to the computational equilibrium, or until fifty periods have passed. The computational equilibrium is defined as the state in which no tenant wants to switch buildings, and therefore no owner wants to change the building's home-sharing policy or increase rents. If the upper bound on the number periods is exceeded, then we say that the market fails to reach an equilibrium.

A.2 Example simulations

To illustrate how our computational model works, we provide the results of a set of example simulations. Figure 3 depicts the time series of the fraction of home-sharing-friendly buildings, the fraction of tenants that are dissatisfied and want to move, and the percentage difference in rents of the two types of buildings until convergence is achieved. Each simulation is represented by a separate line.

For the purposes of our simulation, we consider an ABM with $3,000$ tenants (agents) living in 30 buildings of capacity 100 each. The hosting cost of each tenant is determined through identical and independent draws from a uniform distribution with positive range. As a result, the supply curve is approximately linear and upward sloping. To start, tenants do not incur a moving cost to move apartments. Initial building home-sharing policies are randomly determined. These two factors add stochasticity in our model and hence result in different paths for each simulation. The demand curve for home-sharing is linear and downward sloping. Note that other configurations that we tested did not change the significance or the direction of the results. We use the same simulation parameters in the rest of this section unless otherwise noted.

As expected, the entry of the home-sharing option and the subsequent owners' decisions on building-specific home-sharing policies initially leave some tenants unhappy. Most of the tenant sorting occurs early on in the process, and the number of dissatisfied tenants rapidly drops, with less than 5% being dissatisfied after the second period. The process converges to a state where there is a negligible amount of tenants that are dissatisfied (less than 3%). Note that the number of unhappy tenants is not driven to zero since our computational model is discrete, and the optimal solution need not have an integer number of buildings allowing for home-sharing. Similarly, the number of home-sharing-friendly buildings initially varies but soon converges to one of two values, again due to the discrete nature of our model. Finally, the rent equalization property of the BD policy regime is also satisfied in the example simulations, with equilibrium rents being approximately equal—disparities are again small, and due to the discrete nature of the ABM.

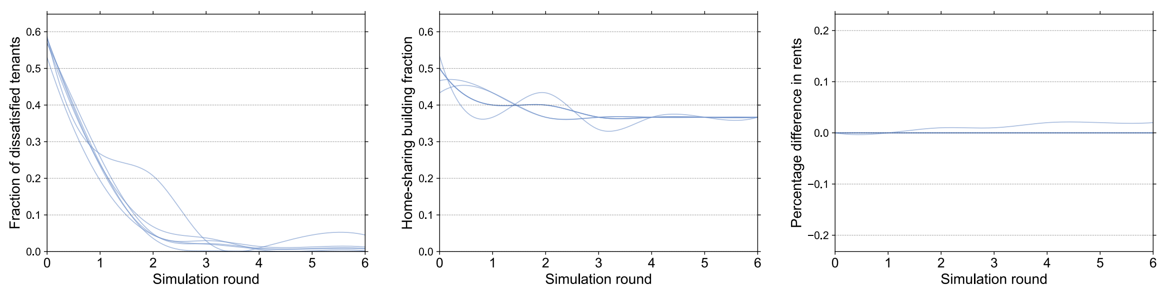

Figure 3: This figure plots the results for a set of example simulations of our agent-based model. Each line indicates a different instance of the ABM. For every round of the simulation, the leftmost panel plots the fraction of tenants that want to move to a building of different type, the middle panel plots the fraction of buildings that allow for home-sharing, and the rightmost panel plots the percentage different of long-term rental rates between buildings of the two types.

A.3 Convergence to the BD equilibrium

Our first question pertains to the time of convergence to equilibrium for a market operating under the BD regime. As we discussed in the beginning of Appendix A, collective behavior of individually rational agents is not guaranteed to result in convergence to equilibrium. In the case of home-sharing, this failure to converge is consequential, as the market may not obtain the positive properties of the BD regime, or it may require a prohibitively long time to obtain them.

To estimate whether the market operating under the BD regime robustly reaches the equilibrium state, our approach is to run a large number of instances of the ABM model starting from different initial conditions. We conduct $20,000$ iterations with parameters chosen as described in Section A.2. The upper bound for convergence is set to $T = 100$ periods; if the market does not reach equilibrium until time T, then we assume that it has failed to converge.

