8-1-2019

Do Rent Increases Reduce the Housing Supply under Rent Control? Evidence from Evictions in San Francisco

Brian J. Asquith

W.E. Upjohn Institute for Employment Research

Upjohn Author(s) ORCID Identifier:

<https://orcid.org/0000-0002-5783-5557>

Upjohn Institute working paper ; 19-296

Follow this and additional works at: https://research.upjohn.org/up_workingpapers

Part of the [Labor Economics Commons](#)

Citation

Asquith, Brian J. 2019. "Do Rent Increases Reduce the Housing Supply Under Rent Control? Evidence from Evictions in San Francisco." Upjohn Institute Working Paper 19-296. Kalamazoo, MI: W.E. Upjohn Institute for Employment Research. <https://doi.org/10.17848/wp19-296>

This title is brought to you by the Upjohn Institute. For more information, please contact repository@upjohn.org.

Do Rent Increases Reduce the Housing Supply under Rent Control? Evidence from Evictions in San Francisco

Upjohn Institute Working Paper 19-296

Brian J. Asquith
W.E. Upjohn Institute for Employment Research
email: asquith@upjohn.org

August 2019

ABSTRACT

Rent control balances strong tenant protections with supply-side incentives for landlords. However, cities with rent control are also some of the United States' most unaffordable, prompting questions about how well these incentives are working. I examine how controlled landlords change their housing supply in response to price increases using a well-identified hyperlocal demand shock: the privately operated commuter shuttle systems in San Francisco. Controlled landlords increased market withdrawal filings and became less likely to create vacancies via evictions in response to a shuttle stop placement. Policies raising barriers to market withdrawals prompted controlled landlords to respond by increasing their at-fault evictions.

JEL Classification Codes: R31, R32, R52, K11

Key Words: rent control, evictions, private transportation, LASSO

Acknowledgments:

I'd like to thank Samantha A. Spallone and Eileen Z. Raney for their excellent research assistance. Further assistance is gratefully acknowledged from Yael S. Katz, Kathleen Mastrogiacomo, John S. Asquith, Lauren S. Asquith, Gabrielle A. Leibowitz, Reuven Markov, and Avram Markov. I'd like to thank Jan K. Brueckner, Marianne P. Bitler, David Neumark, Nanneh Chehras, Esther Volchek, Chris Severen, and seminar participants at UC Irvine, the Urban Economics Association, the Southern Economics Association, the W.E. Upjohn Institute for Employment Research, the University of Pennsylvania's Public Policy Initiative, Freddie Mac, Moody's, the Western Economics Association, and New York University's Furman Center. I would also like to thank the many San Francisco city officials who helped me acquire data and better understand their process, including Robert Collins and Edward McCaffrey. Further thanks go to Kate Pennington and the San Francisco City Planning Department. This paper was funded by generous grants from the UC Irvine Economics department. Any errors are my own.

Upjohn Institute working papers are meant to stimulate discussion and criticism among the policy research community. Content and opinions are the sole responsibility of the author.

Introduction

Rising housing unaffordability in several U.S. coastal cities have reignited public interest in rent control. While recent work by Sims (2007), Autor, Palmer, and Pathak (2014, 2019) and Diamond, McQuade, and Qian (2019a,b) has examined the effects of imposing or removing rent control on the housing stock as a whole, these shifts of large numbers of units from uncontrolled to controlled are very rare. In most cities, rent controls are long-standing regimes commanding popular support. Policymakers interested in improving mature systems unfortunately have little empirical evidence to draw on from the economics literature. Given that the major U.S. cities that have had rent control for 30 years or more—such as New York, Los Angeles, and San Francisco—are renown for their unaffordable rents, it is worth examining how controlled landlords change their supply in response soaring rents.

The disparity between the intention (rent regulation) and the outcome (sky-high rents) 40 years into the policy raises two questions: 1) Does rent control dampen landlords’ willingness to supply units? Or worse, 2) could controlled landlords even decrease in supply as price rise? Modern rent-control systems work by strictly capping within-tenancy rent increases, restricting evictions, and automatically renewing leases, so that rents can come to lag far below market for long-staying tenants. All U.S. cities with rent-control-exempt new construction and owner-occupied housing,² so when it is cost-effective, landlords have incentives to convert to owner-occupied housing or demolish and rebuild. In almost all cases, this represents a permanent loss to the city’s controlled housing stock and puts additional price pressure on the remaining apartments. These rules are relatively straightforward, but also create incentives for landlords to circumvent the system. These cities have thus layered on additional rules regulating how landlords can change their housing supply so that determining whether landlords change their supply controlled units when prices rise is often not

²New York City alone has a program that grants a tax exemption to owners of new buildings who voluntarily choose to be rent stabilized —the NYC term for what is called rent control in California. This distinction in New York City exists because there are still buildings that have WWII-vintage, first-generation hard rent controls, which go by “rent control” locally. After the credit’s expiry, all new vacancies are automatically decontrolled.

clear even to policymakers themselves.

In this paper, I examine how controlled landlords change their housing supply in response to identifiable demand increases. Four research questions regarding rent controls are addressed. First, do controlled landlords respond to price signals by increasing or decreasing their controlled housing supply? Second, how do supply changes vary across different sub-markets? Third, even if landlords supply the same or more number of units, do they respond to price increases by seeking to maintain housing quality? Lastly, how do policy attempts to shape controlled landlord behavior change their willingness to supply units?

I overcome two constraints to answer these questions. The first is constructing treatment and control groups that did or did not experience a local demand shocks. San Francisco is this paper’s research setting because it had a unique, spatially-varying, well-identified demand shock: the network of commuter corporate shuttles stops operated by Google, Apple, Facebook, and Electronic Arts (EA), which transport employees from various sites around the city to Silicon Valley and is a highly valued employee benefit (Dai and Weinzimmer 2014). I show in this paper that the value of the shuttle’s transit amenity is capitalized into housing prices, and large enough that controlled landlords were unable to fully price the amenity into rents for existing tenants. In the absence of a comprehensive month-by-building rents dataset, shuttle stop placements serve instead to proxy (but not instrument for) for time- and area-varying demand increases.

The second overcome constraint is that housing supply changes in the controlled market cannot be measured by housing starts, namely because of the new construction exemption.³ Landlords can still change their supply along the intensive margin, namely by increasing (up conversion) or decreasing (down conversion) their unit counts. The landlord’s main obstacle to up or down conversions is winning approval from the city planning department, which may give greater scrutiny to the permits they believe threaten the controlled housing supply. Another means of changing their housing supply is via a legal mechanism called no-

³See also Asquith (2019a).

fault evictions. These no-fault evictions allow a landlord to withdraw one unit or an entire building from the controlled market for several years. Landlords being more likely to perform a no-fault eviction when faced with a demand increase is strongly suggestive evidence that they believe that their medium-term returns will worsen.

In theory, San Francisco and other cities attempt to design the controlled market to give landlords incentives to continue to supply and maintain their units. Landlords are allowed to pass-through the costs of capital improvements, maintenance, and utilities to their tenants. Landlords are also incentivized to return units to market via vacancy bonuses, which allow landlords to freely negotiate the base rent with a new tenant. This can range from a 10 percent increase over the rent paid by the vacating tenant in Washington, D.C., to whatever the market can bear in all California cities.⁴

However, these markets may nonetheless experience supply failures for two main reasons. The first is that returns are almost always lower than in the uncontrolled market, and these losses compound with tenancy duration. Maintenance cost pass-through rules and annual rent increase caps may be insufficiently generous, so that landlords receiving very low rents from long-standing tenants may calculate it is better to exit the rent control market rather than continue to accept substandard returns.

The second reason is that modern rent-control systems create adverse selection problems for landlords (Miron 1990).⁵ Controlled landlords prefer short-stayers but cannot directly observe the tenant's duration preference. Long-staying tenants, knowing this, are incentivized to present themselves as short-staying tenants. Landlords must then maximize their profits in the face of uncertainty engendered by imperfect information on the tenant's true type (Basu and Emerson 2000, 2003; Arnott and Igarashi 2000; Hubert 1995; Iwata 2002), raising costs for everyone.

Like in any other market with adverse selection, price increases may cause landlords to

⁴See Table B3.

⁵Miron's paper, like all others in the literature, considers only the extreme case that evictions are entirely prohibited.

seek to exit the controlled market (Akerlof 1970) and discourage them from seeking a new tenant. Basu and Emerson (2000, 2003) point out that landlords’ adverse selection problem worsens when rents increase because only tenants preferring long durations are willing to pay the controlled unit premium.⁶ The crux of the problem is that unlike nuisances and rent delinquents, controlled landlords lack any direct remedy for long-stayers. Tenants have indeed been repeatedly found to disproportionately have long tenures in all forms of rent control, and specifically so in the case of San Francisco by Diamond, McQuade, and Qian (2019b).⁷ All these problems will be exacerbated as rents rise, and in fact, might also “push” more landlords to exit to avoid an increasingly adversely selected pool of tenants (Miron 1990; Basu and Emerson 2003). The net effect is that in practice, it is difficult to know *ex ante* whether the controlled housing supply will rise or fall.

This is not the only attempt to quantify housing supply changes under rent control in San Francisco. Diamond, McQuade, and Qian (2019a,b) also uses San Francisco data and purport to find that landlords decrease their housing supply generally under rent control. Distinguishing my contribution from theirs is difficult because in those places where their paper most substantively overlaps with mine, their paper makes a series of policy misinterpretations that cause the authors to arrive at inaccurate conclusions.⁸ A theoretical paper by McFarlane (2003) also addressed this question directly, modeling a controlled market with vacancy decontrol and showing that under stable price growth, the long-run supply would be essentially unchanged with or without rent control. However, that paper’s findings hinge critically on assuming that if a building’s capital depreciates, it can be demolished, rebuilt, and have controls reimposed. Since no current market in the United States operates under these rules, its policy relevance in the United States is limited.

⁶Several papers have empirically and theoretically found that tenants do pay a premium for rent control and security of tenancy, including Raess and von Ungern-Sternberg (2002), Skelley (1998), Arnott (1995), Arnott and Igarashi (2000), and Nagy (1997).

⁷A list of studies finding longer durations in other systems include Linneman (1987); Gyourko and Linneman (1989); Munch and Svarer (2002); Nagy (1995, 1997); Ault, Jackson, and Saba (1994); Krol and Svorny (2005); Svarer et al. (2005).

⁸I discuss the problems of their paper in some depth in Asquith (2019b), and I refer the reader to that paper for clarity on my contribution to the literature vis-à-vis theirs.

The findings here are informative to policymakers interested in how well the controlled housing stock absorbs demand shocks. If the evidence shows that the system absorbs price increases by expanding supply, even weakly, this would speak well to its policy design and may point toward other culprits for San Francisco’s affordability crisis. If instead the evidence suggests housing quantity and quality decline when prices rise, this would indicate that there are serious problems in the design of modern rent controls that policymakers may have to address by enacting reforms to restore affordability.

I first find evidence the shuttle raised local prices. For every 1km a stop gets placed closer to a condominium (out to 2km), the condo’s sales value increases by 18 percent, which represents a very large increase. Landlords are most likely to respond to this large demand shock by temporarily reducing their housing supply via no-fault evictions, chiefly via a mechanism that singles out one tenant per building. This response is strongest among landlords with three to six units, who comprise 38.5 percent of the controlled stock (29.7 percent of total controlled units). Landlords of very small buildings (two-unit buildings) are particularly likely to respond to large local demand shocks by withdrawing their entire buildings, although this is partially offset by a decline in these kinds of evictions among landlords of large buildings (six or more units). These results are substantially weakened when excluding “historic” buildings—defined as built before 1906—which speaks to how San Francisco’s restrictive rules for passing through capital improvement costs to tenants may be driving landlords of older buildings out of the controlled market.

There is no evidence that controlled landlords respond to these large shocks by filing to increase their unit count (up-converting), but neither is there any evidence that they respond by down-converting. The most holistic explanation of my findings is that in a market with adverse selection, the landlords least confident of being able to properly screen tenants or get rid of ones who become long-stayers will choose instead a form of market exit. Further evidence for adverse selection’s impact is that controlled landlords appear less likely to turn out tenants via at-fault evictions in response to a local demand shock, in spite

of strong incentives to do so. Lastly, there is evidence that landlords of two-unit buildings do respond to a demand shock by increasing their repair permit filing, which is evidence in favor of Arnott and Shevyakhova (2014)’s hypothesis that tenancy rent control gives landlords incentives to spruce up their buildings in between tenancies to take advantage of vacancy bonuses. However, their hypothesis is that most maintenance would occur between tenancies, and taken together with the evidence above that involuntary turnover declines and increased withdrawals of these buildings, it is more likely that increased repair permit filings reflect that landlords are sprucing up their buildings to take advantage of the fact that these newly vacant buildings will sell for a higher price.⁹

The instrumental variables (IV) results indicate that the shuttle-induced price increase is large, and the IV supply results are also very large relative to the dependent variable values. These large magnitudes mean that applying a straightforward policy analysis would likely overstate net supply effects that city officials could actually expect. The OLS results are much more modest, and the results from the hedonic price analysis clearly show that the OLS results are downwardly biased, but still have price effects just strong enough to induce some sort supply response. Comparing the OLS to the IV results can be thought of comparing what would occur after a modest price shock that can be largely priced into rents against what would occur after a large price shock that mostly could not. These results indicate that a small price shock creates effects most along the lines of what policymakers would hope for: an increase in filing permits that would up-convert (add units to) the building and an increase in repair permit filings, and no significant changes in evictions. These results speak to the power of the annual rent increase cap, the keystone of rent regulation. San Francisco’s policy restricts rent increases to 60 percent of the regional CPI, so that there may well be strong ameliorative supply effects from the relatively minor reform of allowing rent increases to be indexed to inflation. Instead, if the city is reluctant to reduce any tenant benefit, trying instead to restrict market withdrawals will likely redound against tenants in other

⁹The San Francisco Tenants Union reported that realtors claim that vacant buildings are worth 20 percent more. Source: San Francisco Tenants Union, <https://www.sftu.org/ellis/> (accessed July 15, 2019).

ways. I show that when the city tried to force landlords to distribute relocation payments to tenants after a whole building withdrawal eviction, or restrict the ability of landlords to do no-fault evictions during the school year, landlords responded by increasing their at-fault eviction rate.

Section 1 sketches the San Francisco policy environment. Section 2 details the data used to test the hypothesis. Section 3 explains the empirical strategy. Section 4 reviews the shuttle’s estimated price impact. Section 5 overviews the paper’s empirical supply analysis. Section 6 reviews the results generated from executing the empirical strategy. Section 7 concludes.

1 Policy Setting

This section gives a brief overview of San Francisco’s rent control and evictions policies. Policy details will be referenced throughout the rest of the paper to help interpret the results. The four key provisions of San Francisco’s Rent Ordinance are:

1. **Rent increases are capped at 60 percent the rate of inflation**, but certain costs and hardship provisions exist to ensure that landlords can earn a profit.
2. **Security of tenancy.** Landlords cannot refuse to renew the lease of a tenant in compliance. Any eviction must be for “just-cause” and approved by the San Francisco Rent Board.
3. **Vacancy decontrol-recontrol (or vacancy decontrol).** Tenants and landlords are free to negotiate the base rent, and only subsequent increases are controlled (capped). After a tenant vacates, the landlord can freely negotiate a new base rent with the next tenant and once again controls only apply to subsequent rent increases.
4. **No new controlled buildings.** Only buildings built before June 13, 1979, and with two or more units are subject to rent control.

In summary, rent control works by locking landlords and tenants into a base rent that slowly erodes in real terms with each passing year. Base rents, set for a new tenant, are unregulated.

There are only three paths to full decontrol: demolition, condo conversion, or substantial rehabilitation. Each of these are tightly regulated by the city. Asquith (2019b) discusses all three at length, but the reason these are not considered here is that while they change the housing supply, they are inappropriate for use as dependent variables. In the case of demolitions or substantial rehabilitations, they are simply too rare. San Francisco only granted 114 demolition permits for these buildings between 2003 and 2013, or less than one per month. Further, only two buildings were granted a rent control exemption for substantial rehabilitation over this time.

In the case of condo conversions, these were only allowed via an annual lottery for two-to-six-unit buildings that was suspended in 2013. Two-unit buildings are allowed an unlimited lottery bypass, subject to individual level ownership rules. For these reasons, they are not included in the present study, but the lottery did shape landlord behavior, which I further discuss below.

As stated in Provision (2), rent-controlled landlords must have a “just cause” for an eviction. The 15 grounds for a “just-cause” eviction are given in Table B1 and cover various lease breaches, like nonpayment of rent. There are seven grounds for which a tenant is “at-fault”,¹⁰ and eight grounds where the tenant is not at fault, or “no-fault.” Three of these no-fault evictions are not part of this study because they are either for temporary improvements made to the building or very rarely permitted by the city.¹¹

I focus on two no-fault eviction types that involve withdrawing housing from the controlled market. Owner move-in (OMI) evictions allow a landlord to reclaim a unit for their

¹⁰These are 1) on-payment or late payment on rent, 2) general breach of the lease, 3) nuisance, 4) illegal usage of unit, 5) refusing to quit after previously agreeing to end tenancy, 6) refusing to grant landlord lawful access, and 7) sole remaining tenant is an unapproved subtenant.

¹¹These include temporary eviction for lead abatement or capital improvements; and revoking of “Good Samaritan” status for tenants who are fleeing natural disasters.

use or for a close family member. Ellis Act evictions allow landlords to withdraw all of their units from the market. These are advantageous to landlords because, unlike demolitions, there are more minimal bureaucratic delays. However, demolitions are most preferred because the property can be more immediately converted to uncontrolled housing, because of the new building exemption (Provision (4)).

One important rent-control exit channel, condo conversion, is restricted around how many units the building has. For buildings with seven or more units (hereafter “large” buildings), the only path to full decontrol is demolition and reconstruction. For these buildings, the more general question of whether controlled landlords will seek to increase their long-run supply either by up conversion or creating a vacancy via tactical at-fault eviction or buyout. Since landlords are rational, forward-looking actors, they can price in expected tenancy duration into base rents (Nagy 1997). The empirical prediction would be that for large landlords, they will respond to price increases by increasing their at-fault eviction rate, which is observed in this paper, or their buyout rate, which is not.

These rules represent the “base” of how rent control works, but landlords over the years have challenged or circumvented many of them so that the city has superimposed a second, more complicated level of rules to constrain landlord behavior. These rules are explained in more detail in the AppendixB. These two sets of rules give the system a cantilevered quality that creates complicated interactions between the all the regulation. Predicting how supply will respond to these interactions is thus difficult. For the reader’s ease, I will reference the AppendixB where the additional policy details are most relevant.

2 Data

Data for controlled, “just-cause” evictions come courtesy of the Rent Board of the City of San Francisco. The Rent Board publishes information on all just-cause evictions since January 1, 1993, and are publicly available by request. The dataset includes the eviction

address, the date served, and the eviction reason. This dataset totaled 30,992 evictions. I worked in conjunction with Kate Pennington of UC Berkeley to correct errors in the recorded addresses, and the resulting dataset is her proprietary work which is available upon request. The final dataset spanning July 2003–December 2013 includes 10,004 at-fault and 4,032 no-fault eviction notices.

Data on issued eviction notices in uncontrolled units were obtained from San Francisco Superior Court unlawful detainer filings. Court documents include scanned copies of eviction notices and other relevant information. Of these scanned documents covering July 2003 through January 2014, 40,997 were manually examined for whether the notice was sent to a residential, uncontrolled unit and for the eviction cause. A total of 1,249 evictions from uncontrolled units were recorded, including those otherwise eligible buildings exempted by the landlord from the ordinance.¹² These were appended with the controlled evictions dataset to create a master dataset of eviction notices.

Up conversions, down conversions, and repair permit filings were tabulated from the universe of building and demolition permits information made available by the San Francisco Planning Department via data.sfgov.org. Due to vagaries in how property owners filed permits, conversions could be reported as either a building change or a demolition. In both permit datasets, filing date, issue date, and existing and proposed unit counts are given, so that changes at the property lot level could be observed. This came to 1,325 up conversions in 944 properties between June 2003 and December 2013, and 526 down conversions in 418 properties over the same period.

Evictions and conversions were then matched to buildings. Data for building characteristics come from the City and County of San Francisco’s Office of the Assessor-Recorder. The same dataset also has transaction values and dates.¹³ Supplemental information on addresses, latitude and longitude, and other property parcel characteristics was obtained

¹²A list of these buildings can be made available upon request.

¹³Discrepancies in reported sales prices were compared with publicly reported information from the real estate website www.zillow.com

from the website of the City and County of San Francisco.¹⁴ The dataset covers assessments made yearly from June 2003 to June 2014, for 2,373,721 observations, ranging from 188,333 parcels in 2003 to 205,229 parcels in 2014. Some records were self-evidently corrupted, and were cleaned through manual inspection and by comparing other records for the same parcel. All identified public housing complexes are dropped.

I refer throughout this paper to buildings and parcels interchangeably, but there is a distinction for a few properties. Individual properties are assigned a unique block and lot ID. In some cases, properties that are clearly part of the same development, such as condominiums, assign block and lot IDs to each individual unit in the same building. A second property ID, called a map lot ID, was used to identify these cases, although for most buildings, the map lot IDs and block and lot IDs were the same. I use the map lot ID as the unit of observation and for merging on the controlled evictions data onto my property panel. Before collapsing to the map lot ID level, I retained only those block and lot parcels with years built up through 2002, and had at two or more units at the beginning of the study period. I also defined eligible buildings as those with property codes for apartments (“A”, “AC”, and “OA”), mixed apartments and dwellings (“DA” and “DD”), and flats (“DF”, “F”, “FA”, and “F2”). This leaves 34,708 controlled and 1,433 uncontrolled parcels. Any parcel or eviction that is in Golden Gate Park, the Presidio, or Yerba Buena and Treasure Islands is dropped. The former two are dropped because they are park properties and the latter two because they are small, noncontiguous parts of the city.

Table 1 presents information about the housing market in San Francisco. San Francisco’s housing stock is disproportionately older, and apartment buildings comprise just over half of the stock. The majority of housing units are also rent controlled, and the decade between 2003 and 2013 saw a decrease in the number of rent-controlled housing units, apartment units, and residences. As the housing stock modernized, rental units were replaced with owner-occupied single-family housing. The other notable feature is that on average, the

¹⁴data.sfgov.org

nearest shuttle stop is less than a half-mile away, meaning that condos did in fact experience significant “treatment” from the placement of shuttle stops.