Our results are reported in Figure 4. Convergence times appear to be following a truncated normal distribution. Importantly, we do not find any case where the market does not reach equilibrium. The error bar of the number of tenants who want to move as a function of time is presented as a ribbon on the mean, and shows that the number of dissatisfied tenants quickly drops to near-zero values. Furthermore, the equilibrium number of tenants home-sharing is on average within 0.1% of the BD equilibrium quantity (standard deviation=0.005). Accounting for the discrete nature of our computational model, the results of our experiment indicate that the market both reaches equilibrium within a reasonable time limit, and that this equilibrium always coincides with the theoretical prediction for the BD regime.

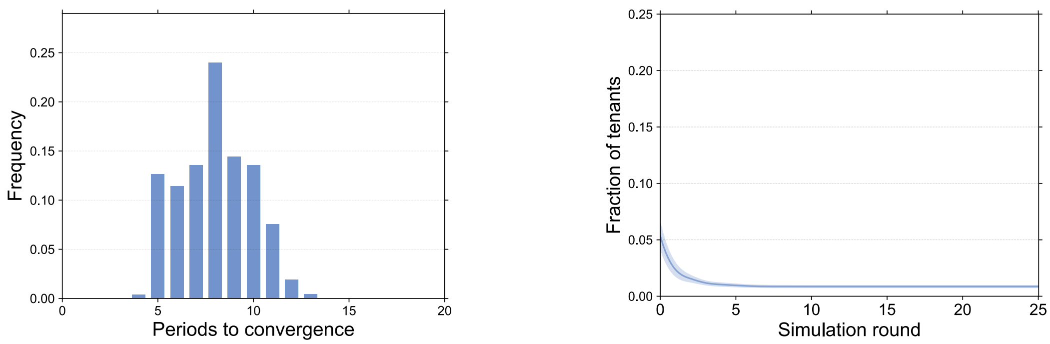

Figure 4: Distribution of equilibrium convergence times (left) and fraction of tenants who want to move as a function of time (right).

A.4 Moving costs

An important factor that is not captured by our theoretical analysis are the costs associated with moving: tenants who are dissatisfied with their building's home-sharing policy have to incur substantial costs to move to a building of the appropriate "type." As a result, some tenants may elect to stay in their current building even if they would be better off elsewhere; the market is then pushed to a sub-optimal state with a number of home-sharing rentals different than what is predicted by the BD equilibrium without moving costs.

To assess the impact of the costs of moving on the market operating under the BD regime, we employ our computational model and carry out simulations while varying moving costs. Moving costs are set equal to 10% of the annual rent, with the results remaining qualitatively similar for different values that we tried. Figure 5 reports error bars depicting the (normalized) mean tenant surplus and the average fraction of tenants that host in home-sharing-friendly buildings as a function of moving costs, reported as the ratio with respect to the annual rent. We notice a considerable decrease in tenant surplus. However, we also observe that almost every tenant in the home-sharing-friendly buildings hosts for even large values of moving costs, but the percentage starts decreasing as moving costs become very large; this indicates that some tenants are dissatisfied but cannot change buildings.

To examine the underlying effects further, we report in Figure 6 the percentage change effect of costs on the amount of sorting required, the home-sharing market supply and the tenant surplus. Home-sharing market supply is barely affected, and is equal to the BD equilibrium value for a wide range of tenant costs. However, both the tenant surplus and the sorting required for convergence to equilibrium decrease as moving costs increase. This implies that, while the home-sharing supply remains efficient, tenants with high hosting costs are "locked into" home-sharing-friendly buildings; these tenants see their utility decrease but cannot move. Among them, those tenants for whom the individual rationality condition is satisfied will list their apartments, although the internalization condition (Equation 1) does not hold. As a result, market price decreases, and tenants with lower hosting costs are no longer willing to incur the cost to move to home-sharing-friendly buildings. This effect is welfare-reducing, with a 10% increase in moving costs resulting in an average of 10% decrease in tenant surplus on a yearly basis.