3 Empirical Overview

This paper asks whether rent-controlled landlords decrease their housing supply via evictions or conversions in response to an increase in local demand. The ideal empirical strategy would be to estimate a reduced-form equation such as

$$\begin{aligned} Outcome_{it} = & \zeta_0 + \zeta_1 \ln(Rent_{it}^{RC}) + \zeta_2 \ln(Rent_{it}^{UC}) + \zeta_3 \ln(Rent_{it}^{RC}) \times \mathbb{1}\{RentControl_i\} \\ & + \zeta_4 \ln(Rent_{it}^{UC}) \times \mathbb{1}\{RentControl_i\} + \delta_t \lambda_t + \delta_i \gamma_i + \Theta X_{it} + \epsilon_{it}, \end{aligned} \quad (1)$$

where $\log(Rent_{it}^{RC})$ is the log of the contract rent for a vacant unit in building i during time period t if it were controlled; $\log(Rent_{it}^{UC})$ is the log in the contract rent for a vacant unit in the same building if it were uncontrolled; λ_t is a fixed effect for time period t ; γ_i is a fixed effect for the building; and X_{it} are other building, neighborhood, or temporal characteristics. I use the rent-control status of the building at time $t = 0$ (July 2003) above as $RentControl_i$ to avoid endogeneity bias and with the building fixed effect absorbing the uninteracted rent-control term.

It is important that both rents be included here, because controlled landlords are very likely responding to signals from both markets when making their supply decisions. However, including $\log(Rent_{it}^{RC})$ and $\log(Rent_{it}^{UC})$ creates prohibitive challenges. Even if an exhaustive dataset of unit-by-month rents existed, and there was some way to ascertain what the controlled rents would be in uncontrolled units (and vice versa), rents are endogenous to evictions and conversions. This makes consistent estimation of Equation (1) impossible. Since consistent rents are both systematically unobservable and endogenous, this paper instead proxies for changes in rents using changes in local transit amenities created by the roll-out of Google, Apple, Facebook, and Electronic Arts’ commuter shuttle stop program.

The key assumption I make here is that the commuter shuttles change rents in both markets by the same magnitude, so that Equation (1) is refitted to become

$$\begin{aligned} Outcome_{it} = & \zeta_0 + (\zeta_1 + \zeta_2) Shuttle_{it} + (\zeta_3 + \zeta_4) Shuttle_{it} \times \mathbb{1}\{RentControl_i\} \\ & + \delta_t \lambda_t + \delta_i \gamma_i + \Theta X_{it} + \epsilon_{it}, \end{aligned} \quad (2)$$

where $Shuttle_{it}$ is one of two measures for proximity to a shuttle stop discussed below. While future research will have to tease out how landlords respond differentially to signals from the two markets (ζ_1 versus ζ_2 and ζ_3 versus ζ_4), this approach can clearly identify the net effect of the two price signals on eviction, maintenance, and conversion probabilities across the two markets.¹⁵

It is worth stressing that the shuttles are only a proxy for localized rent increases and not intended to be an instrument for them, largely because the shuttle stops were not randomly distributed throughout the city. Figure 1 shows how the shuttle stops spread throughout San Francisco, starting from September 2004, when Google placed the first two stops, to the study period’s end in December 2013. A more extensive description of the history and data collection behind the shuttles can be found in the AppendixC, but a brief overview illustrates why they are a useful proxy for localized rent increases. The shuttles were first initiated by Google in September 2004 and went to just two locations: Glen Park BART Station and a park and ride stop near Candlestick Park. Apple started its own system in 2007, Facebook in 2009, and Electronic Arts in 2012.¹⁶ By 2009, the shuttles had come to cover many city neighborhoods either completely or in part, particularly in the eastern

¹⁵This assumption is almost surely wrong, but because of data limitations, this is the assumption that must hold for empirical tractability. However, if anything, this assumption slightly downwardly biases the stated impact among controlled landlords. Rents for vacant units in the controlled market likely rise more since that market is more supply-constrained. $Shuttle_{it}$ represents the average change between the two markets, and since the controlled market is much larger than the uncontrolled market, $Shuttle_{it}$ is going to mostly represent a slight understatement of the effect from the change in the rent signal in the controlled market.

¹⁶Yahoo initiated its own service in 2005, followed by Genentech in 2006 and others after the recession (Dai and Weinzimmer 2014). Genentech and Yahoo! shuttle stops are not included in this study, but existing information indicates that they overlapped with Google, Facebook, and Apple stops almost completely.

half of the city. The shuttle stops clearly favored the city’s east and northeast, and so like $Rent_{it}$, the shuttle stops’ placements are likely endogenous to local conditions. Therefore, the paper instruments for shuttle stop locations by exploiting exogenous constraints on their placements. The commuter shuttle stops can only be located at large public stops, called “bus zones” hereafter, because the shuttles are otherwise too large for San Francisco’s streets. Patterns for shuttle stop placement are detailed below in Section 3.1.

The empirical investigation is thus conducted in two steps. The first step establishes that the commuter shuttles’ transit amenity caused rent pressures plausibly strong enough to have induced the marginal (and supramarginal) landlord to change their housing supply because prices changed by more than the values given in Conditions 1 and 2. I demonstrate this via an IV approach for shuttle stop placement. The second step of the investigation tests for supply responses in controlled units by exploring how supply changed in response to greater commuter shuttle coverage and thus greater demand and higher prices. Estimates on changes in evictions after exposure to both the observed and fitted shuttle placements are then presented.

AI also report results for two samples: one with and without buildings built before 1906, the year San Francisco experienced an earthquake and a fire that razed a large section of the city. The rebuilt areas were significantly more dense than the untouched neighborhoods (Siodla 2015). The year 1906 represents thus a clean cutoff between what are more and less likely to be deemed historic properties, and further, buildings whose capital is likely to be far more depreciated than others. One might expect at least two important differences between buildings built before and after 1906. The first is that demand for more historical housing units might be fundamentally different. This means that the shuttle amenity might be stronger (or weaker) for older buildings. Second, older buildings’ capital has depreciated far more than the buildings built in the uncontrolled sample (built 1979–2002), or even the “younger” controlled buildings ($\sim$ 1945–1978). The capital improvements restrictions discussed in the AppendixB.4 mean that older buildings have perhaps even stronger incentives

to exit the controlled market than newer ones. Thus, breaking out the results into different samples based on the age of the building is important for understanding the impact of capitalization and redevelopment rules on the controlled housing supply.

Since an IV estimation strategy is needed, generating unbiased estimates can only occur if a strong first-stage regression can be identified. I will use a constant coefficients approach that draws on exogenous constraints on placement and placement timing shocks, along with existing features of the city. To improve the chances that I properly measure the shuttles' pricing impact via a strong first-stage instrument, I use two measures for proximity to a shuttle stop. The first is an indicator for being within a half mile of a shuttle stop, $\mathbb{1}\{Shuttle_{it}\}$. The second is *Shuttle2km*, which measures how far a condo is to the nearest stop's two kilometer outer radius, calculated as

$$Shuttle2km_{it} = \frac{(2000 - Distance_{it})}{2000} \times \mathbb{1}\{Distance_{it} \leq 2000\}, \quad (3)$$

where $Distance_{it}$ is the raw distance (in meters) from the closest shuttle stop and building (or condo) i in month t . The idea here is that it gives a continuous measure for how far a given building is from a stop up to a distance of two kilometers, after which its assumed in a city as geographically compact as San Francisco that price effect would have largely petered out. By taking the additive inverse over this distance, the measure now has the largest values closest to the stops and a value of zero at the boundary and beyond. The hedonic model coefficient will thus be positive in the event of a positive price effect and vice versa, which eases interpretation for the reader. Lastly, the overall interpretation is (relatively) simple: all else equal, this measure captures the price effect of a stop being placed one kilometer closer to the building, given that the building is within two kilometers of a stop. Both measures are assigned a value of zero in the period before the first shuttle stop in September 2004.

3.1 Instrumenting for Shuttle Placement

Because the shuttle stops were not randomly distributed across space and time, the tech companies may have chosen where to place shuttles in response to local price changes; for example, by placing them where tech company employees demand was already locating. I thus instrument for shuttle stop placement via a constant coefficients approach, whose identification strategy is based on exploiting two sources of exogenous shuttle stop placement predictors. The first predictor is that the shuttles could only use public bus stops long enough to accommodate the 50-foot plus motorcoaches, called “bus zones.”¹⁷ Figure 2 shows the location of the 870 bus zones that match this description, which are about 25 percent of the SFMTA’s total public bus stops. The greatest concentration of eligible bus zones is in the far northeast part of the city, near the central business district, and extending directly west to Golden Gate Park and directly southwest into the Inner Mission. Outside of these areas, eligible bus zones are much sparser.

Satisfying the exclusion criterion for using the eligible bus zones rests on the assumption that while local residents’ location decisions may be influenced by being near a public bus stop, residents are almost surely indifferent to the closest bus stop’s *length*. The 50-foot constraint is relevant for shuttle placement but not for local prices or any other factor correlated with evictions.

Yet, the bus zones’ distribution alone is not enough to predict shuttle placement in the city for two reasons. The first reason is that even with the 50-foot restriction, there are many more eligible bus zones than shuttle stops. For my instrument, I thus generate many different bus zone characteristics measures that could have potentially influenced the companies’ choices *among* the set of eligible stops.¹⁸ The idea is that conditional on the fact that I am already

¹⁷Information on the length of motorcoaches used by the companies was unavailable, but MCI motorcoaches that seat roughly the same number of people that Google’s shuttles allegedly do are just over 45 feet long. Anecdotally, Google shuttles appear to be a bit longer, so that 50 feet was selected as the cutoff. The shuttles need room to maneuver in and out of the stop, so a more reasonable cutoff might be more like 70 feet. Source: <http://www.mcicoach.com/luxury-coaches/passengerJ4500.htm>, last accessed August 28, 2016.

¹⁸Greater detail on how this was done can be found in the AppendixD

controlling for the building/condominium’s transit-proximity characteristics, then variation in the eligible bus zone’s proximity to these other transit corridors will predict the building’s or condo’s shuttle exposure.

The second reason that the bus zone distribution and their characteristics are insufficient is that the bus zone distribution was almost completely static over the course of the study period. To resolve this problem, I exploit the second exogenous predictor on shuttle stop placement: the value of Google’s (GOOG) stock price. This is advantageous for several reasons. First, Google’s initial public offering (IPO) was August 19, 2004, just a month before the shuttle service started. GOOG’s IPO was an exogenous-to-local-conditions driver of San Francisco housing demand, created by Google employees newly able to publicly trade their stock holdings and buy property. Further, because the value of GOOG is zero before August 19, 2004, the excluded instruments have a value of zero in an almost perfect (and probably not entirely coincidental) overlap with when the shuttle measures have a value of zero. Second, inasmuch as employees used stock windfalls to buy property in San Francisco or finance living in the very expensive city, it is logical that stock price surges would have exogenously boosted demand for city living and thus for expanding the shuttle network. Lastly, Google was the leader of expanding the shuttle network, and the other companies clearly sought to keep up with Google’s offerings, so it is reasonable and parsimonious to use just post-IPO fluctuations in the value of GOOG as the main source of exogenous time variation.

More information on the construction of all candidate bus zone characteristics is in the AppendixD, which expands on the information given above. Once all interactions are generated, the total number of candidate explanatory variables is 1,578, a number of excluded instruments almost guaranteed to experience the sort of poor asymptotic properties associated with a many-IV design (Chao et al. 2012). For that reason, I use the LASSO-based IV estimator first developed in Belloni et al. (2012), and expanded further in Belloni et al. (2016) to include cluster-robust standard errors in a panel data setup. The LASSO

procedure here works in two steps. In the first step, the LASSO procedure chooses the GOOG-interacted bus zone characteristics that best predict my endogenous shuttle stop variable. Importantly, the estimator allows me to select certain variables to always be included, which in this case would be confounding condominium or building characteristics plus time and space fixed effects. In the second step, I estimate a classic first stage OLS regression that I can use an IV regression (post-LASSO regression). Since this explicitly a constant coefficients approach, as long as the variables chosen by the LASSO procedure are exogenous to the condominium price but do predict the presence of a commuter shuttle stop in a interpretable manner, then doing model selection via machine learning will still yield consistent second stage estimates I can draw causal inferences from (Belloni et al. 2016).

Before proceeding, it is worth noting that this same procedure will be used to build instruments for estimating the hedonic equation and for the housing supply responses, but the LASSO algorithm will return different predictors in each case. This is mainly because I can include building fixed effects in the panel setting which will absorb almost all of the underlying static transit and bus zone characteristics, whereas in the hedonic regression, I will have to control for observable characteristics directly so that it is more credible to claim that instrument really is orthogonal to price conditional on controlling for the bus zone and transit characteristics that residents actually value. I control for observable bus zone and transit characteristics in various ways to buttress my claim that my LASSO-generated instrument is valid and the resulting hedonic effect is a credible estimate that can be used for the housing supply regressions. When I turn to the housing regressions, the building fixed effects organically reinforce the LASSO-generated instrument's validity so I will focus on other forms of robustness.

4 The Hedonic Price Effect of the Commuter Shuttles

In order to deploy the shuttles as a proxy, but not an instrument for, highly-localized rent increases in San Francisco, I first demonstrate that the shuttles did have a positive price effect and attempt to quantify the magnitude. In order for the shuttles to be a suitable proxy, they must have a hedonic effect that exceeds the highest observed annual rent increase cap—2.2 percent—over the course of the study period. The marginal landlord would then be left with some unrealized amenity capitalization that might then tempt them to use other mechanisms to increase their returns. Further, the per-unit level increase should also ideally exceed the relocation costs, of about \$20,000, so that it covers the fixed costs of a no-fault eviction.

In the absence of consistently-observed rent data, condominium sales in San Francisco from July 2003 to December 2013 are used to estimate the effect of the shuttle stop placement on housing prices. The hypothesis being tested is that the shuttles provide a transit amenity that technology company employees might value highly and are thus likely to induce employees to strongly prefer shuttle-adjacent homes. The distance from San Francisco to worksites in Silicon Valley can be as long as 50 miles, so that many employees can now use their practically door-to-door, free commutes productively via the on-board wifi (Dai and Weinzimmer 2014). Although the shuttles are a private good, San Francisco’s housing market is very tight, so that it is quite likely that a privately provided transit amenity would still have been aggressively capitalized into housing prices.¹⁹

In the property records, the most recent transaction date is recorded each year. I selected only those single-unit condominiums that reported an arms-length transaction, for a total of about 32,843 transactions spanning July 2003 to December 2013. I further restricted the sample to include only those with a reported price between \$100,000 and \$12,000,000 (in

¹⁹Several news stories have reported that IPOs prompt housing demand spikes, particularly in San Francisco. See, for example, Gebel, Meira, February 6, 2019. “Home Prices in San Francisco are Already Among the Most Expensive in the US, and They’re Expected to Skyrocket When Tech Companies Like Uber and Slack Go Public in 2019.” *Business Insider*. Butler, Fauver, and Spyridopoulos (2019) document this phenomenon more broadly across a variety of outcomes.

2013 dollars), and a minimum of 291 square feet.²⁰ This restricted the sample size down to 24,969 for condos built between 1906 and 2013, and to 26,884 when including all condos.

The first assumption, for proper identification, is that there are no other unobserved shocks to the outcomes coincident with shuttle placement that affected pricing outcomes. The second is that changes in condominium prices offer persuasive evidence that rents also increased. This second assumption is buttressed by noting that condominium and apartment unit supply are endogenous (Sinai 2008),²¹ even when there is rent control (Häckner and Nyberg 2000). Further, condominium sales prices also provide a useful way of appraising the net present value of the transit amenity, which is germane because landlords weigh the increased sales value of the unit against the relocation payment. Thus, while I cannot determine exactly how the shuttles change controlled rents, quantifying it from condominium sales is sufficient for this paper’s purposes to establish that the premium almost surely exists in apartment rents.

The hedonic price equation to estimate the shuttle amenity in condominium sales takes the following form:

$$\begin{aligned} \ln(\text{Price}_{it}) = & \alpha + \beta_1 \text{Shuttle}_{it} + \beta_2 \text{New}_{it} + \beta_3 \text{SqFt}_{it} + \beta_4 \text{YearBuilt}_{it} + \beta_5 \text{FullBaths}_{it} \\ & + \beta_6 \text{Beds}_{it} + \beta_7 \text{OtherTransit}_{it} + \beta_8 \mathbb{1}\{\text{HalfBath}_{it}\} + \nu \text{Nbrhood}_i + \delta_t + \epsilon_{it}, \end{aligned} \quad (4)$$

where (as above in Equation (2)) Shuttle_{it} is a measure for condominium i ’s exposure to the nearest Google, Apple, Facebook, or Electronic Arts shuttle stop at time t . This is

²⁰Buildings built prior to 1906 have special historic status, due to surviving the 1905 Earthquake and subsequent fires. Many, but not all, have historic preservation status. Since consumers may have altogether different preferences for historic versus non-historic buildings, the older condos are excluded from the sample. The 291-square-foot minimum was chosen from a 2015 news story that claimed that this was the smallest condo sold in San Francisco: CBS SF Bay Area, “San Francisco’s Tiniest Closet-Like Condo Just Sold For \$415K”, April 15, 2015. URL: <http://sanfrancisco.cbslocal.com/2015/04/15/san-franciscos-smallest-condo-just-sold-for-415k-soma-south-of-market/>, last accessed September 15, 2016.

²¹In my sample, the correlation in the sales price between condominiums and apartment buildings was 0.645.

the treatment variable of interest. Control variables include $OtherTransit_i$, a collection of measures of how far the condo is from other forms of transit defined in various ways. This includes information on how far the condominium is from the BART and Caltrain, to the light rail transit stops, and to major north-south thoroughfares. Specifically, these are the distances in miles to the condominium's nearest BART station, Caltrain station, MUNI Metro station (but not cable cars), and nearest major thoroughfare segment. The major north-south thoroughfares are 19th Ave/State Highway 1, and U.S. Highway 101 and extensions I-80, and I-280. I also control for the distance to the central business district, under the assumption from the monocentric city model that prices are highest closest to it, and distance to the geographical center of the city, allowing for the possibility that prices might rise more away from the center (and closer to the coast). $SqFt_{it}$, the total area of the condo in square feet; $YearBuilt_{it}$, the year the condominium was built; and New_{it} , whether the sale occurred within the first two years after the unit was built. I also include dummies for the number of bedrooms in the condo, $Beds_{it}$, and dummies for the number of full bathrooms in the condo, $FullBaths_{it}$, plus a dummy ($\mathbb{1}\{HalfBath_{it}\}$) for whether the condo has a half bathroom. I also include for neighborhood fixed effects ($Nbrhood_i$), to control for fixed neighborhood characteristics that influence prices, and year/month fixed effects (δ_t) to control for time-specific housing demand or supply shocks experienced across the entire city.

Panel A in Figure 3 shows the distribution of condominiums by neighborhood. But-tressing the claim that a condo hedonic price analysis is indicative of rental responses, a comparison between Panels A and B in Figure 3 shows at least heavy spatial overlap between where condos and rent-controlled apartments are. Table 2 shows selected descriptive statistics of the condominiums. There are important differences between mean and median characteristic measures, namely that the average condo is more expensive, older, and further from a shuttle stop than the median. This in particular means that outliers in prices might bias the coefficients upward. I will thus also run an IV quantile regression to estimate the

effect at the median, subject to the caveat that reliable inference after first stage model selection via LASSO is unproven. Thus, the quantile regression estimates will largely be reported to given a sense of the bias from outliers rather than as a replacement for the standard 2SLS.

4.1 Hedonic Results

Table 3 shows robust evidence that the shuttle stops increased nearby prices. I estimate the coefficient for first the indicator for being within a half mile of a shuttle stop ($\mathbb{1}\{\text{Shuttle}\}$) and also the continuous shuttle measure for distance from the outer edge of a two kilometer radius around each stop (Shuttle2km_{it}). As explained above, this formulation is designed to ensure that the sign on the coefficient is positive in distance if the price effect is positive and vice versa, and can be interpreted as the price change for each kilometer a condo is closer to a shuttle stop. In keeping with concerns explained above about how landlords of “historical” buildings have reasons to behave differently than those with younger buildings, I also break out the results for whether they do or do not include “historical” condos.

Column (1) presents the point estimates from estimate Equation (4), but I am implicitly assuming here that residents do not value access to the eligible bus zones themselves. I relax that assumption in Column (2) by including additional characteristics in my exogenous controls that might covary with the instrument, such as the Easting and Northing measures for the condo (see the Appendix D for more information on Easting and Northing); number of nearby eligible bus zones; the average length of nearby bus zones; and the fraction of nearby bus zones have a shelter. The magnitude and precision both increase after adding these variables, likely because the specification in Column (1) cannot distinguish well between bus zone amenities and shuttle amenities. Since bus zones are relatively more plentiful and only serve the immediate city limits of San Francisco, they almost surely have lower amenity values than the shuttles. The omitted variable bias would push the shuttle coefficient downward.

It is further possible that the results are being confounded by time-varying changes in commuter demand for the other forms of transit (BART, Caltrain, etc.) since there is strong

geographic overlap in where shuttle stops and transit colocate. Surges in employees looking to work in Silicon Valley while living in the city would also cause demand increases in transit-proximate housing. This might mean that since the instrument is sensitive to the location of other transit options, the IV estimates in particular might pick up the influence of coincident demand shocks to these other transit forms rather than the actual price impact of the shuttles. Column (3) presents the results from including interactions between year fixed effects and indicators for whether the condo is within a half mile of a BART stop; a Caltrain stop; a north-south thoroughfare; a MUNI light rail stop; or the Central Business District. I refer to these hereafter as corridor-by-year or transit corridor-by-year fixed effects.

In every OLS specification in both samples, the reported coefficient is greater than the 2.2 percent benchmark established above, ranging from 3.6 percent to 5.5 percent for the indicator and 5.7 percent to 6.9 percent for the continuous measure, and all are significant at at least the 10 percent level. As half a mile is about 80 percent of a kilometer, it is unsurprising that the half-mile estimates are about 67–79 percent of the continuous distance measure. Given that the average condo is valued near \$1,000,000 (2013 dollars, Table 2), the in-levels change easily exceeds the \$20,000 relocation payment benchmark.

Across both samples, the point estimates for both OLS and IV specifications increase as the model gets more saturated for both shuttle measures. This bears out the concern that omitting bus zone characteristics leads to downward bias. Focusing on the sample excluding the historical condos, I present the second-stage IV estimates in the second panel of Table 3.²² My first stage F-statistics, calculated according to the method of Montiel-Olea and Pflueger (2013), show that I cannot build a strong first stage for my shuttle indicator measure, but can at least at the 10 percent level for my continuous shuttle measure (*Shuttle2km_{it}*) so I focus my discussion on just the continuous measure.

²²I reserve discussion about the first stages for the AppendixD.4, because the first-stage selection is not integral to this paper’s thesis, and because inference after LASSO is not possible in this setting in any case. The post-LASSO first stage is mostly important in terms of how strongly it predicts the endogenous variable, and whether the variables selected qualitatively support my instrument construction strategy so it is more appropriately dealt with in this Appendix which covers the instrument construction process more generally.

The three IV results offer collectively robust evidence that the shuttles had a positive price effect, and that condos closer to the stops had substantially more price appreciation. Even under very generous assumptions on the relationship between condo price appreciation and rent appreciation—that is, that rent appreciation only occurs at one-half or one-third the rate of condo price appreciation, a new shuttle stop for a building or condo previously one kilometer away from a stop would cause rental appreciation that would well exceed even the highest annual rent increase cap of 2.2 percent. Column (3) suggests that the appreciation could be $8.3 \times$ that amount (18.2 percent versus 2.2 percent). While reliable inferences or interpretation cannot be done with the indicator, the point estimates given more or less agree that exposure within a half mile radius would cause appreciation well above the annual cap.

The last panel in Table 3 displays the same series of results now generated by including the historical condos. The OLS point estimates shrink very slightly, but the IV point estimates are larger in the expanded sample. In the most saturated model (Column (12)), every kilometer closer you draw to a shuttle stop increases your condo price by 18.4 percent, conditional on being within two kilometers. This suggests that the shuttle stops were selected to near particularly less expensive historical buildings, since the magnitude of the corrected endogeneity bias is much larger than in the restricted sample.

These are strong results, and may in fact understate the equilibrium impact of shuttle amenity exposure because buyers and sellers may need additional time to respond to changes in shuttle stop placement. Appendix A includes Table A1, which shows the same specifications estimated in Table 3 but with the shuttle measures (and their instruments) lagged back six months. Comparing OLS and IV results from the six-month lag and the contemporaneous effects estimates shows very few meaningful differences. On balance, the lagged results yield slightly lower positive price effects, but not enough of a difference to believe that a better estimate of price effects require significant empirical lags.²³ I will continue to present additional results with the shuttle measure lagged back six months, but the general

²³Since both shuttle measures and all values for the instruments are zero before September 2004 and July 2004 (respectively), estimating the lagged results in this case does not require any observation loss.

pattern is that there is no particular reason to favor a lagged model over a contemporaneous effects model.

I now perform some brief exercises to show that the results are generally robust to simple sample and estimation perturbations. One concern with these results might be that since the mean sales price exceeds the median by almost 20 percent, the OLS and IV results may be upwardly biased due to right-tail outliers. Table 4 presents quantile regression results at the 25th, 50th, and 75th percentiles to examine which market segments were most impacted by the shuttles. The quantile regressions show that by far the strongest shuttle effects were in the lower half of the condo price distribution. Focusing chiefly on the contemporaneous continuous shuttle measure, shuttle-proximate condominiums at the 25th percentile in price saw a 7.6 percent increase in price and those at the median saw a 6.9 percent price increase in the most enriched specification. Those in the 75th percentile saw the smallest price increase, about 5.3 percent in the most enriched specification.²⁴ Expanding the sample to include historical condos slightly decreases the point estimates, as it does in Table 3, but does not otherwise change the fundamental conclusion that the shuttles raised prices by more than max annual rent increase cap and that the steepest price increases were observed in condos valued in the bottom quartile.

My next robustness check examines the valid instrument assumption. This is an important assumption to test because it is admittedly difficult to ensure even by controlling for all observable bus zone and transit characteristics that the instrument is relying only on variation exogenous to prices. Amenities generated by public transit and topographical features of the city can come from complex interactions based on proximity and where these networks form nodes. If the LASSO algorithm populates the first stage with regressors that also influence price, even while directly controlling for basic forms of transit exposure

²⁴Quantile IV estimates are not reported because the consistency of the estimates under a LASSO selection approach for IV quantile regressions has not been demonstrated. Since the direction of the bias is clearly downward, the quantile results help establish the lower bound. As all of these results presented in Table 4 have point estimates above 2.2 percent, it is reasonable to infer that a valid quantile IV approach would yield results even higher above the 2.2 percent benchmark.

and geographic position, then second-stage estimates (and inference on these estimates) are unreliable.

I address this concern via the method of Conley, Hansen, and Rossi (2012), which the authors refer to as the “plausibly exogenous” method. Borrowing their formulation, after partialling out all exogenous variables, the simultaneous equations problem two-stage least squares solves is

$$Y = X\beta + Z\gamma + \epsilon \quad (5)$$

$$X = Z\Pi + \nu, \quad (6)$$

where Y is the ultimate outcome of interest; X is the endogenous variable; and Z represents the first-stage instruments. In the special case that β and γ are scalars, the 2SLS estimate of β is asymptotically equal to

$$\hat{\beta} = (Z'X)^{-1}Z'Y \xrightarrow{p} \beta + \frac{\gamma}{\Pi} \quad (7)$$

$$Var(\hat{\beta}) = (Z'X)^{-1}Z'\epsilon\epsilon'Z(Z'X)^{-1} \xrightarrow{p} Var(\beta) + (Z'X)^{-1}Z'\Omega_\gamma Z(Z'X)^{-1}. \quad (8)$$

This highlights the importance of ensuring that the first stage is exogenous to prices, because both the point estimates and standard errors would be biased when $\gamma \neq 0$. It also underscores the importance of building a strong first stage, because small values of Π can create a divide-by-zero problem that can cause the bias on the 2SLS estimates to blow up. I use the local-to-zero (LTZ) approach to test how violations of the exogeneity assumption bias my results in Table 3, mostly because of its transparency and ease of implementation. In the LTZ approach, the user supplies the prior distribution of γ and then $\hat{\beta}_{2SLS}$ and $Var(\hat{\beta}_{2SLS})$ are recalculated.

Drawing in part on Conley, Hansen, and Rossi (2012), I assume $\gamma \sim N(0, \Omega_\gamma)$. In their suggested specifications, they usually retain the $\gamma = 0$ approach but give little guidance for

specifying the variance when there are multiple instruments. I thus use $V(\hat{\gamma})$ generated from estimating Equation (5), which will give a sense of how much downward bias there is in the confidence intervals around the estimates in Table 3.

The results are presented in Table 5. The standard errors grow by a relatively modest amount of 13.6–17.9 percent for the restricted sample, but by a more significant 21.1–27.5 percent for the expanded sample. This is unsurprising since the instrument is weaker in the expanded sample, and the bias in the standard errors shrinks with instrument strength. The key conclusion from this exercise is that the significance of the results remains robust to even a more relaxed assumption about the instruments’ exogeneity. Thus, shuttle stop placement is almost surely a sufficiently large shock to local prices to induce the marginal (and supramarginal) landlord to pursue changing his or her housing supply via nonprice mechanisms.

5 Estimating Controlled Landlord Supply Responses

Having established above that the shuttles provide a large price shock that landlords cannot fully capitalize into rents, I now investigate how landlords change their housing supply in response. I do this via five outcomes: OMI evictions, Ellis Act evictions, up conversions, down conversions, and repair permit filings.

The panel IV fixed effects model is employed to estimate the impact of price shocks on OMIs and Ellis Act evictions at the individual building level. This equation takes the form of

$$\begin{aligned} Outcome_{iut} = & \zeta_0 + \zeta_1 Shuttle2km_{it} + \zeta_2 LightRail_{it} + \zeta_3 Policies_{it} \\ & + \delta_t YYMM_t + \delta_{ut} YYMM_t \times UnitGroup_u + \delta_i Building_i + \theta_{it} OtherTransit_i \times t + \epsilon_{iut}, \end{aligned} \quad (9)$$

where $Outcome_{iut}$ is an indicator for whether OMI or Ellis Act evictions occurred in building i in month t . $Shuttle2km_{it}$ is defined the same as in Section 3. $LightRail_{it}$ differs from Section

4 in that it only has measures for distances to the MUNI Metro and interactions with rent control status—all other distance measures are dropped in this specification because they are not time-varying. $OtherTransit_{it} \times t$ interacts the transit half-mile dummies with a linear time trend, to control for housing supply and demand trends in transit-adjacent corridors. $Policies_{it}$ are city policies that applied to only buildings with specific unit counts, or in the case of relocation payments, captures how costs may have escalated nonlinearly with the number of units. $YYMM_t$ is a vector of year/month fixed effects, controlling for month-specific shocks to eviction rates. Note that this also nonparametrically controls for policies affecting only controlled buildings. $Building_i$ is a vector of building-specific fixed effects to control for all time-invariant building characteristics. Standard errors are clustered at the neighborhood level, under the assumption that housing demand and supply is often neighborhood-specific but usually independent across them.²⁵

Lastly, $YYMM_t \times UnitGroup_u$ controls for monthly shocks to buildings grouped into three unit group categories u : two-unit (“small”) buildings, three-to-six-unit (“midsized”) buildings, and seven-plus-unit (“large”) buildings.²⁶ These categories reflect that the city allowed three-to-six-unit landlords to condo convert via an annual lottery until 2013. Two-unit landlords could exit via the lottery as well as an unlimited lottery bypass to condo convert upon satisfying certain ownership requirements even after the lottery’s suspension.²⁷ Reconstructing and parameterizing the complex lottery rules is very difficult, because many eligibility rules hinged on complex ownership and length-of-stay requirements (Asquith 2019b). Nonetheless, controlling for landlord lottery access is important for this study, because a landlord’s private assessment of their exit chances via the lottery very likely influenced their decision to use other means to change their housing supply. Thus, these unit group-specific time shocks capture common demand shocks to large- versus medium- versus small-sized

²⁵While this is something of a strong assumption, it is not unreasonable to assume for tractability that buildings are much more independent across neighborhoods than not, while being possibly highly correlated within the same neighborhood.

²⁶To avoid endogeneity with changes in housing supply, these categories are determined from the initial unit counts in the first study period, July 2003.

²⁷More details are in the Appendix. B.3

buildings, but also any changes in lottery (or other) rules impacting exit chances via this mechanism.

The instruments used in the first stage for Equation (9) is created via the same process described in Section 3.1, and is populated as before via the LASSO <https://www.overleaf.com/project/5d41dee> algorithm.

6 Results

I begin by examining supply effects via five model specifications. The first uses the base specification given in Equation (9). Specification two incorporates neighborhood-by-year fixed effects which recognizes that renters often target specific neighborhoods rather than locations, so that neighborhood demand shocks may confound identification of supply responses via shuttle stop placement. In specification three, I add corridor-by-year fixed effects similar to those added in Section 4 to further control for any year-specific demand or supply shocks in these transit- and business-proximate corridors. Specification four adds neighborhood monthly time trends to reflect that long-term neighborhood trends might also be dictating housing demand and supply. The last specification substitutes neighborhood monthly time trends with neighborhood dummies interacted with *GOOG*, which captures how changes in Google’s stock price might have spurred housing demand across neighborhoods.

I start by estimating no-fault responses with and without historical controlled apartments. Table 6 shows the OLS and IV fixed effects results from estimating Equation (9) for Ellis Act evictions and OMIs. Excluding historical residences, the point estimates are small, insignificant, and unstable in sign. Going from OLS to IV estimates almost uniformly moves the point estimates rightward along the real line, in line with the idea from Section 4.1 that removing endogeneity bias will uncover stronger landlord responses, because the “true” unrealized amenity capitalization is higher than the naive estimate. In this case, the IV

estimates show that for both Ellis Act and OMI evictions, landlords are *more* likely to use these when prices go up, although none of the results are statistically significant.

The first stage here is quite strong across all specifications. These regressions are not only nonparametrically controlling for not only all the bus zone characteristics that troubled the results in Section 4.1, but also the data here are rich enough to factor out other neighborhood-level confounds. These results help validate the general approach because the residual variation is sufficiently predictive to confidently reject the weak instrument null hypothesis.

Adding in the historical buildings validates the hypothesis that older landlords may be especially responsive to price increases, likely because of capital improvements restrictions discussed in Appendix Section B.4. In the IV results in the bottom right panel, all the point estimates increase markedly, particularly in the three most saturated models. In that panel's Column (5), landlords are found to be 0.057 percentage points more likely to perform an OMI when prices increase.

This is an objectively small number, but it represents a 158 percent increase over the average monthly OMI rate of 0.036 percent. The pattern of the IV estimates is that they almost always double or more the prevailing average monthly rate. This in turn reflects the very high estimated value of the shuttle stops in Section 4.1, and imply that the large gap between the value of stop-proximity and the rent cap leaves landlords with strong incentives to change their supply response. I will argue for the rest of this section that while the magnitudes are almost surely overstating the true effect, the sign on the coefficient is what indicates whether controlled landlords are increasing or decreasing their supply. Nonetheless, I will report only IV results for the remainder of this section because they are based off a stronger implied price effect, and thus are more informative of how landlords respond to a large demand shock that cannot be fully priced into rents under the annual cap.

Analyzing the entire market may mask substantial supply response heterogeneity across at least two dimensions. The first dimension is along initial, preshuttle rent levels. The

results in Section 4.1 indicate both that the shuttles disproportionately targeted areas with relatively less expensive condos and that the greatest price appreciation occurred among the cheaper condominiums.

The second dimension is along building unit counts, because of the condominium conversion lottery that ran in San Francisco from 1997 to 2013 discussed in Section 1.²⁸ These divisions also map at least loosely onto differing abilities by larger versus smaller landlords to pay court fees, buyout tenants, and utilize property managers to collect rents and control tenants that may be relevant for outcomes under rent control.

6.1 Results by Neighborhood Price Quartiles and Building Size

I reutilize the condominium transaction data to identify which neighborhoods were in the bottom three quartiles for transaction prices between July 2002 and July 2004, the month before Google went public. I was able to estimate averages for 66 neighborhoods, leaving 49 neighborhoods in my bottom three quartiles. Table 7 uses buildings in just these 49 neighborhoods, which are where one would expect, on average, the strongest landlord responses might be found. It also includes the full sample (historic and nonhistoric buildings, as defined above), because the evidence in Table 6 shows stronger responses across the full sample than the restricted sample. The primary purpose of this paper is to establish the general pattern of behavior (increase or decrease of housing supply), I now proceed by attempting to identify where the strongest responses might be, describe the pattern identified, and then discuss the findings' significance in the context of how rent control works in San Francisco.

The top panel has IV results covering all in-sample buildings, first for Ellis Act evictions and then for OMIs. As expected, restricting the sample to the bottom three condo price quartiles improved the results' precision. Now, Ellis eviction probability increases by 0.025 percentage points in specification three, significant at the 10 percent level. OMI evictions are

²⁸See Asquith (2019a,b) for more information.

shown to increase by 0.067–0.081 percentage points across specifications two through five, with estimates statistically significant at the 5 percent level. For context on how meaningfully landlords’ responses changed after the shuttles were put in, the smallest significant coefficient, 0.067 percent represents a 172 percent increase in the OMI eviction probability when a shuttle stop changes from two kilometers (or more) away to one kilometer away. This is again almost surely too high, but the important point is that landlords in the neighborhoods where prices most likely experienced the greatest increases when stops were put in *decreased* their medium-term housing supply in response. This is a very substantive finding, because it implies that the supply of controlled housing in San Francisco may not be upward sloping in price.

Exploring effects across submarkets by size, the second panel thus looks at buildings with 7 or more units. Ellis Act eviction probabilities appear to very substantially decrease once a building comes within range of a shuttle stop: the estimates in specifications two through five range from -0.104 percent to -0.067 percent, and all are at least significant at the 10 percent level. All the shuttle coefficient estimates in the OMI regressions are positive, but outside of specification one, none are statistically significant. This means that at least for the large buildings, the best evidence is that landlords become less likely to temporarily withdraw their housing supply as prices increase.

This is perhaps a surprising result, because if adverse selection is a problem in the controlled market, these evictions represent some of the only means landlords of large buildings have to reduce their housing supply. This evidence suggests either that landlords of larger buildings are less afraid of adverse selection, or that they instead turn to other means of increasing their returns, such as at-fault evictions or buyouts. Both factors may be at play. These landlords can reduce adverse selection more easily by stronger tenant screens, and if they do choose a long-staying tenant, these landlords have the deeper pockets needed to get rid of that tenant either via a trumped up at-fault eviction or a buyout.

Smaller landlords may thus respond differently, so the third panel in Table 7 has results on

three-to-six-unit buildings, which are eligible to condo convert only via the lottery. Midsized landlords are more likely to use OMIs, with statistically significant results in specifications two through five. Landlords of these buildings are 0.099–0.112 percent more likely to perform an OMI eviction in a given month for each one kilometer decrease in the distance from their building to a shuttle stop, given that they are within two kilometers.

The bottom panel in Table 7 has results on two-unit buildings. The last panel of Table 7 shows robust evidence that Ellis Act evictions increase in two-unit buildings after coming within two kilometers of a shuttle stop. This finding is significant across specifications three through five, and ranges between 0.033 to 0.046 percentage points. The coefficient in the OMI eviction regressions is positive, but not statistically significant in any regression.

This is also a remarkable finding because performing either an OMI or an Ellis Act eviction effectively renders the odds the landlord wins the lottery close to zero.²⁹ Two-unit landlords can use Ellis Act evictions as a means to eventually exiting via the lottery bypass, but three-to-six-unit landlords are effectively closing the door to permanently exiting rent control. Even for small building-landlords, an Ellis Act eviction still forecloses on their ability to exit via the lottery, in those cases where for some reason they cannot meet the bypass’ requirements. Thus, at least some landlords are choosing to withdraw units from the market for at least three to five years *after* a price increase.

These are strong results. In Appendix A, I present and discuss two robustness checks on the findings in Table 7, both of which further underscore the findings above. Overall, the results are generally clear: bigger controlled buildings become less likely to withdraw housing units in response to housing price increases. Smaller landlords, by contrast, become more likely to do so, either via OMIs (three-to-six-unit landlords) or the Ellis Act (two-unit landlords), even though doing so lessens their ability to permanently exit rent control by converting to condos and fully capitalizing on the new transit amenity.

²⁹See Asquith (2019b) for more information.

6.2 Results on Conversions and Repairs

I now turn to analyzing whether landlords attempt to change supply along either the intensive margin—that is, increasing or decreasing the number of units in their building, or increase their willingness to maintain their units.

It is clear from Section 6.1 that the first-stage estimation and the precision of the second stage results significantly improves once both neighborhood-by-year and corridor-by-year fixed effects are added. For the remainder of the paper, I will include these fixed effects as part of the base specification. I thus estimate the up and down conversions and repairs model as

$$\begin{aligned}
Outcome_{it} = & \zeta_0 + \zeta_1 Shuttle2km_{it} + \zeta_2 (Shuttle2km_{it} \times \mathbb{1}\{RentControl_i\}) \\
& + \zeta_3 (Policies_t^{all} \times \mathbb{1}\{RentControl_i\}) + \zeta_4 (Policies_t^{spec} \times Units_i \times \mathbb{1}\{RentControl_i\}) \\
& + \zeta_5 LightRail_{it} + \zeta_6 (LightRail_{it} \times \mathbb{1}\{RentControl_i\}) + \theta_{it} OtherTransit_i \times t \quad (10) \\
& + \delta_t YYMM_t + \delta_{ut} YYMM_t \times UnitGroup_u + \delta_i Building_i + \delta_{year, nbrhd} (Year_t \times Nbrhd_i) \\
& + \delta_{year, transit} (OtherTransit_i \times Year_t) + \epsilon_{it}.
\end{aligned}$$

This adds a new key variable, $Shuttle2km_{it} \times \mathbb{1}\{RentControl_i\}$, which reflects the change in conversion or repair filing probability in controlled buildings vis-à-vis the change in uncontrolled buildings. Since the sample includes uncontrolled and controlled buildings, the regression now includes $Policies_t^{all}$, which refers to policies applied to all controlled buildings.³⁰ $Policies_t^{spec}$ refers to policies that are unit-count specific. These are interacted with rent control status, which captures how these policies were intended to change controlled landlords' behavior, and are analogous to $Policies_{it}$ in Equation (9).³¹ While these policies are not usually targeted to change landlords' willingness to do up or down conversions,

³⁰Details can be found in Online Appendix Table A2.

³¹For the imposition of relocation payments on owner move-in evictions in 2006 and the extension of relocation payments to all Ellis'd tenants in 2005 described in Online Appendix Table A3, I create dummy variables also interacted with rent control status for both policy changes.

landlords may change their likelihood to do so if withdrawing controlled units becomes more expensive. The empirical assumption here is that eviction policy changes are endogenous with respect to the city’s overall eviction rate, but exogenous to changes in any individual building or neighborhood.

The interaction between rent control and the shuttles ($Shuttle2km_{it} \times \mathbb{1}\{RentControl_i\}$) could also be endogenous. Stops may have favored controlled apartments because technology employees are willing to pay the premium, anticipating that high job security necessitates prolonged tenures. Figure 3b shows the fraction rent controlled per neighborhood, and a comparison with Figures 1 and 2 shows that areas with a high rent control fraction are both well-covered by the shuttles and have many bus zones. Thus, the LASSO-selected instrument is augmented by variable interactions with the rent control dummy.

Table 8 shows estimation results for just three specifications: the base model given in Equation (10); add neighborhood-specific time trends; substitute neighborhood-specific time trends with neighborhood-specific demand shocks from *GOOG*. As stated above, the controlled sample excludes historic buildings so as to better compare effects between buildings of a more similar vintage.³²

Intriguingly, the first panel showing results for all up conversions seem to decrease when the building comes within two kilometer of a shuttle stop. Rent-controlled landlords are also less likely to up-convert their building after a localized demand shock, but the decrease is higher for uncontrolled than for controlled landlords. The second panel shows that this decrease among uncontrolled buildings appears to be driven by large buildings. Otherwise, there seems to be little significant change in conversion permits filed in response to coming within range of a shuttle stop.

Similarly, Table 9 shows that repair permit filing does not appear to significantly change with the placement of a shuttle stop. The general pattern is that for large landlords (seven

³²Results on up and down conversions and repair permit filings that include the historic controlled buildings generally arrive at the same results as above, but often report a statistically significant margin because the differences between controlled and uncontrolled landlords are sharpened. These results are available upon request.

or more units), the coefficient estimates in the uncontrolled units are unstable in sign, but in the controlled units the coefficients are uniformly negative but not significant. Landlords of these buildings have the highest permit filing rate (0.11 percent of these landlords file at least one permit per month), so even the smallest marginal effect (specification two) represents a 72.2 percent decrease. Among three-to-six-unit landlords, the point estimates are too small and unstable in sign and with confidence intervals too wide to draw any sort of conclusions.