It is important to note that the discount rate of tenants and the amount of "organic" moving that occurs can matter in the interpretation of the results. If tenants have a low discount rate, moving costs would become less important relative to the long-term benefits of being in the "right" building. Similarly, if tenants move frequently anyway, the cost of being in the "wrong" building can be fairly small, especially with a high discount rate. We view the simulation of market adjustment with moving costs as an illustration of the mechanisms by which welfare-relevant outcomes arise.

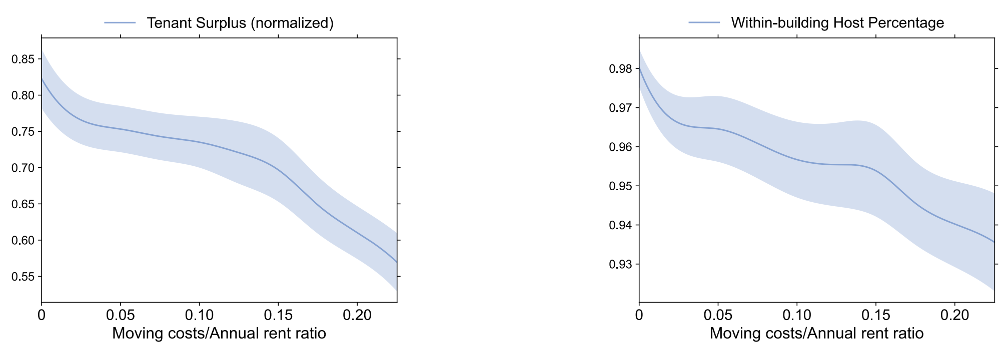

Figure 5: Tenant Surplus (left) and percentage of tenants living in home-sharing-friendly building that host (right) as a function of moving costs.

A.5 Correlation in tenant types

An additional concern with the BD equilibrium is the amount of sorting that needs to take place before the equilibrium is reached. However, the amount of moving required for that to happen can in fact be less than one might initially think: rather than full mixing, it seems likely that in practice similar tenants live in the same building, hence tenant "types" are

Figure 6: Percentage changes with respect to the zero moving cost case.

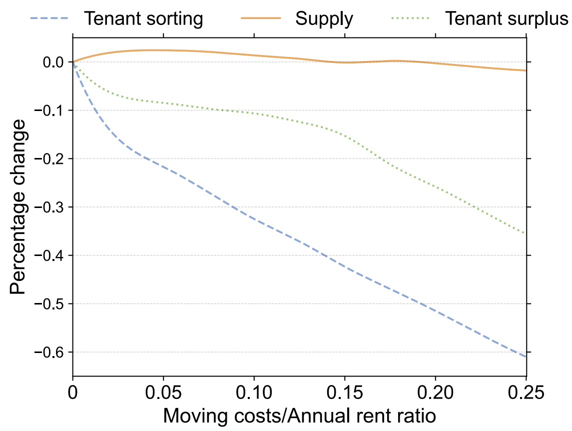

correlated within buildings. In our model, this intuition translates into tenants with high hosting costs (e.g. high opportunity cost, wealthier individuals) to be more likely to reside in some buildings at the time of the introduction of home-sharing in the market, while tenants with lower hosting costs in others. Since tenants are already "sorted," we would expect that the sorting necessary for the process to reach equilibrium may be less than if tenants were fully mixed.

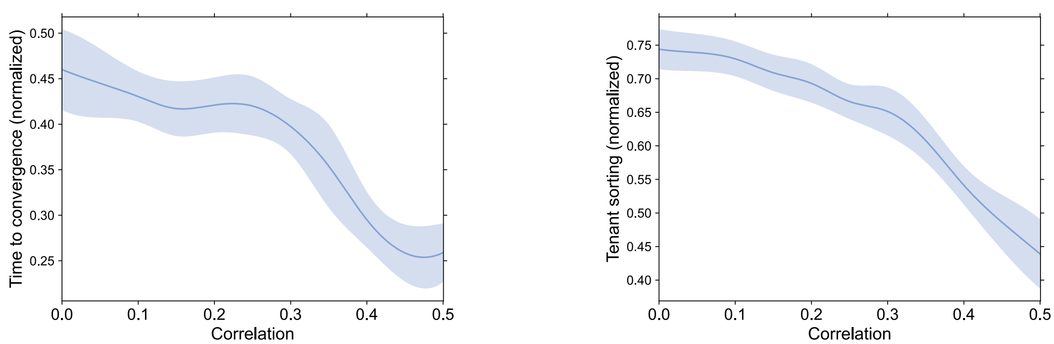

Figure 7: Time to convergence (left) and tenant sorting required to converge (right) as a function of within-building hosting cost correlation.

We incorporate the above intuition in our computation model by adding within-building hosting cost correlations. The hosting costs of tenants within a building can be independently drawn (corr=0) or completely correlated (corr=1). The results of our experiments are shown in Figure 7, and the percentage change effects are reported in Figure 8. Initially, correlation has a small but negative effect on both tenant sorting and time to convergence. As the value of correlation further increases, we observe a large reduction in both quantities, with a 10% increase in correlation resulting in an average of 14% decrease in tenant sorting and a 10%

Figure 8: Percentage changes with respect to the zero correlation case.

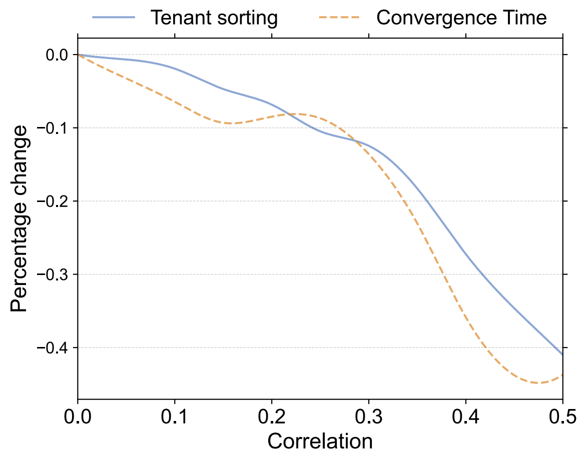

decrease in convergence time.

B Details on the empirical analysis

B.1 Data

Our data set consists of 21,262 New York city apartment listings across 13,243 buildings collected in February 2017 from StreetEasy, one of the leading online rental advertising platforms. We collected the data over a one-week period in by building a crawler in Python. Our data set contains information on every NYC rental listing on the web page during that period, as well as all information contained in each listing's respective building page. Our data collection was completed successfully, in that the entire collection of rentals was parsed—we cross-referenced a large number of samples with the available listings on the website at that time, and found no discrepancies.

For each listing, we have access to 87 attributes, including attributes of the listing's building (e.g., building age, doorman service, concierge service, and so on), as well ass geographical information (e.g., zip code and borough information, latitude and longitude, whether the building is on the waterfront, and so on). Table 2 provides summary statistics for several data attributes., and Figure 9 depicts for listings in our data.

Figure 9: Heatmap of the spacial density of rentals

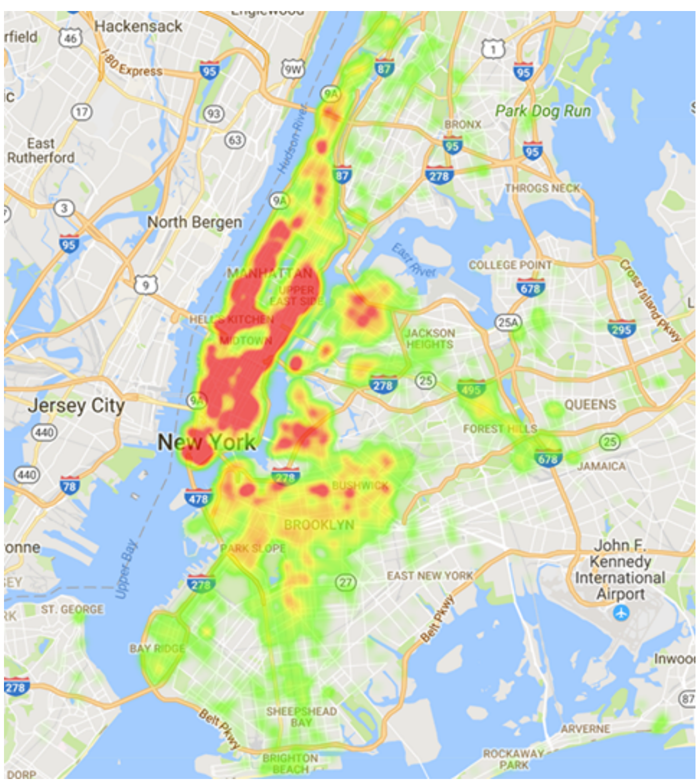

Notes: This figure plots a heatmap indicating the density of apartment listings in our data by geographical location. Redder hues indicate a higher number of rentals in that area, and greener hues a lower number.