By contrast, there is a robust finding that two-unit controlled buildings became much more likely to file repair permits. These estimates range from 0.053 to 0.066 percentage points and are statistically significant at the 1 percent level across all three specifications. One interpretation is that this is a sign that at least some controlled landlords hope to increase their returns by maintaining their housing when the unit next becomes vacant. The other interpretation is that this may be driven by controlled landlords seeking to sell a building after an Ellis Act eviction or spruce it up for sale after condo converting via the lottery bypass. I explore this question more below.

6.3 Additional Evidence from At-Fault Evictions

The strongest evidence so far has been that smaller landlords are more likely to at least temporarily withdraw either the entire building or a single unit in response to a local demand shock. However, there is also evidence that these same buildings are more likely to have their landlords file repair permits in response to the same shock. One explanation is that these results reflect heterogeneity among landlords: some withdraw because of adverse selection fears, whereas perhaps others who have more confidence in their ability to screen and turnover tenants, and thus increase their maintenance. Another explanation is that the two results are complementary: small landlords withdraw their units, then spruce them up in preparation for a sale as vacant units.

This paper's last analysis will use at-fault evictions to attempt to give additional evidence on how controlled landlords are responding to market pressures. Reiterating briefly the facts

covered in Section 1, rent control in California cities comes with a vacancy bonus, because rents in between tenancies are not regulated. This gives landlords strong incentives to turn over tenants when prices rise, and when they cannot be bought out, an at-fault eviction may turn out a recalcitrant tenant. In comparison to no-fault evictions, at-fault evictions return the unit to the controlled market and can be understood both as a rent-control workaround and also a sign that the controlled market is incentivizing landlords to keep supplying units to it (Asquith 2019a). If there is evidence that at-fault evictions increase, especially among two-unit landlords, then that is evidence that there are likely heterogeneous landlord responses within building groups: some choose to exit, some choose to try to take advantage of higher prices. If instead at-fault evictions are unchanged or fall, then it is likely that permit filing is increasing in this group to try to make the post-Ellis vacant building more attractive for sale.

Table 10 presents the results of estimating Equation (10), but like with the repairs regressions, weighting the buildings by unit count. The strongest evidence in Columns (2) and (3) is that uncontrolled large landlords (seven or more units) increase their propensity to do an at-fault eviction, although these coefficients have p-values consistently a bit above even the 10 percent significance level so the null hypothesis of a zero coefficient cannot be rejected. That said, the interaction terms for these landlords are uniformly negative and significant at the 10 percent level across all three specifications. Interestingly, the point values on the interaction terms is very close to the same absolute size as for the shuttle term itself, leading the marginal effect to be very close to zero. In fact, the t-statistics for the marginal effect in Column(3) is itself just 0.01. The most reasonable interpretation here is that unlike their uncontrolled peers, at-fault evictions do not increase in large, controlled buildings.

Among two-unit landlords, Column (3) has evidence that uncontrolled landlords decreased their at-fault eviction rate by 0.109 percentage points (significant at the 10 percent level), but the interaction term, while positive, is not significant and not larger in absolute terms than the coefficient on the uninteracted shuttle term. Thus, the marginal effect is

negative and not significant for all three specifications. The most reasonable inference here is that controlled landlords do not increase their tenant turnover rate after a local demand shock. If anything, these landlords may be less inclined to do so, although the probability decrease is possibly lower in controlled than in uncontrolled buildings. This is commensurate with the interpretation that landlords of the smallest buildings are withdrawing their housing supply, at least over the medium term. There is no evidence that they take meaningful steps to increase supply, either by up-converting (all results were negative but not significant) or increasing turnover via at-fault evictions. Taken altogether, the fact that this group is most likely to both use the Ellis Act and increase their filing of repair permits is best understood as their landlords preparing to sell them as vacant buildings.

6.4 Supply Results Summary

The point estimates across the specifications in Tables 7–10 are generally too high to be plausibly used as part of a conventional policy analysis. Nonetheless, the clear direction of these estimates is still very informative, and I can use these results instead to create a rough range of expected outcomes. The IV results can be thought to represent what would occur from a strong local demand shock. By contrast, the OLS results in Section 4.1 generally have point estimates only modestly above the 2.2 percent maximum annual allowable rent increase, and thus an OLS regression of Equations (9) and (10) can be viewed as what would occur under a smaller local demand shock. Taken together, the two results can be viewed as a range of outcomes, where the IV results represent what would occur if a very large fraction of the shuttle amenity’s valuation cannot be capitalized into rents, and the OLS results represent what would occur if only a relatively minimal fraction of the valuation cannot be capitalized.

Table 11 shows the OLS estimates of the shuttles’ effect on controlled buildings for just the specification with neighborhood-by-year, corridor-by-year, and neighborhood-by-GOOG fixed effects. This corresponds to column (5) in Table 7 for Ellis and OMI evictions; column

(3) in Table 8 for up and down conversion; column (3) in Table 9 for repair permit filings; and column (3) in Table 10 for at-fault evictions. The assumption here is that the results in this table reflect the downward bias in the hedonic effect, and do not bias the results through any other mechanism. This is unlikely to be completely true, so this exercise is done merely as an attempt to contrast how landlords respond under two different kinds of shocks, and is not an attempt to formally plot the supply responses in price.

Most of the OLS results are not significant, but the general pattern is for the point estimates to move consistently across building groups either right or left along the real line. Ellis Acts, OMIs, up conversions, and repair permit filings decrease; at-fault rates increase; down conversions do not exhibit a clear pattern. There are three statistically significant results. Up conversion filings increase by 71 percent (0.00005/0.00007) across all buildings, which appears to be largely driven by a 133 percent (0.00012/0.00008) increase in filings from mid-sized buildings. Lastly, repair permit filings increase by 64 percent in 2 unit buildings.

Compared to the results in Tables 7–10, the OLS results would lead one to think that the rent-controlled system is not working too badly at all. Evictions do not increase, plus there are increases in the number of units and building maintenance are all positive developments. The IV results, by contrast, speak to a worse set of outcomes, namely that there are increases in no-fault evictions among smaller buildings that at least temporarily decrease the housing supply, and which is not offset at all by increases in the physical supply of units via up conversions.

There are a few policy-relevant conclusions that can be drawn from these results. The contrast between the IV and OLS results suggests (but does not dispositively prove) that the annual allowable rent increase cap plays a critical role in how well the market functions. San Francisco’s rules are unusually stingy, mandating that landlords lose money on their rents in real terms by setting the cap at 60 percent the regional CPI rate. Even allowing the cap to match the rate of inflation would probably lead to better supply outcomes, while minimally reducing tenants’ benefits.

This is a somewhat modest suggestion, because it does not fundamentally change the architecture of rent control in San Francisco. Nonetheless, even increasing the annual rent cap may be politically infeasible. Tenants and their advocates may simply not be willing to countenance allowing landlords to effectively index rents to inflation. Thus, for those whose chief interest is in maintaining the tenants' status quo while looking to mitigate the worst effects of landlords attempts to circumvent controls, clamping down on no-fault evictions may seem like the path of least resistance.

As mentioned when describing Equations (9) and (10), San Francisco in fact attempted exactly this via a series of restrictions and costs no-fault evictions. Table 12 shows how these policies impacted controlled landlords' willingness to file a repair permit or perform an at-fault eviction. Rule 1 made it somewhat easier for controlled landlords to perform an at-fault eviction if the tenant had violated a lease change that had been mandated by law. Unsurprisingly, this raised the monthly at-fault eviction rate by 2.3 percentage points and by 2.6 percentage points among large buildings.

Rule 2 forbade landlords from performing an OMI eviction if there was a minor in the unit during the school year. Rule 3 mandated relocation payments to tenants evicted under the Ellis Act, and boosted the monthly at-fault eviction rate by 4.1 percentage points. Rule 4 extended the relocation payment mandate to other no-fault evictions (including OMIs), which appears to have had little effect. Rule 5 applied to uncontrolled landlords and allowed them to only give 30 days notice, as opposed to 60 days, before ending the lease.

Discussing additional responses by unit count, the largest landlords appear to be most responsive. While they were the least likely to use a no-fault or at-fault eviction, the results in Tables 7 and 10 may mislead the reader into thinking that large building landlords are not responsive to local price trends. This paper has no results on tenant buyouts, which are almost surely most used by large landlords and which may be absorbing most their response. That said, if they do want to turn over a tenant who will not accept a buyout agreement, they have greater resources to win the at-fault eviction in court. It is therefore less surprising that

large landlords became significantly more likely to engage in at-fault evictions when the city made no-fault evictions more difficult and expensive. In terms of Rule 5, lessening the time to terminating the lease lowers the at-fault eviction rate for large landlords, likely because making it less costly to terminate the lease, rather than evict the tenant, unsurprisingly leads to a lowering of the eviction rate. While the coefficients on the midsized and small landlords are smaller and not significant, the coefficients do generally retain the same sign for the smaller buildings. The conclusion here is that making it harder or more expensive to withdraw units from the controlled market does induce landlords, particularly landlords of larger buildings, to substitute into turning over tenants presumably to increase their returns by other means.

7 Conclusion

In this paper I tested how controlled landlords change their housing supply in response to localized demand shocks. Exploiting unique properties of San Francisco’s housing market—namely an identifiable, locally differentiated change in prices—I tested eviction and permitting responses. Landlords were found to engage in evictions that allowed them to directly or indirectly begin the process of switching to the uncontrolled market, even at the expense of decreasing their medium-term housing supply. There is no evidence that landlords attempt to increase (or decrease) their buildings’ unit count in response to a demand shock. Landlords of small buildings (two units) seem much more likely to file repair permits, but given that they also increase their filings to withdraw their units, it is just as likely they are engaging in maintenance as they are simply sprucing up vacant buildings for sale. Further, while I find no evidence that landlords increase their propensity to reset base rents higher via at-fault evictions in response to the price shock, I find evidence they do so when market withdrawal evictions become more expensive. In light of these findings, future research could revisit the distributional and housing market aspects of rent-control policy by accounting for

landlord's ability to shrink the size of the controlled market.

The bulk of the evidence suggests that the system poorly absorbs demand shocks. The linchpin of San Francisco's rent-control law is the annual rent increase cap of 60 percent of the regional CPI rate, which is low enough that landlords cannot use prices to respond to significant demand shocks. The IV results suggest that when rent control leaves landlords with large, unrealizable price gains, an increasing number of landlords choose to withdraw some or all of their units from the market, even at the cost of leaving them vacant. While this paper does not perform a welfare analysis comparing the results from changing different aspects of the rent control law, it does highlight the supply-side consequences of constraining landlords from being able to use prices to allocate their units.

References

- Akerlof, George A., 1970. The Market for “Lemons”: Quality Uncertainty and the Market Mechanism. *The Quarterly Journal of Economics* 84, 488-500.
- Arnott, Richard, 1995. Time for Revisionism on Rent Control? *Journal of Economic Perspectives* 9, 99-120.
- Arnott, Richard, and Masahiro Igarashi, 2000. Rent Control, Mismatch Costs, and Search Efficiency. *Regional Science and Urban Economics* 30, 249-288.
- Arnott, Richard and Elizaveta Shevyakhova, 2014. Tenancy Rent Control and Credible Commitment in Maintenance. *Regional Science and Urban Economics* 47, 42-55.
- Ault, Richard W., John D. Jackson, and Richard P. Saba, 1994. The Effect of Long-Term Rent Control on Tenant Mobility. *Journal of Urban Economics* 35, 140-158.
- Autor, David H., Christopher J. Palmer, Parag A. Pathak, 2014. Housing Market Spillovers: Evidence from the End of Rent Control in Cambridge, Massachusetts. *Journal of Political Economy* 122, 661-717.
- Autor, David H., Christopher J. Palmer, Parag A. Pathak, 2019. Ending Rent Control Reduced Crime in Cambridge. *AEA Papers and Proceedings* 109, 381-384.
- Basu, Kaushik and Patrick M. Emerson, 2000. The Economics of Tenancy Rent Control. *The Economic Journal* 110, 939-962.
- Basu, Kaushik and Patrick M. Emerson, 2003. Efficiency Pricing, Tenancy Rent Control, and Monopolistic Landlords. *Economica* 70, 223-232.
- Belloni, Alexandre, Daniel L. Chen, Victor Chernozhukov, and Christian Hansen. Sparse Models and Methods for Optimal Instruments with an Application to Eminent Domain. *Econometrica* 80, 2369-2429.
- Belloni, Alexandre, Victor Chernozhukov, Christian Hansen, and Damian Kozbur, 2016. Inference in High-Dimensional Panel Models with an Application to Gun Control. *Journal of Business & Economic Statistics* 34, 590-605.
- Butler, Alexander W., Larry Fauver, and Ioannis Spyridopoulos, 2019. “Local Economic Spillover Effects of Stock Market Listings,” *Journal of Financial and Quantitative Analysis* 54, 1025-1050.
- Chao, John C., Norman R. Swanson, Jerry Hausman, Whitney Newey, and Tiemen Woutersen, 2012. “Asymptotic Distribution of JIVE in a Heteroskedastic IV Regression with Many Instruments,” *Econometric Theory* 28, 42-86.
- Conley, Timothy G., Christian B. Hansen, and Peter E. Rossi, 2012. Plausibly Exogenous. *The Review of Economics and Statistics* 94, 260-272.
- Dai, Danielle and David Weinzimmer, 2014. Riding First Class: Impacts of Silicon Valley Shuttles on Commute & Residential Location Choice. Working Paper, UCB-ITS-WP-2014-01.
- Diamond, Rebecca, Timothy McQuade, and Franklin Qian, 2019a. “Who Pays for Rent Control? Heterogeneous Landlord Response to San Francisco’s Rent Control Expansion.” *AEA Papers and Proceedings* 109, 377-380.
- Diamond, Rebecca, Timothy McQuade, and Franklin Qian, 2019b. The Effects of Rent Control Expansion on Tenants, Landlords, and Inequality: Evidence from San Francisco. *American Economic Review*, forthcoming.

- Gyourko, Joseph and Peter Linneman, 1990. Rent Controls and Rental Housing Quality: A Note on the Effects of New York City's Old Controls. *Journal of Urban Economics* 27, 398-409.
- Häckner, Jonas and Sten Nyberg, 2000. Rent Control and Prices of Owner-Occupied Housing. *The Scandinavian Journal of Economics* 102, 311-324.
- Hubert, Franz, 1995. Contracting with Costly Tenants. *Regional Science and Urban Economics* 25, 631-654.
- Iwata, Shinichiro, 2002. The Japanese Tenant Protection Law and Asymmetric Information on Tenure Length. *The Journal of Housing Economics* 11, 125-151.
- Krol, Robert and Shirley Svorny, 2005. The Effect of Rent Control on Commute Times. *Journal of Urban Economics* 58, 421-436.
- Linneman, Peter, 1987. The Effect of Rent Control on the Distribution of Income Among New York City Renters. *Journal of Urban Economics* 22, 14-34.
- McFarlane, Alistair, 2003. Rent Stabilization and the Long-Run Supply of Housing. *Regional Science and Urban Economics* 33, 305-333.
- Miron, J. R., 1990. Security of Tenure, Costly Tenants and Rent Regulation. *Urban Studies* 27, 167-184.
- Munch, Jakob Roland and Michael Svarer, 2002. Rent Control and Tenancy Duration. *Journal of Urban Economics* 52, 542-560.
- Nagy, John, 1995. Increased Duration and Sample Attrition in New York City's Rent-Controlled Sector. *Journal of Urban Economics* 38, 127-137.
- Nagy, John, 1997. Do Vacancy Decontrol Provisions Undo Rent Control? *Journal of Urban Economics* 42, 64-78.
- Olea, J. L. M. and Pflueger, C., 2013. A Robust Test for Weak Instruments. *Journal of Business and Economic Statistics* 31, 358-369.
- Raess, Pascal and Thomas von Ungern-Sternberg, 2002. A Model of Regulation in the Rental Housing Market. *Regional Science and Urban Economics* 32, 475-500.
- Skelley, Chris, 1998. Rent Control and Complete Contract Equilibrium. *Regional Science and Urban Economics* 28, 711-743.
- Sims, David P., 2007. Out of Control: What Can We Learn from the End of Massachusetts Rent Control? *Journal of Urban Economics* 61, 129-151.
- Siodla, James, 2015. Razing San Francisco: The 1906 Disaster as a Natural Experiment in Urban Redevelopment. *Journal of Urban Economics* 89, 48-61.
- Svarer, Michael, Michael Rosholm, and Jakob Roland Munch, 2005. Rent Control and Unemployment Duration. *Journal of Public Economics* 89, 2165-2181.

Tables

TABLE 1
Key Characteristics of San Francisco's Housing Stock

	2003	2008	2013
Total Housing Units	377,182	389,787	405,021
Average Year Built	1935	1937	1938
<i>As Percent of Total Housing Stock</i>			
Rent-Controlled	55.5%	53.5%	52.0%
Within Half Mile of Shuttle Stop	0.0%	56.5%	59.7%
Condominiums	9.5%	12.1%	14.2%
Apartment Buildings	57.5%	56.2%	55.3%
Residences	29.7%	28.7%	27.7%

SOURCE: Author's calculations from Assessor's data.

NOTE: Table 1 presents averages of key characteristics of the housing stock in San Francisco for three selected years. The sample comprises all buildings in San Francisco in the years specified with a housing property code that indicates it is a condominium, house, apartment building, flat, or townhome, summed by year in Row 1. Condominiums are denoted as buildings in their own right, because they are recorded individually by the city even if they are in a multi-unit building. Average Year Built (Row 2) is the mean year built by property. Rows 4–9 give the fraction of housing units that are rent-controlled; the fraction that are within a half-mile of a shuttle stop as of June of that year; the fraction that are condominiums; apartment buildings; residences; and the fraction that have no-fault evictions, respectively.

TABLE 2
In-Sample Sold Condominium Characteristics ($N = 24,970$)
July 2003–December 2013

	Mean	Median
Sales Price	\$ 919,067	\$ 773,500
# of Baths	1.63	2
# of Beds	1.86	2
Sq. Ft.	1167.3	1096
Year Built	1981	1996
Distance to the Central Business District (km)	2.12	2.51
Distance to BART (km)	1.67	1.41
Distance to CalTrain (km)	2.74	2.54
Distance to MUNI (mi.)	0.97	0.65
Distance to Thoroughfare (mi.)	0.70	0.52
Distance to Any Company Shuttle (mi.) ^a	1.43	0.74
Within a Half-Mile of Shuttle ^a	0.55	1

SOURCE: Author's calculations from Assessor's data.

NOTE: Table 2 presents summary statistics on housing characteristics of condominiums sold in San Francisco between July 2003 and December 2013. The sample is defined as single-use, single-family condominiums with no more than 3.5 bathrooms and 4 bedrooms and a minimum of 291 square feet. "Within a Half-Mile of Shuttle" is coded 0/1 in the underlying data, so that the mean of "Within a Half-Mile of Shuttle" represents the fraction across the entire time period of shuttles that were within a half mile of a shuttle stop.

^a Restricted to the time period after the first shuttle stops were introduced (October 2004 onwards).

TABLE 3
Hedonic Price Regression on Condominium Sales
with and without Historic Condos: July 2003–December 2013

		$\mathbb{1}\{\text{Shuttle}\}$			Shuttle 2km		
Without Historic Condos ($N = 24,969$)							
OLS	<i>Shuttle Measure</i>	(1) 0.0391* (0.0199)	(2) 0.0505** (0.0213)	(3) 0.0546** (0.0230)	(4) 0.0583** (0.0265)	(5) 0.0661** (0.0278)	(6) 0.0692** (0.0291)
	Adj R^2	0.664	0.672	1.000	0.665	0.673	1.000
	F	509.88	695.82	276.03	480.15	612.44	386.02
IV	<i>Shuttle Measure</i>	0.0349 (0.0444)	0.1112** (0.0547)	0.1314** (0.0539)	0.0911** (0.0374)	0.1575*** (0.0453)	0.1815*** (0.0428)
	Hansen's J p-value	0.106	0.147	0.152	0.256	0.285	0.066
	First-stage F stat	13.81	10.53	8.85	16.41*	28.53**	21.52*
	5% Crit Value	27.71	28.43	30.11	27.71	27.53	28.55
	10% Crit Value	15.85	16.37	17.61	15.86	15.70	16.46
With Historic Condos ($N = 26,884$)							
OLS	<i>Shuttle Measure</i>	(1) 0.0364* (0.0193)	(2) 0.0479** (0.0205)	(3) 0.0512** (0.0217)	(4) 0.0571** (0.0261)	(5) 0.0655** (0.0271)	(6) 0.0677** (0.0284)
	Adj R^2	0.662	0.669	1.000	0.663	0.670	1.00
	F	406.45	.	491.09	398.96	.	544.01
IV	<i>Shuttle Measure</i>	0.0127 (0.0490)	0.0917** (0.0445)	0.1281*** (0.0460)	0.1035*** (0.0351)	0.1800*** (0.0375)	0.1843*** (0.0371)
	Hansen's J p-value	0.158	0.159	0.480	0.275	0.380	0.497
	First-stage F stat	10.86	8.35	7.48	15.20	17.75*	16.94*
	5% Crit Value	28.87	28.29	30.79	27.63	27.49	28.69
	10% Crit Value	16.71	16.28	18.11	15.79	15.69	16.57
Bus Zone Vars		N	Y	Y	N	Y	Y
Year-by-Transit FE's		N	N	Y	N	N	Y

SOURCE: Table 3 shows the impact of tech company shuttle stop proximity on condominiums' log sales price generated from Equation (4) in Column (1), with additional specifications reported in Columns (2) and (3).

NOTE: The first shuttle measure is within a half mile of a shuttle stop ($\mathbb{1}\{\text{Shuttle}\}$) or the condo's kilometer-measured distance from a 2km radius ring around the condo's closest shuttle stop. These coefficients can be interpreted as the impact of moving 1km closer to a shuttle stop, given that the condo is within 2km of a stop.

* p<0.10, ** p<0.05, *** p<0.01.

TABLE 4
Hedonic Price Regression on Condominium Sales: July 2003–December 2013
Quantile Regressions by 1st, 2nd, and 3rd Quartiles

	1st Quartile			2nd Quartile			3rd Quartile			
Without Historic Condos		(1)	(2)	(3)	(1)	(2)	(3)	(1)	(2)	(3)
	No Lag	0.0663*	0.0750*	0.0764*	0.0609***	0.0661***	0.0691**	0.0397*	0.0519***	0.0529**
		(0.0411)	(0.0424)	(0.0437)	(0.0232)	(0.0240)	(0.0295)	(0.0204)	(0.0186)	(0.0237)
	R ²	0.651	0.660	0.663	0.652	0.662	0.667	0.649	0.658	0.661
	Lag t – 6	0.0635	0.0768*	0.0856*	0.0595**	0.0652***	0.0774**	0.0358*	0.0498**	0.0537**
		(0.0403)	(0.0448)	(0.0448)	(0.0257)	(0.0249)	(0.0311)	(0.0195)	(0.0229)	(0.0274)
	R ²	0.651	0.660	0.662	0.652	0.662	0.666	0.649	0.658	0.661
With Historic Condos		(1)	(2)	(3)	(1)	(2)	(3)	(1)	(2)	(3)
	No Lag	0.0656*	0.0700*	0.0703	0.598***	0.0655***	0.0670***	0.0400**	0.0495**	0.0489**
		(0.0371)	(0.0411)	(0.0428)	(0.0215)	(0.0246)	(0.0254)	(0.0175)	(0.0196)	(0.0219)
	R ²	0.650	0.657	0.659	0.652	0.660	0.664	0.649	0.654	0.657
	Lag t – 6	0.0650*	0.0765*	0.0842**	0.0584**	0.0646**	0.0762***	0.0357*	0.0492**	0.0505*
		(0.0364)	(0.0415)	(0.0425)	(0.0245)	(0.0262)	(0.0295)	(0.0196)	(0.0237)	(0.0265)
	R ²	0.649	0.659	0.661	0.651	0.661	0.665	0.648	0.657	0.659
Bus Zone Vars				N	Y	Y	N	Y	Y	
Year-by-Transit FE's				N	N	Y	N	Y	Y	

SOURCE: Table 4 shows the impact of tech company shuttle stop proximity on the first, second, and third condominium log sales price generated from Equation (4) (Column (1)), with additional specifications reported in Columns (2) and (3).

NOTE: Results are reported both for contemporaneous shuttle exposure, and shuttle exposure lagged back 6 months. R^2 is computed as the correlation between the fitted values and the dependent variable. These coefficients can be interpreted as the impact of moving 1km closer to a shuttle stop, given that the condo is within 2km of a stop.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE 5
Hedonic Price Regression on Condominium Sales: July 2003–December 2013
Conley, Hansen, and Rossi (2012) Plausibly Exogenous LTZ Method Results

Assuming $\gamma \sim N(0, \Omega_{\hat{\gamma}})$						
	Without Pre-1906 Condos			With Pre-1906 Condos		
No Lag	0.1073** (0.0425)	0.1722*** (0.0507)	0.1815*** (0.0499)	0.1035** (0.0425)	0.1800*** (0.0478)	0.1843*** (0.0462)
<i>Standard Error Inflation</i>	13.6%	17.9%	16.6%	21.1%	27.5%	24.5%
Lag $t - 6$	0.1011** (0.0493)	0.1775*** (0.0506)	0.1790*** (0.0549)	0.0461 (0.0561)	0.1580*** (0.0539)	0.1693*** (0.0569)
<i>Standard Error Inflation</i>	15.5%	10.5%	12.3%	58.5%	14.4%	10.5%
Bus Zone Vars	N	Y	Y	N	Y	Y
Year-by-Transit FE's	N	N	Y	N	N	Y

SOURCE: Table 5 uses the method of Conley, Hansen, and Rossi (2012) to reestimate the IV results presented in Table 3 for the contemporaneous shuttle effect, and the IV results in Table A1 for the six-month lagged shuttle effect.

NOTE: The key assumption under the local-to-zero (LTZ) method of Conley et al. (2012) is to provide a prior distribution for the excluded instruments' impact on the outcome variable. To estimate Ω_{γ} , I use the values generated by estimating Equation (6). These coefficients can be interpreted as the impact of moving 1km closer to a shuttle stop, given that the condo is within 2km of a stop.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE 6
No-Fault Eviction Probability Estimates
Panel Fixed Effects Models, July 2003–December 2013

	Without Pre-1906 Buildings ($N = 3,172,938$)					With Pre-1906 Buildings ($N = 4,228,005$)				
	<i>Ellis Mean DV= 0.00018, OMI Mean DV= 0.00034</i>					<i>Ellis Mean DV= 0.00018, OMI Mean DV= 0.00036</i>				
	(1)	(2)	(3)	(4)	(5)	(1)	(2)	(3)	(4)	(5)
OLS	<i>Ellis</i>	-0.00002 (0.00003)	0.00002 (0.00005)	0.00001 (0.00005)	0.00001 (0.00005)	0.00001 (0.00003)	0.00005 (0.00004)	0.00004 (0.00004)	0.00005 (0.00004)	0.00004 (0.00004)
	Adj R ²	0.006	0.006	0.006	0.006	0.006	0.006	0.006	0.006	0.006
	<i>OMI</i>	0.00001 (0.00004)	-0.00004 (0.00006)	-0.00005 (0.00006)	-0.00004 (0.00006)	0.00003 (0.00003)	-0.00001 (0.00005)	-0.00000 (0.00006)	0.00000 (0.00006)	0.00000 (0.00006)
	Adj R ²	0.004	0.004	0.004	0.004	0.004	0.004	0.004	0.004	0.004
IV	<i>Ellis</i>	-0.00006 (0.00007)	0.00023 (0.00021)	0.00015 (0.00019)	0.00015 (0.00019)	-0.00010 (0.00007)	0.00013 (0.00022)	0.00024 (0.00018)	0.00023 (0.00018)	0.00024 (0.00018)
	<i>OMI</i>	0.00005 (0.00012)	0.00022 (0.00038)	0.00024 (0.00036)	0.00023 (0.00036)	0.00006 (0.00012)	0.00029 (0.00034)	0.00042 (0.00032)	0.00040 (0.00031)	0.00057** (0.00028)
	F stat	71.07**	60.5**	90.17**	84.44**	83.96**	63.51**	80.76**	75.71**	102.85**
	5% CV	30.46	19.74	23.66	24.27	30.11	15.11	23.14	23.85	37.42
	10% CV	18.20	12.55	15.04	15.39	17.95	9.96	14.74	15.15	23.11
Nbrhd×Year	N	Y	Y	Y	Y	N	Y	Y	Y	Y
1{Corridor}×Year	N	N	Y	Y	Y	N	N	Y	Y	Y
Nbrhd×t	N	N	N	Y	N	N	N	N	Y	N
Nbrhd×GOOG	N	N	N	N	Y	N	N	N	N	Y

SOURCE: Table 6 displays the OLS and IV estimates for building-level no-fault eviction probabilities. Point estimates are obtained by estimating coefficients for the model given in Equation (9).

NOTE: “F-stat” refers to the Montiel Olea-Pflueger 1st-stage F statistic computed in the Ellis regressions, and are virtually indistinguishable from those reported from the OMI regression results; “5% CV” refers to the critical value leaving a p-value of 0.05 or less on the hypothesis that the F statistic is zero, with “10% CV” defined similarly.

* p<0.10, ** p<0.05, *** p<0.01.

TABLE 7
No-Fault Eviction Probability Estimates by Building Size in Bottom Three Price Quartiles
IV Fixed Effects, July 2003–December 2013

	Ellis					OMI				
	(1)	(2)	(3)	(4)	(5)	(1)	(2)	(3)	(4)	(5)
All (<i>N</i> = 2,860,046)	<i>Mean DV</i> =0.00017					<i>Mean DV</i> =0.00039				
	-0.00015 (0.00009)	0.00019 (0.00012)	0.00025* (0.00014)	0.00015 (0.00015)	0.00012 (0.00015)	0.00008 (0.00014)	0.00073** (0.00032)	0.00081** (0.00030)	0.00072** (0.00030)	0.00067** (0.00030)
	F stat	50.44**	39.97**	158.14**	108.97**	100.72**	50.44**	39.97**	158.14**	108.97**
7+ Units (<i>N</i> = 371,977)	<i>Mean DV</i> =0.00012					<i>Mean DV</i> =0.00009				
	0.00001 (0.00009)	-0.00067** (0.00033)	-0.00087** (0.00043)	-0.00102* (0.00051)	-0.00104* (0.00054)	0.00028** (0.00014)	0.00169 (0.00129)	0.00131 (0.00133)	0.00121 (0.00127)	0.00121 (0.00135)
	F stat	42.68**	31.70**	61.33**	57.12**	49.89**	42.68**	31.70**	61.33**	57.12**
3–6 Units (<i>N</i> = 1,103,681)	<i>Mean DV</i> =0.00024					<i>Mean DV</i> =0.00041				
	-0.00030** (0.00013)	0.00026 (0.00026)	0.00023 (0.00034)	0.00012 (0.00033)	0.00017 (0.00036)	0.00013 (0.00026)	0.00099** (0.00049)	0.00107** (0.00050)	0.00110*** (0.00041)	0.00112** (0.00045)
	F stat	51.82**	34.71**	139.76**	91.19**	89.14**	51.82**	34.71**	139.76**	91.19**
2 Units (<i>N</i> = 1,384,334)	<i>Mean DV</i> =0.00014					<i>Mean DV</i> =0.00046				
	-0.00010 (0.00014)	0.00034 (0.00022)	0.00046** (0.00019)	0.00040** (0.00018)	0.00033* (0.00019)	0.000016 (0.00016)	0.00040 (0.00056)	0.00058 (0.00061)	0.00042 (0.00058)	0.00033 (0.00057)
	F stat	43.29**	27.53*	175.64**	107.66**	96.68**	43.29**	27.53*	175.64**	107.66**
Nbrhd×Year	N	Y	Y	Y	Y	N	Y	Y	Y	Y
1{Corridor}×Year	N	N	Y	Y	Y	N	N	Y	Y	Y
Nbrhd×t	N	N	N	Y	N	N	N	N	Y	N
Nbrhd×GOOG	N	N	N	N	Y	N	N	N	N	Y

SOURCE: Table 7 displays the IV estimates for building-level no-fault eviction probabilities broken out by the building's unit count.

NOTE: The sample is restricted to neighborhoods whose July 2002–July 2004 average condo sales prices were in the bottom three quartiles, but includes historic (pre-1906) buildings. “F-stat” refers to the Montiel Olea-Pfueger 1st-stage F statistic.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE 8
Up and Down Conversion Probability Estimates by Building Size
IV Fixed Effects, July 2003–December 2013

		Up			Down		
		(1)	(2)	(3)	(1)	(2)	(3)
All (<i>N</i> = 2, 240, 673)		<i>Mean DV</i> = 0.00007			<i>Mean DV</i> = 0.00004		
	Shuttle 2km	-0.00031*	-0.00028**	-0.00026*	-0.00002	-0.00007	-0.00006
		(0.00016)	(0.00013)	(0.00013)	(0.00007)	(0.00005)	(0.00005)
	$\times \mathbb{1}\{RC\}$	0.000209	0.00010	0.00010	0.000073	0.00009	0.00009
		(0.000160)	(0.00011)	(0.00012)	(0.00008)	(0.00006)	(0.00006)
	Margin	-0.0001	-0.00019*	-0.00016	0.00005	0.00002	0.00003
	(0.000118)	(0.00011)	(0.00012)	(0.00006)	(0.00006)	(0.00006)	
	F stat	5.71	60.52	53.63	5.71	60.52	53.63
7+ Units (<i>N</i> = 336, 942)		<i>Mean DV</i> = 0.00008			<i>Mean DV</i> = 0.00004		
	Shuttle 2km	-0.00142*	-0.00120**	-0.00136**	-0.00092	-0.00093	-0.00080
		(0.00073)	(0.00045)	(0.00055)	(0.00108)	(0.00101)	(0.00102)
	$\times \mathbb{1}\{RC\}$	0.00120	0.00086	0.00096	0.00094	0.00095	0.00093
		(0.00085)	(0.00062)	(0.00061)	(0.00104)	(0.00100)	(0.00101)
	Margin	-0.00022	-0.00034	-0.00041	0.00003	0.00001	0.00013
	(0.00050)	(0.00053)	(0.00062)	(0.00017)	(0.00015)	(0.00012)	
	F stat	3.02	3.93	3.79	3.02	3.93	3.79
3–6 Units (<i>N</i> = 818, 368)		<i>Mean DV</i> = 0.00008			<i>Mean DV</i> = 0.00006		
	Shuttle 2km	-0.00071*	-0.00035	-0.00032	0.000076	-0.00007	-0.00008
		(0.00041)	(0.00025)	(0.00025)	(0.00013)	(0.00007)	(0.00008)
	$\times \mathbb{1}\{RC\}$	0.000693	0.00021	0.00023	0.00001	0.00007	0.00006
		(0.000493)	(0.00026)	(0.00026)	(0.00018)	(0.00011)	(0.00011)
	Margin	-0.000014	-0.00014	-0.00009	0.00008	-0.00001	-0.000013
	(0.000284)	(0.00026)	(0.00025)	(0.00010)	(0.00013)	(0.00014)	
	F stat	7.00	62.64	92.81	7.00	62.64	92.81
2 Units (<i>N</i> = 1, 085, 363)		<i>Mean DV</i> = 0.00006			<i>Mean DV</i> = 0.00002		
	Shuttle 2km	-0.00009	-0.00012	-0.00008	0.00005	-0.00000	-0.00001
		(0.00010)	(0.00011)	(0.00011)	(0.00006)	(0.00006)	(0.00005)
	$\times \mathbb{1}\{RC\}$	-0.00001	-0.00002	-0.00002	0.00000	0.00004	0.00004
		(0.00004)	(0.00006)	(0.00006)	(0.00004)	(0.00004)	(0.00004)
	Margin	-0.00010	-0.00014	-0.00010	0.00005	0.00004	0.00004
	(0.00011)	(0.00010)	(0.00009)	(0.00006)	(0.00005)	(0.00005)	
	F stat	75.56	48.04	43.35	75.56	48.04	43.35
Nbrhd×t		N	Y	N	N	Y	N
Nbrhd×GOOG		N	N	Y	N	N	Y

SOURCE: Table 8 displays the IV estimates for building-level up and down conversion probabilities broken out by unit counts. Point estimates are obtained by estimating coefficients for the model given in Equation (10).

NOTE: Historic (pre-1906) buildings are excluded here. “Margin” is the marginal effect of a shuttle stop moving 1km closer (within a 2km radius) to specifically a rent-controlled building. “F-stat” refers to the Kleibergen-Paap rk Wald 1st-stage F statistic.

* p<0.10, ** p<0.05, *** p<0.01.

TABLE 9
Repairs Probability Estimates by Building Size
IV Fixed Effects, July 2003–December 2013

		(1)	(2)	(3)
<i>Mean DV= 0.00071, N = 2,240,673</i>				
All	Shuttle 2km	-0.00134 (0.00168)	-0.00003 (0.00128)	-0.00047 (0.00126)
	× $\mathbb{1}\{RC\}$	-0.00092 (0.00108)	-0.00042 (0.00116)	-0.00075 (0.00125)
	Margin	-0.00225 (0.00214)	-0.00045 (0.00100)	-0.00121 (0.00128)
	F stat	17.58	19.79	23.84
<i>Mean DV= 0.00108, N = 336,942</i>				
7+ Units	Shuttle 2km	-0.00101 (0.00348)	0.00228 (0.00178)	0.00044 (0.00165)
	× $\mathbb{1}\{RC\}$	-0.00355 (0.00240)	-0.00302 (0.00198)	-0.00267 (0.00188)
	Margin	-0.00456 (0.00415)	-0.00073 (0.00217)	-0.00223 (0.00259)
	F stat	26.72	22.47	20.35
<i>Mean DV= 0.00019, N = 818,368</i>				
3–6 Units	Shuttle 2km	-0.00019 (0.00031)	-0.00015 (0.00031)	-0.00009 (0.00029)
	× $\mathbb{1}\{RC\}$	0.00010 (0.00021)	-0.00001 (0.00018)	0.00002 (0.00018)
	Margin	-0.00009 (0.00025)	-0.00016 (0.00027)	-0.00007 (0.00026)
	F stat	6.45	93.05	124.88
<i>Mean DV= 0.00027, N = 1,085,363</i>				
2 Units	Shuttle 2km	0.00058 (0.00035)	0.00030 (0.00037)	0.00025 (0.00039)
	× $\mathbb{1}\{RC\}$	-0.00005 (0.00027)	0.00036 (0.00039)	0.00037 (0.00040)
	Margin	0.00053*** (0.00016)	0.00066*** (0.00020)	0.00063*** (0.00019)
	F stat	75.56	48.04	43.35
<hr/>				
	Nbrhd×t	N	Y	N
	Nbrhd×GOOG	N	N	Y

SOURCE: Table 9 displays the IV estimates for building-level repair permit filing probabilities broken out by unit count. Point estimates are obtained by estimating coefficients for the model given in Equation (10).

NOTE: “Margin” is the marginal effect of a shuttle stop moving 1km closer (within a 2km radius) to specifically a rent-controlled building. “F-stat” refers to the Kleibergen-Paap rk Wald 1st-stage F statistic.

* p<0.10, ** p<0.05, *** p<0.01.

TABLE 10
At-Fault Eviction Probability Estimates by Building Size
IV Fixed Effects, July 2003–December 2013

		(1)	(2)	(3)
<i>Mean DV= 0.01193, N = 2, 240, 673</i>				
All	Shuttle 2km	0.03404 (0.03539)	0.02088 (0.02850)	0.01640 (0.02854)
	$\times \mathbb{1}\{RC\}$	-0.04925 (0.03277)	-0.01888 (0.02127)	-0.01652 (0.02149)
	Margin	-0.01521 (0.01548)	0.00200 (0.02216)	-0.00012 (0.02467)
	F stat	17.58	19.78	23.84
<i>Mean DV= 0.02009, N = 336, 942</i>				
7+ Units	Shuttle 2km	0.08960 (0.07354)	0.15137 (0.10222)	0.13890 (0.09678)
	$\times \mathbb{1}\{RC\}$	-0.12277* (0.06684)	-0.14989* (0.08013)	-0.13940* (0.07123)
	Margin	-0.03317 (0.03335)	0.00148 (0.04382)	-0.00050 (0.05306)
	F stat	26.72	22.47	20.35
<i>Mean DV= 0.00112, N = 818, 368</i>				
3–6 Units	Shuttle 2km	-0.00047 (0.00138)	0.00008 (0.00144)	0.00018 (0.00140)
	$\times \mathbb{1}\{RC\}$	0.00094 (0.00114)	0.00019 (0.00132)	0.00018 (0.00134)
	Margin	0.00047 (0.00056)	0.00027 (0.00051)	0.00036 (0.00049)
	F stat	6.45	93.05	124.88
<i>Mean DV= 0.00045, N = 1, 085, 363</i>				
2 Units	Shuttle 2km	-0.00057 (0.00044)	-0.00107 (0.00065)	-0.00109* (0.00065)
	$\times \mathbb{1}\{RC\}$	0.00009 (0.00049)	0.00089 (0.00065)	0.00087 (0.00065)
	Margin	-0.00048 (0.00032)	-0.00019 (0.00034)	-0.00023 (0.00032)
	F stat	75.56	48.04	43.35
<hr/>				
Nbrhd $\times$ t		N	Y	N
Nbrhd $\times$ GOOG		N	N	Y

SOURCE: Table 10 displays the IV estimates for building-level repair at-fault eviction probabilities broken out by unit count. Point estimates are obtained by estimating coefficients for the model given in Equation (10).

NOTE: “Margin” is the marginal effect of a shuttle stop moving 1km closer (within a 2km radius) to specifically a rent-controlled building. “F-stat” refers to the Kleibergen-Paap rk Wald 1st-stage F statistic.

* p<0.10, ** p<0.05, *** p<0.01.

TABLE 11
Marginal Effects in Controlled Buildings for All Outcomes by Building Size
OLS Fixed Effects, July 2003–December 2013

	Ellis	OMI	Up	Down	Repair Permits	At- Fault
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Mean DV</i>	0.00017	0.00039	0.00007	0.00004	0.00070	0.00873
All	0.00004 (0.00005)	0.00004 (0.00007)	0.00005** (0.00002)	0.00001 (0.00003)	-0.00045 (0.00068)	-0.00155 (0.00567)
<i>Mean DV</i>	0.00012	0.00009	0.00008	0.00003	0.00108	0.01496
7+ Units	-0.00007 (0.00012)	-0.00006 (0.00006)	-0.00001 (0.00007)	-0.00004 (0.00010)	-0.00130 (0.00148)	-0.00393 (0.01051)
<i>Mean DV</i>	0.00024	0.00041	0.00008	0.00006	0.00020	0.00114
3–6 Units	0.00010 (0.00008)	0.00011 (0.00015)	0.00012** (0.00005)	0.00001 (0.00006)	0.00006 (0.00024)	0.00015 (0.00010)
<i>Mean DV</i>	0.00014	0.00046	0.00006	0.00003	0.00028	0.00044
2 Units	0.00003 (0.00007)	0.00002 (0.00012)	0.00003 (0.00002)	0.00002 (0.00002)	0.00018** (0.0009)	0.0005 (0.00012)

SOURCE: Table 11 displays the IV estimates for building-level no-fault eviction probabilities broken out by the building's unit count.

NOTE: Standard errors are clustered at the neighborhood level. "F-stat" refers to the Montiel Olea-Pflueger 1st-stage F statistic.

* p<0.10, ** p<0.05, *** p<0.01.

TABLE 12
Policy Impacts on At-Fault Eviction Probabilities (From IV Estimates), Jul 2003–Dec 2013

	All Buildings	7+ Units	3-6 Units	2 Units
At-Faults				
Rule 1: Landlords allowed to evict for lease changes imposed by law.	0.02310** (0.01123)	0.02593** (0.01105)	0.00032 (0.00025)	0.00003 (0.00032)
Rule 2: Tenants with minor children cannot be OMI evicted during school year.	0.01038 (0.00678)	0.03148* (0.01831)	0.00048 (0.00031)	0.00006 (0.00038)
Rule 3: Landlords must pay relocation payments for Ellis Act evictions.	0.04084** (0.01911)	0.08655*** (0.03056)	-0.00087 (0.00101)	-0.00090 (0.00152)
Rule 4: Landlords must pay relocation payments for other no-fault evictions.	-0.00085 (0.02085)	0.09156* (0.04984)	0.00031 (0.00087)	-0.00040 (0.00047)
Rule 5: Uncontrolled landlords only need to give 30 days notice to terminate lease.	0.00034 (0.01253)	-0.04904** (0.02367)	-0.00042 (0.00049)	-0.00038 (0.00034)

SOURCE: Table 12 displays the IV estimates for various eviction and rent control policy changes' impact on at-fault eviction and repair permit filing probabilities. Point estimates are obtained by estimating coefficients for the model given in Equation (10), corresponding to Column 3 in Tables 9 and 10.