Table 2: Definitions and summary statistics of data attributes

Variable	Definition	Mean
age	building age	76.39
bd	# of bedrooms	1.63
price	monthly rental price	3,3757
stories	# of stories	9.87
sqft	square footage	1,023.9
video	# of videos	0.02
a_ac	has central a/c	0.13
a_balcony	has balcony	0.08
a_bike	has bike room	0.20
a_broker	has broker fee	0.55
a_court	has courtyard	0.07
a_dogs	dogs allowed	0.18
a_elevator	has elevator	0.43
a_fios	fios available	0.34
a_fireplace	has fireplace	0.05
a_furnished	is furnished	0.05
a_garage	has garage	0.14
a_garden	has garden	0.12
a_loft	is a loft	0.03
a_park	park nearby	0.12
a_patio	has patio	0.06
a_storage	has storage room	0.02
a_sublets	sublets allowed	0.01
a_supper	has live-in supper	0.39
a_valet	valet service	0.02
a_watrfrnt	is in the waterfront	0.02
a_wshrdryr	has washer/dryer	0.23

Notes: This table reports summary statistics of features used in our empirical analysis. For binary features, we report only the mean, which is the percentage of apartments with that feature. Omitted features include *board*: board approval required, *borough*: the borough of the apartment, *children*: has children's playroom, *coldStorage*: has cold storage unit, *fullTimeDoorman*: has full-time doorman, *guarantor*: guarantors accepted, *gym*: has gym, *hood*: the neighborhood of the apartment, *hotTub*: has hot tub, *lng*: the longitude of the building, *lat*: the latitude of the building, *leed*: is LEED certified, *mediaRoom*: has a media room, *nycE-vacuation*: the NYC evacuation code of the building, *oceanFront*: is at the ocean front, *recordedSales*: how many apartments have been previously sold in this building, *recreationFacilities*: has recreation facilities, *roofDeck*: has access to roof-deck, *storage*: has storage room, *packageRoom*: has package room, *parking*: has parking space, *partTimeDoorman*: has part-time doorman, *piedATerre*: pied-a-terre's allowed, *protp*: the type of the building (e.g., coop, townhouse), *publicOutdoor*: has access to nearby public outdoor space, *terrace*: has terrace acces, *valetParking*: whether the building has a valet parking service, *virtualDoorman*: whether the building has a virtual doorman, *waterView*: has waterview, *zip*: the zip code of the building.

B.2 Double-debiased Machine Learning

The double-debiased machine learning (DML) family of methods attempts to draw on techniques from the machine learning literature to produce high-quality counterfactual outcome predictions. Following the notation used by Chernozhukov et al. (2017), suppose that

$$Y = \theta D + g(X) + u, \tag{A1}$$

where Y is an outcome variable, D is the treatment variable, g is an unknown and potentially nonlinear function of the high-dimensional vector of observable covariates X, and u is the error term, with $E[u|X, D] = 0$. In the empirical context for our paper, Y is the long-term rental rate, D is the policy variable set by building owners, X is the vector of non-policy attributes, and our goal is to to estimate θ, that is, the effect of the policy choice on rental rates. Suppose now that

$$D = m(X) + v, \tag{A2}$$

that is, that the variation in the treatment is generated by another function of the covariates, where v is the error term, such that $E[v|X] = 0$. Following the "naive" approach that uses a predictive model to estimate θ suffers from bias because of the bias in estimating $\hat{g}$. The DML approach allows us to overcome this problem by utilizing machine learning methods, to obtain estimates of the conditional expectation functions $E[\hat{Y}|X]$ and $E[\hat{D}|X]$, which are then "partialed out" to obtain an estimate for θ.

Chernozhukov et al. (2017) show that the estimator is consistent and unbiased if different samples of the data are used to obtain the estimates of the two conditional expectation functions, and the estimate of the effect is averaged over multiple folds. The main advantage of the DML family of methods is that it allows us to use machine learning methods that combine a large number of covariates, in order to form proxies that predict both the treatment and the outcome variables well—in our context, we found gradient boosting (Chen, He, Benesty et al., 2015) to perform the best. Athey (2018) and Dube, Jacobs, Naidu and Suri (2020) provide excellent discussions of the DML family of methods. In particular, Dube et al. (2020) find that DML estimates of the elasticity of labor supply in an online labor market are very similar to experimental estimates.

B.3 Tables

Table 3: Double ML estimate of effects of sublet policy on long term rentals in NYC.

	Dependent variable:		
	$\log Rent$	$\Delta \log Rent$	
	(1)	(2)	(3)
Building allows subletting (1/0)	0.101^{***}	-0.019^{*}	
	(0.035)	(0.011)	
Δ Building allows subletting			-0.006
			(0.015)
Constant	8.038^{***}	0.001	0.001
	(0.004)	(0.001)	(0.001)
Observations	21,257	21,257	21,257
R^2	0.0004	0.0001	0.00001
Residual Std. Error (df = 21255)	0.520	0.165	0.165

Notes: This table reports the relationship between posted log monthly rental rates and subletting policies. In Column (1), the outcome is the rental rate and the regressor is the subletting policy. In Column (2), the outcome is the residualized log rental rate and the regressor is the subletting policy. In Column (3), the outcome is still the residualized log rental rate and the regressor is the residualized policy variable. The subletting policy variable is residualized with respect to the predictions from an extreme gradient boosting (Chen et al., 2015) using an extensive collection of controls, implementing the DML method developed by (Chernozhukov et al., 2017). Significance indicators: $p \leq 0.10 : *$, $p \leq 0.05 : **$, and $p \leq .01 : ***$.

Table 4: Double ML estimate of effects of dog policy on long term rentals in NYC.

	Dependent variable:		
	$\log Rent$	$\Delta \log Rent$	
	(1)	(2)	(3)
Building allows dogs (1/0)	0.287***	−0.0005	
	(0.009)	(0.003)	
Δ Building allows dogs			0.004
			(0.005)
Constant	7.987***	0.001	0.001
	(0.004)	(0.001)	(0.001)
Observations	21,257	21,257	21,257
R^2	0.046	0.00000	0.00003
Residual Std. Error (df = 21255)	0.508	0.165	0.165

Notes: This table reports the relationship between posted log monthly rental rates and dogs policies. In Column (1), the outcome is the rental rate and the regressor is the dogs policy. In Column (2), the outcome is the residualized log rental rate and the regressor is the dogs policy. In Column (3), the outcome is still the residualized log rental rate and the regressor is the residualized dogs variable. The dogs policy variable is residualized with respect to the predictions from an extreme gradient boosting (Chen et al., 2015) using an extensive collection of controls, implementing the DML method developed by (Chernozhukov et al., 2017). Significance indicators: $p \leq 0.10 : *$, $p \leq 0.05 : **$, and $p \leq .01 : ***$.

Table 5: Effects of having in-apartment washer/dryer on long term rental price in NYC.

	Dependent variable:		
	log *Rent*	Δ log *Rent*	
	(1)	(2)	(3)
In-apartment washer/dryer (1/0)	0.509***	0.058***	
	(0.008)	(0.003)	
Δ In-apartment washer/dryer			0.088***
			(0.004)
Constant	7.919***	−0.014***	−0.0004
	(0.004)	(0.001)	(0.001)
Observations	21,257	21,257	21,257
R^2	0.174	0.022	0.021
Residual Std. Error (df = 21255)	0.473	0.166	0.166

Notes: This table reports the relationship between posted log monthly rental rates and whether the building has an in-apartment washer and dryer. In Column (1), the outcome is the rental rate the regressor is an indicator for an in-apartment washer/dryer. In Column (2), the outcome is the residualized log rental rate and the regressor for an in-apartment washer/dryer. In Column (3), the outcome is still the residualized log rental rate and the regressor is the residualized indicator variable. The indicator variable and the outcome is residualized with respect to the predictions from a extreme gradient boosting (Chen et al., 2015) using an extensive collection of controls, implementing the DML method developed by (Chernozhukov et al., 2017). Significance indicators: $p \leq 0.10 : *$, $p \leq 0.05 : **$, and $p \leq .01 : ***$.