* p<0.10, ** p<0.05, *** p<0.01.

Figures

Figure 1: Evolution of Apple, Electronic Arts, Facebook, and Google Shuttle Stops

September 2004

April 2006

May 2008

September 2009

December 2013

SOURCE: Author's renderings from distribution of shuttle stops over time.

Figure 2: San Francisco Transit Networks as of September 2004

SOURCE: Author's rendering of transit networks in San Francisco, using data from BART, CalTrain, and the SFMTA.

Figure 3: Condominium Transactions versus Rent-Controlled Buildings

SOURCE: Author's renderings using Assessor's data.

A Additional Results Appendix

The first table in this Appendix presents the results of lagging the shuttle measures back six months, and shows that this does little to change the estimates. The second table is a robustness check that attempts to control for ownership and tenant changes in each building nonparametrically with building-by-year fixed effects. Specification one includes building-by-year and year/month-by-unit group fixed effects; specification two adds neighborhood time trends; specification three drops neighborhood time trends and adds neighborhood fixed effects interacted with *GOOG*. Table A2 shows that the general thrust of the results stays the same. Large buildings appear to be less likely to perform Ellis Act evictions, midsized buildings are more likely to perform an OMI eviction (and is significant at the 1 percent level in specification three), and small landlords are more likely to perform Ellis Act evictions. The most notable difference between Tables 7 and A2 is that the point estimates become even larger after adding building-by-year fixed effects. Nonetheless, the fact that the patterns in coefficient signs and significance remain the same helps reinforce the results robustness to additional fixed effects levels.

The third table is a second robustness check that reestimates the specifications discussed in Table 7, but now excludes the historic buildings. It is important to verify that the restricted and unrestricted samples tell the same story, because the up and down conversion, at-fault, and repair filing analysis use the restricted sample only.³³ Table A3 has the results from the restricted sample, and indeed has almost all the same coefficient signs as found in Table 7, particularly for specifications three through five. The standard errors are similar, but the coefficient point estimates are smaller compared to Table 7. The conclusion here is that the older, more historic buildings drive the strength of the main results, but the underlying dynamic is the same for newer and older buildings.

³³See Asquith (2019b) for a more in-depth discussion on the perils of extending the controlled sample back before 1906 when using an uncontrolled sample that almost exclusively includes building built between 1979–2002.

TABLE A1
Hedonic Price Regression on Condominium Sales
Lag $t - 6$ on Both Samples, July 2003–December 2013

		$\mathbb{1}\{\text{Shuttle}\}$			Shuttle 2km		
Without Historic Condos							
OLS	Shuttle Measure	(1) 0.0312 (0.0213)	(2) 0.0528** (0.0257)	(3) 0.0554** (0.0272)	(4) 0.0539* (0.0273)	(5) 0.0707** (0.0301)	(6) 0.0744** (0.0327)
	Adj R^2	0.663	0.673	1.000	0.664	0.674	1.000
	F	519.36	878.34	518.10	520.77	859.17	542.99
IV	Shuttle Measure	-0.0343 (0.0566)	0.0654 (0.0535)	0.0651 (0.0587)	0.1011** (0.0427)	0.1775*** (0.0458)	0.1790*** (0.0489)
	Hansen's J p-value	0.095	0.060	0.138	0.299	0.069	0.111
	First-stage F stat	9.77	11.22	7.04	22.85*	29.99**	27.76*
	5% Crit Value	29.25	28.36	30.07	28.33	27.64	28.05
	10% Crit Value	16.98	16.33	17.58	16.31	15.80	16.10
With Historic Condos							
OLS	Shuttle Measure	(1) 0.0312 (0.0213)	(2) 0.0528** (0.0257)	(3) 0.0554** (0.0272)	(4) 0.0539* (0.0273)	(5) 0.0707** (0.0301)	(6) 0.0744** (0.0327)
	Adj R^2	0.663	0.673	1.000	0.664	0.674	1.000
	F	519.36	878.34	518.10	520.77	859.17	542.99
IV	Shuttle Measure	-0.0391 (0.0545)	0.0537 (0.0522)	0.0444 (0.0593)	0.0462 (0.0354)	0.1580*** (0.0471)	0.1693*** (0.0515)
	Hansen's J p-value	0.081	0.090	0.162	0.074	0.260	0.451
	First-stage F stat	11.89	9.85	6.86	14.02	15.67	13.25
	5% Crit Value	28.78	28.98	31.14	28.61	28.49	29.78
	10% Crit Value	16.64	16.78	18.37	16.51	16.42	17.37
Bus Zone Vars		N	Y	Y	N	Y	Y
Year-by-Transit FE's		N	N	Y	N	N	Y

SOURCE: Table A1 shows the impact of tech company shuttle stop proximity on condominiums' log sales price generated from Equation (4) in Column (1), with additional specifications reported in Columns (2) and (3).

NOTE: The first shuttle measure is within a half mile of a shuttle stop ($\mathbb{1}\{\text{Shuttle}\}$) or the condo's kilometer-measured distance from a 2km radius ring around the condo's closest shuttle stop. These coefficients can be interpreted as the impact of moving 1km closer to a shuttle stop, given that the condo is within 2km of a stop. Both measures are lagged six months.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE A2
No-Fault Eviction Probability Estimates by Building Size
Adding Building×Year Fixed Effects
IV Fixed Effects, July 2003–December 2013

	Ellis			OMI		
	(1)	(2)	(3)	(1)	(2)	(3)
All (<i>N</i> = 2, 860, 046)	0.00045 (0.00053)	0.00069 (0.00044)	0.00098 (0.00078)	0.00089 (0.00077)	0.00102 (0.00064)	0.00084 (0.00221)
F stat	60.59**	30.89**	33.40*	60.59**	30.89**	33.4*
7+ Units (371, 977)	-0.00137* (0.00080)	-0.00153* (0.00088)	-0.00194 (0.00121)	0.00082 (0.00107)	0.00047 (0.00107)	0.00147 (0.00502)
F stat	33.94**	30.83*	50.1**	33.94**	30.83*	50.1**
3–6 Units (1, 103, 681)	-0.00061 (0.00081)	-0.00017 (0.00066)	-0.00029 (0.00165)	0.00256* (0.00129)	0.00236** (0.00098)	0.00536*** (0.00172)
F stat	48.22**	25.17*	25.93*	48.22**	25.17*	25.93*
2 Units (1, 384, 388)	0.00201*** (0.00054)	0.00226*** (0.00041)	0.00272*** (0.00076)	-0.00070 (0.00142)	0.00000 (0.00129)	-0.00250 (0.00310)
F stat	51.13**	31.28*	33.62*	51.13**	31.28*	33.62*
Nbrhd×t	N	Y	N	N	Y	N
Nbrhd×GOOG	N	N	Y	N	N	Y

SOURCE: Table A2 displays the IV estimates for building-level no-fault eviction probabilities with the addition of building-by-year fixed effects. Point estimates are obtained by estimating coefficients for the model given in Equation (9).

NOTE: Historic (pre-1906) buildings are included here. “F-stat” refers to the Montiel Olea-Pfueger 1st-stage F statistic.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE A3
No-Fault Eviction Probability Estimates by Building Size, Excluding Historic (Pre-1906) Buildings
IV Fixed Effects, July 2003–December 2013

	Ellis					OMI				
	(1)	(2)	(3)	(4)	(5)	(1)	(2)	(3)	(4)	(5)
All (<i>N</i> = 2,090,152)	<i>Mean DV</i> =0.00017					<i>Mean DV</i> =0.00036				
	-0.00009 (0.00008)	0.00013 (0.00013)	0.00011 (0.00014)	0.00006 (0.00014)	0.00002 (0.00015)	0.00011 (0.00015)	0.00063 (0.00039)	0.00064* (0.00037)	0.00057 (0.00037)	0.00053 (0.00037)
	F stat	41.43**	39.67**	174.59**	117.80**	109.76**	41.43**	39.67**	174.59**	117.80**
7+ Units (<i>N</i> = 321,384)	<i>Mean DV</i> =0.00011					<i>Mean DV</i> =0.00007				
	0.00003 (0.00006)	-0.00051 (0.00034)	-0.00073 (0.00044)	-0.00083 (0.00050)	-0.00086 (0.00052)	0.00040*** (0.00013)	0.00043 (0.00026)	0.00003 (0.00046)	-0.00001 (0.00046)	-0.00010 (0.00053)
	F stat	39.42**	36.09**	75.34**	70.40**	61.39**	39.42**	36.09**	75.34**	70.40**
3–6 Units (<i>N</i> = 767,952)	<i>Mean DV</i> =0.00026					<i>Mean DV</i> =0.00037				
	-0.00035*** (0.00013)	-0.00003 (0.00027)	0.00001 (0.00029)	-0.00003 (0.00029)	-0.00004 (0.00033)	0.00018 (0.00027)	0.00074 (0.00057)	0.00079 (0.00050)	0.00091* (0.00046)	0.00092* (0.00050)
	F stat	43.64**	35.30**	164.24**	102.25**	101.22**	43.64**	35.30**	164.24**	102.25**
2 Units (<i>N</i> = 1,000,816)	<i>Mean DV</i> =0.00012					<i>Mean DV</i> =0.00044				
	0.00006 (0.00010)	0.00041* (0.00022)	0.00033* (0.00019)	0.00031* (0.00018)	0.00023 (0.00019)	-0.00001 (0.00019)	0.00057 (0.00064)	0.00069 (0.00067)	0.00050 (0.00065)	0.00043 (0.00064)
	F stat	32.92**	26.99*	188.21**	109.51**	100.07**	32.92**	26.99*	188.21**	109.51**
Nbrhd×Year	N	Y	Y	Y	Y	N	Y	Y	Y	Y
1{Corridor}×Year	N	N	Y	Y	Y	N	N	Y	Y	Y
Nbrhd×t	N	N	N	Y	N	N	N	N	Y	N
Nbrhd×GOOG	N	N	N	N	Y	N	N	N	N	Y

SOURCE: Table A3 displays the instrumental variable estimates for building-level no-fault eviction probabilities over the stated time period. Point estimates are obtained by estimating coefficients for the model given in Equation (9).

NOTE: “F-stat” refers to the Montiel Olea-Pflueger 1st-stage F statistic.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

B Additional Policy Appendix

Section 1 gives a brief overview of the main policy points relevant to understanding the base mechanics of rent control in San Francisco. This appendix covers a much broader set of rules that govern landlord and tenant behavior.

B.1 Eviction Regulations

Unlike uncontrolled landlords, controlled landlords must have a “just cause” for an eviction. The 15 grounds for a “just-cause” eviction are given in Table B1. Seven are for an “at-fault” tenant, who is in some way in breach of the lease and eight where the tenant is not at fault, or “no-fault”. Six of the no-fault evictions are not part of the study because they are either temporary or too rarely occur. These include temporary eviction for lead abatement or capital improvements; revoking of “Good Samaritan” status for tenants who are fleeing natural disasters; and converting rental units to condominiums, more on which below in Section B.3. Demolitions are somewhat more commonly granted, but still too rare to generate plausible estimates in a monthly panel.

The two no-fault eviction types used in this study, owner move-in and Ellis Act evictions, are the most common but come with some significant regulations as well. Table B2 shows how San Francisco passed various policy changes between July 2003 and December 2013 to regulate controlled evictions. Other rules include suspending vacancy decontrol on withdrawn units for up to 3 years after an OMI eviction and 10 years after an Ellis Act eviction if the landlords reread the units. Landlords can only do one OMI eviction per building and the set-aside unit is marked on the deed. A post-Ellis Act vacant building faces additional restrictions. A ten year period is marked on the deed where the new building exemption is suspended for the property. If the landlord demolishes the old units and build new ones during this time, rent control will apply until the waiting period expires.

Non-controlled landlords can evict tenants without cause, pursuant only to the lease and

relevant state and city statutes. However, San Francisco is clear that controlled unit evictions should only happen “in good faith” (San Francisco Administrative Code §37.9(8)).³⁴ The good faith requirement also pertains to at-fault evictions. If the city determines that the landlord wrongly took possession of the unit, the city can sue for injunctive relief and monetary damages three times actual damages.³⁵

Evictions can be avoided altogether through a buyout agreement. Court cases for at-fault evictions are costly, time-consuming, and uncertain, so that landlords prefer to buy out a tenant first (Downs, 1988). Legal buyouts do not mean that the economic eviction rate is zero. Evictions are a landlord’s credible threat to achieve a higher rate of successful buyouts, but then can be used as a final resort on recalcitrant tenants. Legal buyouts’ significance for this study is only to artificially lower the observed economic eviction rate.

Buyout agreements could more seriously bias the results if there are reasons to believe that tenants are more (or less) likely to sign them when rents rise. Unfortunately, San Francisco only started regulating and publishing detailed information on buyouts in 2015.³⁶ Further, the literature is silent on how to sign this bias. From first principles, rising rents could make tenants more fearful of eviction because finding comparable units is now more expensive. They would then be more inclined to sign a buyout, biasing downwards the sign on any exogenous local shock to rents in an evictions regression. Thus, the results are likely to understate both the impact of rent control and the impact of a rent shock on controlled landlords.

San Francisco’s rent control and eviction provisions are not unique. Los Angeles, Oakland, San Jose, New York, and Washington D.C. have very similar rules. Table B3 lists the rent control and eviction rules in these cities analogous to the ones in San Francisco listed above. All five cities have at least limited vacancy decontrol coupled with yearly rent increase

³⁴More specifically, landlords can only use the no-fault evictions to “...recover possession in good faith, without ulterior reasons and with honest intent”, San Francisco Administrative Code § 37.9(8)

³⁵San Francisco Administrative Code 37.9(f).

³⁶Currently, only a limited time series is available, although this will be a rich source of information for future researchers in a decade.

restrictions, new building exemptions, and allow tenants to indefinitely renew their leases. Only San Jose does not require a just cause for an eviction, but even there, landlords have to go through a city-mandated arbitration process to evict.

B.2 Eviction Costs

Beyond the usual court costs The next policy consideration is to examine financial costs for evictions. There are at least two benchmarks for whether prices rise enough to incentivize evictions:

- 1. Rent changes should exceed the annual allowable increase**
- 2. Per unit price changes should exceed the fixed cost of a no-fault eviction from relocation payments.**

The first consideration applies to both at-fault and no-fault evictions, as it means that controlled landlords often cannot realize market gains and are incentivized to push out tenants. Between 2003 and 2013, annual allowable rent increases ranged from 0.1% (Mar 2010-Feb 2011) to 2.2% (March 2009-Feb 2010), with an average of 1.3%.³⁷ The second condition means that landlords are disincentivized from performing no-fault evictions, but also that they have reasons to be more careful about giving leases to elderly and disabled tenants or tenants with a minor child. Table B4 shows what landlords would have to pay to different tenant types in the case of an Ellis act or other no-fault evictions, like OMIs. Relocation payments grew with time so that by December 2013, a landlord had to pay roughly \$5,200 for each evicted tenant, capped at about \$15,620, with a protected surcharge of about \$3,470.

³⁷“Allowable Annual Rent Increases”, The Residential Rent Stabilization and Arbitration Board of the City and County of San Francisco, 2016.

B.3 Condo Conversion Regulations

One important rent control exit channel, condo conversion, could only be done via one of two processes. The first was an annual conversion lottery that operated until spring 2013, and capped conversions at 200 units per year. Only buildings with 2-6 unit were eligible to participate, and buildings with 7 or more units had no ability to condo convert at all. The lottery gave very strong preference to landlords with a “clean” eviction history. More information on the lottery is in Asquith (2019a, 2019b), but the most relevant point for this paper is that landlords who felt they had a good chance to win the lottery had a strong disincentive to performing an eviction.

2 unit buildings had two additional means of exiting rent control. The first is that they could bypass the lottery to condo convert if they satisfied an ownership rule where by two, non-related separate owners had at least a 25% stake in each apartment. This lottery bypass was unlimited and allowed to proceed even after the lottery was suspended. Importantly, there is no “clean” eviction requirement for utilizing the lottery bypass to condo convert. A landlord could perform an Ellis Act eviction, occupy one unit, find a buyer for the now-vacant other unit, and then convert the building to a tenancy-in-common. After a year, the TIC units could be converted to condos, with the ability to resell the units. The other means to exit rent control was that if they chose to down convert to a single-family unit, but could only do so if they did not have a tenant who had been continuously in residence since 1996.

B.4 Additional Regulations

Beyond the condo conversion and eviction system, there are some additional, relevant rules that controlled landlords must operate by before changing their supply. Landlords were initially allowed to pass on capital improvements costs to tenants, which was then exploited by some landlords as backdoor means of increasing rents beyond the annual cap. On November 7, 2000, Proposition H passed, which effectively barring landlords from passing through any capital improvements on to their tenants except the bear minimum required to give landlords

their Constitutionally-mandated fair return. After a lawsuit and a permanent injunction, Proposition H went into effect in April 2003, and effectively forced landlords to adopt longer amortization tables and limited increases to 5% of the tenant's base rent as of the time the petition was filed or \$30.00, whichever is greater, in any 12 month period.

All landlords could exit rent control through substantial rehabilitation as well, but the city has let only 6 buildings exit rent control this way since 1992 so it can be safely set aside in this analysis. The main reason it is not more commonly pursued is that only buildings 50 years or older are eligible, and the landlord must perform renovations equivalent to 75% of the cost of constructing the building anew.³⁸

³⁸San Francisco Administrative Code Chapter 37, Section 2(s). Correspondence with the San Francisco Rent Board Executive Director Robert Collins on June 10, 2019 confirmed that the 50 year requirement has been in place since the ordinance was first passed in 1979.

TABLE B1
The 15 Grounds for “Just Cause” Eviction in San Francisco

Reason	Type	Relocation Payments?	Deed Restrictions?^a
<i>Permanent</i>			
Non-payment or habitual late payment on rent	At-Fault	No	No
Breach of lease	At-Fault	No	No
Nuisance or substantial damage to unit	At-Fault	No	No
Conducting Illegal Actions in Unit ^b	At-Fault	No	No
Tenant refuses to quit after tenancy ends	At-Fault	No	No
Tenant refuses to grant landlord lawful access	At-Fault	No	No
Sole remaining tenant is unapproved subtenant	At-Fault	No	No
Owner repossession for primary residence (OMI)	No-Fault	Yes	Yes
Conversion of units to condominiums ^c	No-Fault	Yes	No
Removal of all units from rental use (Ellis Act)	No-Fault	Yes	Yes
Demolition of units	No-Fault	Yes	No
Substantial Rehabilitation	No-Fault	Yes	No
“Good Samaritan” status has expired ^d	No-Fault	No	No
<i>Temporary</i>			
Lead abatement	No-Fault	Yes	No
Capital improvements	No-Fault	Yes	No

SOURCE: San Francisco Administrative Code Chapter 37, Section 9(a)(1)-9(a)(16).

NOTE: Table B1 enumerates the reasons a landlord may reclaim a rent-controlled unit. The “At-Fault” evictions refer to the 7 ways a tenant may be evicted for breaching the rental contract in some fashion, and “No-Fault” refers to the 8 ways a tenant may be evicted even if not in breach of the lease.

^a These include restrictions on how long the landlord must wait before being able to return the units to market, or if the unit is demolished, how long the parcel will remain under the rent ordinance before its provisions are lifted. These range from 3 years for an OMI to 10 years for an Ellis Act eviction.

^b Conversion of rental units to condominiums was previously possible via a permit lottery but was suspended in 2013. However, the city only permitted a handful of these per year prior to its formal suspension.

^c If the tenant is convicted of a crime, the notice to quit is unconditional.

^d “Good Samaritan” status is temporary housing for tenants fleeing a natural disaster.

TABLE B2 Policy Changes Regulating Evictions in San Francisco: July 2003-December 2013		
Description	Start Date	End Date
<u>General Eviction Rules</u>		
Landlords who wish to terminate that tenancy are no longer required to give 60 days notice, only 30-days notice, for tenants who have resided in the premises for one year or more.	1/1/2006	
Owners of properties with two or more residential units must disclose to any prospective purchaser the legal grounds for terminating the tenancy of each unit vacant at the close of escrow and whether the unit was occupied by an elderly or disabled tenant at the time the tenancy was terminated.	6/6/2006	
Reinstated the prior requirement of a 60 day notice to terminate a tenancy without a tenant fault good cause for any tenant or resident residing in the unit for a year or more.	1/1/2007	12/31/2009
A tenant who has resided in the unit for at least one year, and has a child under the age of 18 who also resides in the unit, may not be evicted during the school year for an OMI eviction.	3/14/2010	
Tenant may not be evicted for violation of a unilaterally imposed change in the terms of a tenancy unless the tenant previously accepted it in or the newly imposed term is authorized by the Rent Ordinance.	12/14/2011	2/1/2012
Allows a landlord to evict a tenant for violation of a unilaterally imposed change in terms where the change is required by law	2/1/2012	
Condo conversion evictions are suspended	8/1/2012	
<u>Ellis Act</u>		
Landlords must state in Ellis Act eviction notices that tenants have the right to relocation payments and the amount which the landlord believes to be due.	7/25/2005	1/30/2006
Landlords are no longer required to state the amount of relocation payment the landlord believes to be due to the tenant	1/31/2006	
<u>Owner Move-In</u>		
Landlords seeking to challenge a tenants' protected status for an OMI eviction have to file a petition rather than seeking a court order.	2006	

Source: The Rent Board of the City and County of San Francisco.

TABLE B3
Major City Rent Control and Evictions Policies, October 2016

City	Subject to Controls if the Building is...	Max Annual Allowable Rent Increase	Vacancy Decontrol?	Just-Cause Evictions?	Rental Stock Coverage
Los Angeles ^a	Built Before 10/1/1978 and has 2 or More Units	Regional CPI Rate, Bounded within 3-8%	Yes	Yes	85%
Oakland ^b	Built Before 1/1/1983 and has 4 or More Units	Regional CPI Rate, Max of 10%	Yes	Yes	66%
New York City ^c	Built Before 1/1/1974 and has 6 or More Units	Set by NYC Rent Guidelines Board Annually	No, rent increase for new base rent capped at 20%. ^d	Yes	47%
San José ^f	A Rental Unit Built Before 9/7/1979	Previously 8% 6/2016-: 5%	Yes	No, city-mandated arbitration instead.	33%
Washington, D.C. ^g	An Apartment Building Built Before 1/1/1976	CPI + 2%, Max of 10%	No, rent increase for new base rent capped at 10%. ^d	Yes	66%

^a SOURCES: Los Angeles Municipal Code Chapter 151. Coverage figure comes from the Office of the Mayor of Los Angeles, January 26, 2016. “Mayor Garcetti Announces New Access to Information on L.A.’s Rent-Stabilized Buildings.”, last accessed October 14, 2016.

^b SOURCES: Oakland Municipal Code §8.22 et seq. Coverage figure comes Sam Levin, July, 1, 2015. “When Landlords Target Tenants in Rent-Controlled Buildings.” East Bay Express.

^c SOURCES: The New York State Rent Regulation Reform Act of 1997, 1997 New York Laws 116; The Rent Act of 2015, 2015 New York Laws 20; The New York State Emergency Tenant Protection Act of 1974, 1974 New York Laws 576 §5-a. Coverage figure comes from Sieg, Holger and Chamna Yoon, 2016. “Waiting for Affordable Housing in New York City.” Working Paper.

^d In both cities, landlords can appeal for a rent increase on new base rents of up to 30% if rents in comparable units are shown to be higher.

^f SOURCES: San José Municipal Code, Apartment Ordinance, Chapter 17.23. Coverage figure is from San José Municipal Ordinance No. 29730, p. 1.

^g SOURCES: Code of the District of Columbia, Chapter 42. Coverage figure comes from Tatian, Peter A. and Ashley Williams, 2011. “A Rent Control Report for the District of Columbia.” The Urban Institute, Washington, D.C.

NOTE: All cities exempt new buildings, but developers in NYC can get a tax rebate on new buildings if they agree to controls for 10 years. After 10 years, all new vacancies are fully decontrolled. Source: NYC Rent Stabilization Code, Sections 2520.11(o), (p), (r), and (s).

TABLE B4^a
Relocation Payments for No-Fault Evictions: February 2000-February 2014

		Ellis Act				Other No-Fault ^b		
Start Date	End Date	Low Inc. Tenant	General Tenant	Max Payment	Special Surcharge ^c	General Tenant	Max Payment	Special Surcharge ^b
2/13/2000	8/9/2004	4,500	0	0	3,000	1,000	0	0
8/10/2004	4/24/2005	4,500	4,500	13,500	3,000	1,000	0	0
4/25/2005	5/25/2005	4,500	0	0	3,000	1,000	0	0
5/26/2005	2/28/2006	4,503	4,503	13,510	3,047	1,000	0	0
3/1/2006	8/9/2006	4,503	4,503	13,510	3,047	1,000	0	0
8/10/2006	2/28/2007	4,503	4,503	13,510	3,047	4,500	13,500	3,000
3/1/2007	2/28/2009	4,572	4,572	13,716	3,048	4,568	13,705	3,046
3/1/2009	2/28/2010	4,945	4,945	14,836	3,297	4,941	14,825	3,295
3/1/2010	2/28/2011	5,105	5,105	15,316	3,403	5,101	15,304	3,401
3/1/2011	2/29/2012	5,105	5,105	15,316	3,403	5,101	15,304	3,401
3/1/2012	2/28/2013	5,175	5,175	15,472	3,438	5,153	15,460	3,436
3/1/2013	2/28/2014	5,211	5,211	15,633	3,474	5,207	15,621	3,472

SOURCES: The Rent Board of the City and County of San Francisco.

NOTE: Table B4 shows the mandated relocation payments given to tenants for Ellis Act evictions and all other no-fault evictions. “Low Income Tenants” are the payments originally only given to poor tenants before August 2004 for Ellis Act evictions before being extended to all tenants. “General Tenants” are the relocation payments that were given to any controlled tenant. All amounts are in nominal US dollars..

^a From March 2006 onwards, payments were adjusted each March (at the discretion of the Rent Board) using the Consumer Price Index calculated for the San Francisco-Oakland-San Jose Combined Statistical Area.

^b “Other No-Fault” includes Owner move-in, demolitions, temporary capital improvement work, or substantial rehabilitation.

^c “Protected Surcharge” refers to the extra relocation payment the landlord pays if one of the evicted tenants is a minor, an elderly adult aged 60+, or who is disabled within the meaning of §12955.3 of the California Government Code.

Data Appendix

C Commuter Shuttle Stop Data Appendix

This Appendix has a list of citations for the sources used for shuttle stops and routes, and lists any assumptions made on routes. Information on shuttle routes comes chiefly from Google,³⁹ Stamen Design,⁴⁰ and the publicly available sources detailed below. There are some gaps in the publicly-available information on where the shuttle stops are, so some assumptions were made about timing and placement. These are detailed in Section C.4.

C.1 Websites

Stamen Design's Map of Selected Company's Shuttle Routes as of August 2012, <http://stamen.com/zero1/>.

The Anti-Eviction Mapping Project's Map of Tech Bus Stops and No-Fault Evictions, <http://www.antievictionmappingproject.net/techbusevictions.html>.

Map of Google Shuttle Stops as of April 2010 Maintained by Anonymous Google Maps User, https://www.google.com/maps/d/viewer?mid=zGkojhWojNBo.kH7J2er3ffro&hl=en_US.

Map of Google Shuttle Stops as of October 2011 Maintained by Anonymous Google Maps User, https://www.google.com/maps/d/viewer?mid=zM2GuPjAzei0.kZZnxCxSIWAg&hl=en_US.

Map of Google Shuttle Stops as of June 2012 Maintained by Anonymous Google Maps User, <https://www.google.com/maps/d/viewer?mid=z2hLFvZ7Lg5A.k00paEK693Qc&hl=en>.

Map of Google Shuttle Stops as of January 2013, <http://www.lookingformaps.com/mapa.php?mapa=Shuttle-Commuter-Stops-Effective-1-3-13>.

Map of Google Shuttle Stops as of July 2013 Maintained by Anonymous Google Maps User, https://www.google.com/maps/d/viewer?mid=z6pdx1VOR3mU.kAHGFAmzMOVo&hl=en_US.

Map of Apple Shuttle Stops as of September 2013 Maintained by Google Maps User Gtok, <https://maps.google.com/maps/ms?gl=US&ptab=1&ie=UTF8&oe=UTF8&msa=0&msid=21581207767043290004e5e26e576e99b6e0b&dg=feature>.

³⁹Interview with Brendon Harrington conducted on May 30, 2014

⁴⁰<http://content.stamen.com/zero1>

Map of Yahoo! Shuttle Stops as of August 2009 Maintained by Google Maps User Chris, <https://www.google.com/maps/d/viewer?dg=feature&msa=0&mid=z1TFSs454VJc.kQ1HwuIQzKDg>.

A list of Electronic Art's Stops as of January 17, 2012 is available at <http://www.etc.cmu.edu/siliconvalley/blog/faq/>.

A list of Shuttle Stops Entered Into Foursquare Curated by User Zach as of July 2012 is available at <http://dotspotting.org/u/939/sheets/2227/#c=11.00/37.7550/-122.4328>

C.2 News Stories and Publications

Anders, Corrie M. Google Shifts Bus Stop to Church Street. *The Noe Valley Voice* (October 2007), 2.

Brousseau, Fred, 2014. Impact of Private Shuttles. City and County of San Francisco Board of Supervisors, 1-36.

Dai, David and Danielle Dai, 2014. Riding First Class: Impacts of Silicon Valley Shuttles on Commute & Residential Location Choice. Working Paper.

EA Staff. EA's Bay Area Shuttle Increases Ridership. *EA News* (September 28, 2012), 2.

Farivar, Cyrus. Apple Launches Employee Shuttle This Week. *MacUser* (October 25, 2007).

Helft, Miguel. Google's Buses Help Its Workers Beat the Rush. *New York Times* (March 10, 2007), 4.

Poletti, Therese. Could Tech Shuttles Solve Bay Area's Transit Problem? *MarketWatch* (Jan 15, 2014), 2.

San Francisco County Transportation Authority, 2011. Strategic Analysis Report: The Role of Shuttle Services in San Francisco's Transportation System. 1-20.

Spivack, Cari. Worth the Drive. *Google: Office Blog* (September 13, 2004).

Thomas, Owen. Google's First Shuttle Bus Made Just Two Stops. *Business Insider* (October 12, 2012), 2.

Walker, Joseph. Google, Facebook, Genentech's Silicon Valley Bus Mania. *Technology and IT Jobs News and Advice* (April 3, 2012), 3.

C.3 Miscellaneous Sources

Thanks to David Dai and Danielle Weinzimmer, who privately provided me their maps of the shuttle routes. I would also like to thank an anonymous employee of a company under study who provided me with their shuttle stops. Lastly, I would like to thank Brendon Harrington of Google, who allowed me to interview him about his company’s bus operations.

C.4 Assumptions

- Apple shuttle service started in October 2007. Assume that stops observed in Winter 2009 were in place by then, as there is no evidence that stops ever changed.
- Google shuttle service started in September 2004, and expanded throughout 2005 and 2006 with very little documentation. Helft (2007) discusses riders moving to the Pacific Heights stops, which refers presumably the Van Ness corridor, in 2005. Anders (2007) mentions that the stops in Noe Valley had been in place since “early 2006”. Google maps from Fall 2006 confirm that there were Noe Valley stops, in addition to a stop at a park and ride center near Lake Merced in the extreme southwest of the city. These maps also indicate that there were no stops directly in the Van Ness Corridor, but spread elsewhere throughout Pacific Heights, in accordance with Helft. Earliest extant full map is from January 2009, and has the Van Ness Corridor, 19th Avenue Corridor, Noe Valley/Bernal Heights/Castro Districts and Haight-Ashbury stops in place.
 - Assume that the Lake Merced Stop was replaced by the 19th Avenue Stops and the Cow Hollow/Pacific Heights stops were replaced by the Van Ness Corridor Stops in May 2008.
 - Assume that Noe Valley, Bernal Heights, Haight-Ashbury, and Castro Valley stops were in place since February 2006.
- EA began their service in 2011 according to EA Staff (2012). Assume June 2011.

D Instrument Appendix

D.1 Data

A panel of all public bus, train, and light rail stops was assembled to provide an instrument for the placement of commuter shuttle stops. The agency directly provided a series of PDFs on all the SFMTA bus routes as of February 2008 and then a supplementary dataset of stop characteristics as of 2015 for cross-checking. From the 2008 PDFs, a list of stops was assembled and basic information such as stop length was collected. Data prior to 2008 was not forthcoming, but representatives of the agency claim that public bus stops were virtually unchanged between 2003 and 2008. Thus, I use bus zone distribution and characteristics from the 2008 PDFs as my fixed distribution of eligible zones for my instrument.

From `data.sfgov.org`, the main repository for any publicly available information published by the City and County of San Francisco, a March 2012 list of stops with their characteristics was made available. This was merged on to the 2008 routes data to confirm whether stops had a shelter or not. While it necessarily misses any changes in stop characteristics between February 2008 and March 2012, it is likely that most stops changed very little over time.⁴¹

The result is a panel of 880 eligible bus stops. The vast majority are concentrated around the central business district, but also throughout the city clustered around significant arteries. Figure 2 displays the locations of the eligible public bus stops, Caltrain, BART, and SFMTA light rail stops as of September 2004. The most significant change that occurs over the study’s time period is the opening of the SFMTA’s T Line on 3rd Street on the far east side of the city in January 2007.

The next three subsections outline how the bus zone characteristics were used to form the instrument described in Section 3.1.

⁴¹It is very hard to get a sense of how much change there is. A comparison of the March 2012 stops characteristics spreadsheet with the Fall 2015 spreadsheet is impractical because few fields in common are consistently filled in. However, stop length is reported in both, and the number of reported changes in length was less than 2%.

D.2 Eligible Bus Zone Variables per Building

Inspection of the shuttle stop distribution and the underlying transit and topographical infrastructure of the city quickly shows that the shuttles essentially integrated into pre-existing well-traveled corridors. Figure D2 overlays the commuter shuttle stop maps (streamlined from Figure 1), with the underlying transit options in San Francisco at each point in time. The first shuttle stops were along the BART system at Glen Park and a park and ride lot near Candlestick Park (Figure 1a), but the orientation of the system changed quickly once it became more popular. In Figure D2, an eligible bus zone is defined as being thoroughfare-adjacent if it is within a half-mile of a north-south thoroughfare. Figure D2 unmistakably shows that the shuttles clearly prioritized access to the thoroughfares in every phase of their growth.

Thus, the set of bus zone characteristics should include the basic variables (presence of a shelter, length), proximity variables (close to public transit, north/south thoroughfares), interactions between the basic variables and proximity variables, and then a last set of interactions that also interact with whether the bus zone is within a half-mile of a north/south thoroughfare.

The core of the spatial aspect of the instrument is computing the spatially-weighted average bus zone characteristics. Iterating over every in-sample building i , for each bus zone j within 1/2 mile of the building (804.67 meters) is assigned a spatial weight of:

$$spweight_{ij} = (1 - (Distance_{ij}/804.67)^2)^2.$$

Then characteristics like length and each bus zone's distance to relevant transit nodes are calculated as straightforward weighted averages. Shelter characteristics were computed in a similar fashion, after assigning to each bus zone an indicator for whether it had a shelter or not. For those bus zones where only one side of the street had a shelter, a value of 0.5 is given.

A variety of distance measures were included. These included a dummy on each bus zone on whether it was within a half mile of a given transit node (A north/south thoroughfare, MUNI Metro, BART, CalTrain, or BART/Caltrain combined) or the CBD; a reverse 2 kilometer measure similar to the shuttle measure used above; and the raw distance to the nearest transit stop, north-south thoroughfare, or CBD.

Another series of bus zone variables were added that included plausibly important interactions between characteristics. For example, it's likely that buildings where the average nearby bus zone was close to a north-south thoroughfare and had shelters was probably more likely to be near a shuttle stop than one whose average nearby bus zone had a shelter or was close to a north-south thoroughfare. Similarly true for bus zone length, where longer bus zones could conceivably allow a technology company to have more than one shuttle stop at a time. I thus include these spatially-weighted interaction terms with all of the bus zone distance variables and the key bus zone characteristics, bus zone length and shelter presence.

While bus zones are most heavily concentrated in some city quadrants versus others, these characteristics could plausibly found in bus zones throughout the city. However, the shuttles seemed more likely to favor certain areas over others, possibly due to exogenous factors such as physical infrastructure restraints, congestion fears, or particularly zealous city council members in a given ward. Another round of variables is thus introduced that interacts the characteristics above with measures for how far east or north the average nearby bus zone is.

These measures are created via the Universal Transverse Mercator (UTM)'s easting and northing coordinates. These are used in place of latitude and longitude in UTM projected coordinate system, which overlays the earth with unique ellipsoid zones (San Francisco is in UTM Zone 10) so that distances on the globe can be measured directly in meters. Easting in this case is how many kilometers east the point is from the zone's point of origin, which is the intersection between the zone's central meridian and the Equator. Likewise, northing is how many meters north the point is from the zone's point of origin. An arbitrary point in

the ocean just southwest of the peninsula was chosen as a new point of origin (specifically Easting=543500, Northing=4173500 in UTM Zone 10) and the easting and northing coordinates were rescaled and converted to kilometers so that interaction terms would not become arbitrarily large. There is one series of interactions for an "easting" measure, which is defined as the distance in kilometers from the longitudinal coordinate above and the average nearby bus zone longitudinal coordinate; a similar "northing" measure between the latitudinal coordinate and the average nearby bus zone latitudinal coordinate; a "hypotenuse" measure that is the crow-fly distance in kilometers between the above point and the average nearby bus zone; and a "central" measure that is the crow-fly distance in kilometers between the average nearby bus zone and the geographic center of the city (excluding Treasure and Yerba Buena Islands).

The way to think about these interactions is that they allow the LASSO algorithm to say that certain stop characteristics may have been favored overall, and then as the system spread west (or south) then stops with those favored characteristics would have been the most likely to be subsequently chosen along the given gradient. This permits a straightforward interpretation of the results. For example, the expectation is that thoroughfare-adjacent bus zones in the east (having a high "east" value) are going to be assigned the highest probability of getting early shuttle coverage. As more western stops get selected for a shuttle, the average east value of a shuttle stop will decline over time. Thus, the expected coefficient on the thoroughfare-adjacent flag interacted with the east value will be negative.

Rounding out the pool of candidate first-stage regressors is to then add for all of the variables described above two new sets of variables: one interacted with the total number of nearby bus zones and then another set interacted with the minimum distance from the building to a bus zone. While this list is not exhaustive, it yielded 1,578 candidate first-stage regressors, which did prove sufficient for generating strong first stages.

D.3 Introducing Time Variation

The absence of time variation is the most important obstacle to overcome when constructing this instrument, because the bus zone distribution is essentially static. Even if it were not, valid concerns would exist that changes in bus zones were done by the city in response to changes in commuting demand to Silicon Valley. I thus interact the constructed bus zone variables given above with the time series of Google’s stock, (ticker symbol GOOG), to model when demand shocks from tech workers. The value of GOOG is set to zero prior to August 2004 (the date of the IPO) and normalized prices thereafter to the average closing price for August 2004 (\$50.82). In spirit, this is similar to constructing a Bartik (1991) instrument, whereby the initial distribution of a covariate are interacted with an exogenous-to-local-conditions timeseries to instrument for where shuttle stop demand was most likely to cluster based on changes in demand by Google employees.

Figure D1 shows the normalized changes in GOOG from July 2003 to December 2013 along with the fraction of units that were within a 1/2 mile of a shuttle stop. Unsurprisingly, large changes in GOOG often precede expansions of the shuttle network, even those driven by other companies.

D.4 Selected Bus Zone Characteristics For Shuttle First-Stages in Eviction Regressions

Since the first stage variables change from regression to regression, and their interpretations are largely immaterial to testing the hypothesis, I present the results in this section. Included here are first-stage results covering Tables 3 and 7. For Tables 6, 8, 9, 10, first stage results are not presented largely to keep this appendix relatively brief. Table 6’s results are not used to drive the main conclusions, and Tables 8-10 do not include different selected variables than Table 7, but simply include the extra interaction terms with $\mathbb{1}\{\text{RC}\}$.

The variable nomenclature here is very similar to the rest of the paper, with some

simplifications. “min BZ Dist” is the shortest distance between the building/condo and all nearby bus zones. All variables representing averaged bus zone distances (in km) from the eligible zones and the city’s features are prefixed by “Dist”, with “Dist to Center” referring specifically to the distance to the geographic center of the city described in Section D.2. “Hypotenuse” refers to the average Euclidean distance from the nearby bus zones to the offshore point used to calculate Easting and Northing. Variables suffixed with “2 km” are variables constructed like the Shuttle 2km measure (described in Section 3). “Total BZs” refers to the count of nearby eligible BZs.

Table D1 shows the first-stage results of the hedonic regressions for the Shuttle 2km variable. The LASSO algorithm selected 64 different variables across six specifications to predict *Shuttle_2km*, but for interpretation purposes, I focus here on the 15 variables that were selected across all specifications. The first such variable is “min BZ Dist $\times$ $\mathbb{1}\{1/2 \text{ Mile to Light Rail}\}$ ”, which has a uniformly positive coefficient. This means in the aftermath of GOOG demand shocks, buildings close to light rail stops, particularly those a bit further away from the nearest eligible bus zone, were more likely to be close to a shuttle stop. Another variable selected across specifications is “Northing $\times$ min BZ Dist $\times$ Dist to Center $\times$ Light Rail 2km”. This variable is negative across most specifications, and essentially says that buildings in the city’s north, who were close to a light rail stop, but far from the geographic center and the nearest bus zone is relatively far are less likely to be close a shuttle stop after a GOOG-induced demand shock. Similarly, buildings in the far east of the city, close to many eligible bus zones, but are similarly far from the city center and within a half mile of either the BART/CalTrain were also less likely to be close to a shuttle stop after a GOOG-induced demand shock.

What both variables reveal is how the instrument derives its power. Within broad geographical areas (North vs East), or close versus near to the geographical center, there is variation building-by-building in terms of proximity to other features that it picks up on. This helps reinforce the general intuition behind constructing the instrument this way, with

several levels of interactions that the algorithm can then detect perhaps better than a human researcher could.

Bus zone physical characteristics are also important stop placement predictors, although not nearly as many were consistently chosen as location-oriented variables. “Length $\times$ BART/CalTrain 2km” and “Length $\times$ Light Rail 2km” were also selected across all six specifications, with the former significant in 5 of these. The coefficient is positive in all cases, which indicates that buildings close to long bus zones near the BART or CalTrain were much more likely to be near a stop after a GOOG-induced demand shock.

One problem with the hedonic sample is that it is too small to saturate with fixed effects, which could be used to better control for neighborhood- or transit-specific trends that might strongly co-vary with the candidate instruments. Further, because it is a repeated cross-sections sample, I cannot control for location fixed effects that would absorb all of the observed and unobserved location attributes that I instead have to try to control for directly. The panel set-up for the evictions, conversions, and permit filing regressions can sidestep both of these problems. One major expected change would be that the number of selected variables would greatly shrink, which is exactly what occurred. Table D2 bears out that being able to add panel fixed effects and richer time/neighborhood fixed effects greatly reduces the set of selected variables.

The results highlight how important these fixed effects are. The least saturated version (Column (1)) selected 5 variables. Adding neighborhood-by-year fixed effects reduces the selected variables to 4, with no selection overlap between Columns (1) and (2). Column (3) has two and Columns (4) and (5) has three. Columns (3)-(5) all control for both neighborhood and transit corridor-specific yearly shocks, with (4) and (5) including neighborhood-specific time trends, either linear (Column (4)) or neighborhood-specific GOOG shocks (Column (5)).

Two predictors were selected across Columns (3)-(5): “ $\mathbb{1}\{1 \text{ Mile of Thrufare}\} \times \text{Dist to BART/Caltrain}$ ” and “ $\text{min BZ Dist} \times \mathbb{1}\{1 \text{ Mile of Thrufare}\} \times \text{Dist to BART/Caltrain}$ ”. Both

coefficients are positive and significant at 1% level. The interpretation of the first is that for buildings whose nearby bus zones are within 1 mile of a north/south thoroughfare, those relatively further from the BART/CalTrain were more likely to wind up close to a shuttle stop after a GOOG-induced demand shock. The second variable further says that distance to the closest nearby bus zone also further strengthens the chance the building is close to a stop.

In the last two specifications (Column (4) and (5)), an additional variable was selected: “min BZ Dist $\times \mathbb{1}\{1/2 \text{ Mile to Light Rail}\}$ ”. The interpretation here is that for buildings whose nearby bus zones were within a half mile of a shuttle stop, those whose nearby bus zones were somewhat further away were more likely to be closer to a stop.

TABLE D1
Post-LASSO First-Stage Instrument Regression Results, July 2003-December 2013

	Without Historical Condos			With Historical Condos		
min BZ Dist $\times$ 1{1/2 Mile to Light Rail}	0.00453 (0.00835)	0.02335** (0.01048)	0.00802 (0.00781)	0.01041 (0.01174)	0.01811* (0.01081)	0.00358 (0.00797)
min BZ Dist $\times$ Dist to Center $\times$ Thrufare 2km	-0.00164* (0.00089)	-0.00075 (0.00104)			0.00016 (0.00074)	
Northring $\times$ min BZ Dist $\times$ Dist to Center $\times$ BART/CalTrain 2km	-0.00001 (0.00021)		0.00063 (0.00069)	-0.00013 (0.00021)	-0.00032 (0.00020)	-0.00026 (0.00018)
Northring $\times$ min BZ Dist $\times$ Dist to Center	-0.00014	0.00008	-0.00011	-0.00021**	-0.00032***	-0.00026***

SOURCE: Table D1 show the coefficient estimate for predicting how far a condominium is from the nearest stop's 2km outer ring generated from the method of Belloni et al. (2016)

* p<0.10, ** p<0.05, *** p<0.01, presuming standard method for computing confidence intervals.

TABLE D1
Post-LASSO First-Stage Instrument Regression Results, July 2003-December 2013

	Without Historical Condos			With Historical Condos		
×Light Rail 2km	(0.00010)	(0.00044)	(0.00010)	(0.00009)	(0.00010)	(0.00009)
min BZ Dist×Dist to Center×Dist to Thrufare	0.00038				-0.00047	-0.00053
	(0.00056)				(0.00047)	(0.00038)
Northring×min BZ Dist×Dist to Center	-0.00020*	0.00001	-0.00017*			
×Dist to Thrufare	(0.00010)	(0.00006)	(0.00009)			
Northring×min BZ Dist×Dist to Center	-0.00001	-0.00018***	-0.00003		-0.00010***	-0.00010***
×Dist to BART/CalTrain	(0.00005)	(0.00006)	(0.00004)		(0.00004)	(0.00003)

SOURCE: Table D1 show the coefficient estimate for predicting how far a condominium is from the nearest stop's 2km outer ring generated from the method of Belloni et al. (2016)

* p<0.10, ** p<0.05, *** p<0.01, presuming standard method for computing confidence intervals.

TABLE D1
Post-LASSO First-Stage Instrument Regression Results, July 2003-December 2013

	Without Historical Condos		With Historical Condos		
min BZ Dist $\times$ length $\times$ Dist to BART	-0.00248	-0.00118	-0.00056	-0.00082	0.00006
	(0.00263)	(0.00252)	(0.00372)	(0.00568)	(0.00355)
min BZ Dist $\times \mathbb{1}\{\text{Shelter}\}$	0.00751		0.01142**		
	(0.00469)		(0.00465)		
Easting $\times$ Total BZs $\times$ Dist to Center	-0.00020***				
$\times \mathbb{1}\{1/2 \text{ Mile to Light Rail}\}$	(0.00004)				
Northing $\times$ Total BZs $\times \mathbb{1}\{1 \text{ Mile of Thrufare}\}$	-0.00000	-0.00000	-0.00000	-0.000003*	-0.00000

SOURCE: Table D1 show the coefficient estimate for predicting how far a condominium is from the nearest stop's 2km outer ring generated from the method of Belloni et al. (2016)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, presuming standard method for computing confidence intervals.

TABLE D1
Post-LASSO First-Stage Instrument Regression Results, July 2003-December 2013

	Without Historical Condos			With Historical Condos		
×Dist to Center× $\mathbb{1}\{1/2 \text{ Mile to BART/CalTrain}\}$	(0.00000)		(0.00000)	(0.00000)	(0.000002)	(0.00000)
Easting×Total BZs×Dist to Center	0.00015***	-0.00002**	-0.00005***	-0.00004***	-0.00005***	-0.00004***
× $\mathbb{1}\{1/2 \text{ Mile to BART/CalTrain}\}$	(0.00004)	(0.00001)	(0.00001)	(0.00001)	(0.00001)	(0.00001)
Northing×Total BZs×Dist to Center	0.00001			-0.00001	0	
× $\mathbb{1}\{1 \text{ Mile of Thrufare}\}$ ×BART/CalTrain 2km	(0.00001)			(0.00001)	(0.00001)	
Northing×Total BZs×Dist to Center	0.00011	-0.00021	-0.00002	-0.00040**	-0.00028	-0.00023
×Light Rail 2km	(0.00024)	(0.00020)	(0.00020)	(0.00016)	(0.00018)	(0.00016)

SOURCE: Table D1 show the coefficient estimate for predicting how far a condominium is from the nearest stop's 2km outer ring generated from the method of Belloni et al. (2016)

* p<0.10, ** p<0.05, *** p<0.01, presuming standard method for computing confidence intervals.

TABLE D1
Post-LASSO First-Stage Instrument Regression Results, July 2003-December 2013

	Without Historical Condos			With Historical Condos		
Northings×Total BZs×Dist to Center	-0.00003***	-0.00002**	-0.00001	-0.00001	-0.00001	-0.00001
× $\mathbb{1}\{1 \text{ Mile of Thrufare}\}$ ×Light Rail 2km	(0.00001)	(0.00001)	(0.00001)	(0.00001)	(0.00001)	(0.00001)
Easting×Total BZs×Dist to Center	-0.00050**	-0.00026	-0.00032**	-0.00010	-0.00012	-0.00021*
×BART/CalTrain 2km	(0.00025)	(0.00018)	(0.00015)	(0.00018)	(0.00019)	(0.00013)
Easting×Total BZs×Dist to Center	0.00010	-0.00076	0.00113*	0.00028	-0.00087	0.00088
×CalTrain 2km	(0.00198)	(0.00175)	(0.00059)	(0.00187)	(0.00164)	(0.00068)
Total BZs×Length×CalTrain 2km	0.19216	0.35514**		0.14100	0.33549**	

SOURCE: Table D1 show the coefficient estimate for predicting how far a condominium is from the nearest stop's 2km outer ring generated from the method of Belloni et al. (2016)

* p<0.10, ** p<0.05, *** p<0.01, presuming standard method for computing confidence intervals.

TABLE D1
Post-LASSO First-Stage Instrument Regression Results, July 2003-December 2013

	Without Historical Condos			With Historical Condos		
	(0.20531)	(0.14480)		(0.16302)	(0.13694)	
Total BZs×Dist to Center×Dist to Thrufare	-0.00129***		-0.00113**			
	(0.00044)		(0.00050)			
Total BZs×Dist to Center×Dist to CBD	-0.00022	-0.00135***	0.00005	-0.00146***	-0.00183***	-0.00140***
	(0.00041)	(0.00040)	(0.00041)	(0.00034)	(0.00037)	(0.00034)
Easting×Total BZs×Dist to CBD	0.00080	0.00227***	0.00093	0.00169***	0.00255***	0.00152**
	(0.00090)	(0.00059)	(0.00067)	(0.00049)	(0.00061)	(0.00065)

SOURCE: Table D1 show the coefficient estimate for predicting how far a condominium is from the nearest stop's 2km outer ring generated from the method of Belloni et al. (2016)

* p<0.10, ** p<0.05, *** p<0.01, presuming standard method for computing confidence intervals.

TABLE D1
Post-LASSO First-Stage Instrument Regression Results, July 2003-December 2013

	Without Historical Condos			With Historical Condos		
Nothing×Total BZs×Dist to Center	-0.00015**			-0.00014**		
×Dist to CBD	(0.00007)			(0.00007)		
Nothing×Total BZs×Dist to CalTrain	0.00219**	0.00036	0.00189***	0.00123***	0.00043	0.00165***
	(0.00091)	(0.00049)	(0.00053)	(0.00038)	(0.00053)	(0.00040)
Nothing×Total BZs×Dist to Center	-0.00002	-0.00001	-0.00001	-0.00003***	-0.00001	-0.00003***
× $\mathbb{1}\{\text{Shelter}\}$	(0.00001)	(0.00001)	(0.00001)	(0.00001)	(0.00001)	(0.00001)
$\mathbb{1}\{1/2 \text{ Mile to Light Rail}\}$	0.00840***			0.00574***	0.00540***	0.00624***

SOURCE: Table D1 show the coefficient estimate for predicting how far a condominium is from the nearest stop's 2km outer ring generated from the method of Belloni et al. (2016)

* p<0.10, ** p<0.05, *** p<0.01, presuming standard method for computing confidence intervals.

TABLE D1
Post-LASSO First-Stage Instrument Regression Results, July 2003-December 2013

	Without Historical Condos			With Historical Condos		
Easting×Thrufare 2km	(0.00190)		(0.00158)	(0.00196)	(0.00175)	(0.00168)
	-0.00274	-0.00294	-0.00019	-0.00021	0.00163	0.00018
	(0.00211)	(0.00218)	(0.00146)	(0.00226)	(0.00212)	(0.00126)
Easting×Dist to Center×Thrufare 2km	0.00013	0.00014		0.00010	0.00006	
	(0.00012)	(0.00012)		(0.00013)	(0.00014)	
Length×Thrufare 2km	0.03845***	0.03200*	0.03177**	0.01797		0.02676**
	(0.01274)	(0.01664)	(0.01435)	(0.01272)		(0.01275)

SOURCE: Table D1 show the coefficient estimate for predicting how far a condominium is from the nearest stop's 2km outer ring generated from the method of Belloni et al. (2016)

* p<0.10, ** p<0.05, *** p<0.01, presuming standard method for computing confidence intervals.

TABLE D1
Post-LASSO First-Stage Instrument Regression Results, July 2003-December 2013

	Without Historical Condos			With Historical Condos		
Length×BART/CalTrain 2km	0.02000	0.03106**	0.04182***	0.02733**	0.02835**	0.03515***
	(0.01224)	(0.01318)	(0.01107)	(0.01161)	(0.01265)	(0.00997)
Easting×Light Rail 2km	0.00466**	0.00602**	0.00390*	0.00446**	0.00296	0.00327*
	(0.00223)	(0.00239)	(0.00229)	(0.00194)	(0.00198)	(0.00195)
Length×Light Rail 2km	0.00655	0.01127	0.00441	0.02180	0.03410**	0.01340
	(0.01537)	(0.01417)	(0.01563)	(0.01467)	(0.01635)	(0.01351)
Northing×Dist to Light Rail	-0.00007		0.00193***	0.00136***	0.00145***	0.00189***

SOURCE: Table D1 show the coefficient estimate for predicting how far a condominium is from the nearest stop's 2km outer ring generated from the method of Belloni et al. (2016)

* p<0.10, ** p<0.05, *** p<0.01, presuming standard method for computing confidence intervals.

TABLE D1
Post-LASSO First-Stage Instrument Regression Results, July 2003-December 2013

	Without Historical Condos		With Historical Condos		
	(0.00083)		(0.00023)	(0.00045)	(0.00042)
Northring×Dist to Center×Dist to Light Rail	0.00010***	0.00008***			(0.00026)
	(0.00003)	(0.00002)			
Easting×Dist to Center× $\mathbb{1}\{1 \text{ Mile of Thrufare}\}$	-0.00000				
×Dist to Light Rail	(0.00000)				
Dist to Center×Dist to Thrufare	-0.00008		-0.00018		
	(0.00029)		(0.00032)		

SOURCE: Table D1 show the coefficient estimate for predicting how far a condominium is from the nearest stop's 2km outer ring generated from the method of Belloni et al. (2016)

* p<0.10, ** p<0.05, *** p<0.01, presuming standard method for computing confidence intervals.

TABLE D1
Post-LASSO First-Stage Instrument Regression Results, July 2003-December 2013

	Without Historical Condos		With Historical Condos	
Northing×Dist to Thrufare	0.00104		0.00153*	
	(0.00078)		(0.00085)	
Easting×Dist to Thrufare	0.00080		0.00086	0.00050
	(0.00082)		(0.00093)	(0.00044)
Easting×Dist to CBD	0.00038			
	(0.00035)			
$\mathbb{1}\{1 \text{ Mile of Thrufare}\} \times \text{Dist to BART/CalTrain}$	0.00070***	0.00016	0.00052***	

SOURCE: Table D1 show the coefficient estimate for predicting how far a condominium is from the nearest stop’s 2km outer ring generated from the method of Belloni et al. (2016)

* p<0.10, ** p<0.05, *** p<0.01, presuming standard method for computing confidence intervals.

TABLE D1
Post-LASSO First-Stage Instrument Regression Results, July 2003-December 2013

	Without Historical Condos			With Historical Condos		
	(0.00015)	(0.00019)	(0.00011)			
Length×Dist to BART/CalTrain	-0.00592***	-0.00207**	-0.00276*			
	(0.00166)	(0.00103)	(0.00150)			
min BZ Dist× $\mathbb{1}\{1/2 \text{ Mile to Light Rail}\}$	0.00453	0.02335**	0.00802	0.01041	0.01811*	0.00358
	(0.00835)	(0.01048)	(0.00781)	(0.01174)	(0.01081)	(0.00797)
min BZ Dist×Dist to Center		-0.00261				
×BART/CalTrain 2km		(0.00191)				

SOURCE: Table D1 show the coefficient estimate for predicting how far a condominium is from the nearest stop’s 2km outer ring generated from the method of Belloni et al. (2016)

* p<0.10, ** p<0.05, *** p<0.01, presuming standard method for computing confidence intervals.

TABLE D1
Post-LASSO First-Stage Instrument Regression Results, July 2003-December 2013

	Without Historical Condos	With Historical Condos
min BZ Dist×Dist to Center×Light Rail 2km	-0.00407 (0.00406)	
min BZ Dist×Dist to Center×BART 2km	-0.00710 (0.00477)	
Easting×min BZ Dist×Dist to Center × $\mathbb{1}\{1 \text{ Mile of Thrufare}\}$ ×Dist to CBD	0.00001 (0.00001)	
min BZ Dist×Length×Dist to CBD	0.00075	

SOURCE: Table D1 show the coefficient estimate for predicting how far a condominium is from the nearest stop's 2km outer ring generated from the method of Belloni et al. (2016)

* p<0.10, ** p<0.05, *** p<0.01, presuming standard method for computing confidence intervals.

TABLE D1
Post-LASSO First-Stage Instrument Regression Results, July 2003-December 2013

	Without Historical Condos	With Historical Condos		
	(0.00201)			
Easting×Total BZs×Dist to Center	0.00000	0.00001***	0.00001***	0.00001***
× $\mathbb{1}\{1 \text{ Mile of Thrufare}\}$ ×Dist to BART	(0.00000)	(0.00000)	(0.00000)	(0.00000)
Northing×Dist to CalTrain	0.00099**			
	(0.00038)			
Northing×min BZ Dist×Dist to Center		-0.00019*	-0.00015*	-0.00019**
×Thrufare 2km		(0.00010)	(0.00008)	(0.00009)

SOURCE: Table D1 show the coefficient estimate for predicting how far a condominium is from the nearest stop’s 2km outer ring generated from the method of Belloni et al. (2016)

* p<0.10, ** p<0.05, *** p<0.01, presuming standard method for computing confidence intervals.

TABLE D1
Post-LASSO First-Stage Instrument Regression Results, July 2003-December 2013

	Without Historical Condos	With Historical Condos	
Hypotenuse×min BZ Dist×Dist to Center	-0.00087		
×BART/CalTrain 2km	(0.00062)		
Length×CBD 2km	-0.02197		
	(0.01429)		
Easting×min BZ Dist×Dist to Center		-0.00101	-0.00088
×BART 2km		(0.00074)	(0.00074)
Hypotenuse×min BZ Dist×Dist to Center		-0.00005	-0.00005

SOURCE: Table D1 show the coefficient estimate for predicting how far a condominium is from the nearest stop's 2km outer ring generated from the method of Belloni et al. (2016)

* p<0.10, ** p<0.05, *** p<0.01, presuming standard method for computing confidence intervals.

TABLE D1
Post-LASSO First-Stage Instrument Regression Results, July 2003-December 2013

	Without Historical Condos	With Historical Condos
×Dist to BART/CalTrain	(0.00004)	(0.00003)
min BZ Dist×Length×Dist to BART/CalTrain	-0.00139	
	(0.00538)	
Total BZs×Length×CBD 2km	-0.00782	
	(0.01221)	
Dist to Center×1{1 Mile of Thrufare}	0.00007	-0.00001
×Dist to Light Rail	(0.00006)	(0.00002)

SOURCE: Table D1 show the coefficient estimate for predicting how far a condominium is from the nearest stop’s 2km outer ring generated from the method of Belloni et al. (2016)

* p<0.10, ** p<0.05, *** p<0.01, presuming standard method for computing confidence intervals.

TABLE D1
 Post-LASSO First-Stage Instrument Regression Results, July 2003-December 2013

	Without Historical Condos	With Historical Condos
Easting× $\mathbb{1}\{1 \text{ Mile of Thrufare}\}$ ×Dist to Light Rail	-0.00009 (0.00010)	
Hypotenuse× $\mathbb{1}\{1 \text{ Mile of Thrufare}\}$ ×Dist to CalTrain	-0.00001 (0.00002)	-0.00001 (0.00001)
min BZ Dist×Dist to Center×Dist to CBD		0.00015 (0.00010)
Easting×min BZ Dist× $\mathbb{1}\{1 \text{ Mile of Thrufare}\}$		0.00012

SOURCE: Table D1 show the coefficient estimate for predicting how far a condominium is from the nearest stop’s 2km outer ring generated from the method of Belloni et al. (2016)

* p<0.10, ** p<0.05, *** p<0.01, presuming standard method for computing confidence intervals.

TABLE D1
Post-LASSO First-Stage Instrument Regression Results, July 2003-December 2013

	Without Historical Condos	With Historical Condos
$\times \text{Dist to BART}$		(0.00007)
$\mathbb{1}\{1 \text{ Mile of Thrufare}\} \times \text{Dist to CalTrain}$		0.00005
		(0.00015)
$\text{Northing} \times \text{Dist to Center} \times \text{Dist to CalTrain}$		0.00003**
		(0.00002)
$\text{Easting} \times \min \text{ BZ Dist} \times \text{Dist to Center}$		-0.00001
$\times \mathbb{1}\{1 \text{ Mile of Thrufare}\} \times \text{BART 2km}$		(0.00002)

SOURCE: Table D1 show the coefficient estimate for predicting how far a condominium is from the nearest stop's 2km outer ring generated from the method of Belloni et al. (2016)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, presuming standard method for computing confidence intervals.

TABLE D1
Post-LASSO First-Stage Instrument Regression Results, July 2003-December 2013

	Without Historical Condos			With Historical Condos		
Easting×min BZ Dist× $\mathbb{1}\{1 \text{ Mile of Thrufare}\}$						0.00007
×Dist to CBD						(0.00004)
Northring×Total BZs×Dist to CBD						0.00037
						(0.00067)
Adj R ²	0.85	0.85	0.95	0.85	0.85	0.95
N	24,975	24,975	24,975	26,884	26,884	26,884

SOURCE: Table D1 show the coefficient estimate for predicting how far a condominium is from the nearest stop’s 2km outer ring generated from the method of Belloni et al. (2016)

* p<0.10, ** p<0.05, *** p<0.01, presuming standard method for computing confidence intervals.

TABLE D2
No-Fault Eviction Sample First Stage Estimates of Shuttle 2km
Neighborhoods in Bottom 3 Quartiles, All Buildings, Jul 2003-Dec 2013

	(1)	(2)	(3)	(4)	(5)
min BZ dist×BZ Dist to Center×BZ Dist to CBD	-0.00010 (0.00012)				
BZ Dist to Center×CBD 2km	-0.00029*** (0.00004)				
min BZ dist×Dist to Center×Dist to Light Rail	-0.00094* (0.00047)				
min BZ dist×Dist to Center×Dist to BART/CalTrain	0.00013 (0.00027)				
Easting×Dist to Center×CBD 2km	0.00017** (0.00008)				
Total BZs×Dist to Center× $\mathbb{1}\{1 \text{ Mile of Thrufare}\}$ ×Dist to BART/CalTrain		-0.00005 (0.00004)			
Easting×Dist to Center× $\mathbb{1}\{1 \text{ Mile of Thrufare}\}$ ×Dist to BART/Caltrain		0.00001*** (0.00000)			
Easting×min BZ Dist×Dist to Center× $\mathbb{1}\{1 \text{ Mile of Thrufare}\}$ ×Dist to BART/Caltrain		0.00002*** (0.00000)			
$\mathbb{1}\{\text{Shuttle}\}$		0.00313*** (0.00091)			
$\mathbb{1}\{1 \text{ Mile of Thrufare}\} \times \text{Dist to BART/Caltrain}$			0.00068*** (0.00007)	0.00073*** (0.00007)	0.000739*** (0.00008)
min BZ Dist× $\mathbb{1}\{1 \text{ Mile of Thrufare}\} \times \text{Dist to BART/Caltrain}$			0.00146*** (0.00023)	0.00118*** (0.00027)	0.001198*** (0.00028)
min BZ Dist× $\mathbb{1}\{1/2 \text{ Mile to Light Rail}\}$				0.02401*** (0.00490)	0.02543*** (0.00485)
Adj R^2	0.83	0.93	0.93	0.93	0.97
F	69.90	60.42	33.56	60.41	19.99

SOURCE: Table D2 displays the first stage estimates for indicator for whether building i is within a half-mile of a shuttle stop. Point estimates are obtained via the procedure discussed in Section 3.1.

NOTE: For more information on the LASSO procedure, please refer to Belloni et al. (2012) and Belloni et al. (2016).

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, presuming standard method for computing confidence intervals.

Figure D1: Normalized Value of Google's Stock (GOOG) and the Fraction of Buildings within 1/2 Mile of a Shuttle Stop

SOURCE: Author's rendering based on GOOG price normalized to the IPO price and the average in-sample shuttle coverage.

Figure D2: Growth of Shuttle Stops Along Thoroughfare-Adjacent Bus Zones, December 2004-December 2013.

SOURCE: Author's rendering based on data about the shuttle network, and BART, Cal-Train, and SFMTA data.