

# The Long-Term Impacts of Rent Control

Matthew Gross<sup>\*†</sup>

*Draft date:* November 18, 2020

## Abstract

This paper is the first to provide estimates of the long-term benefits to children of growing up with rent controlled housing. Rent control is a common policy tool enacted to limit the growth of rents and allow tenants to remain in their homes for longer. Prior empirical research has mainly focused on rent control's impact on neighborhoods and housing markets while ignoring the potential long-term impacts of rent control for the people directly affected by the policy, particularly children. Using nearest neighbor matching at the census tract level, I estimate the effects of rent control on average long-term outcomes for children. Consistent with the literature, I first provide evidence that rent control leads to increases in the average housing tenure. I also show evidence implying that rent control improves economic mobility for those who receive it while also creating negative spillover effects for those that do not but live in cities with rent control policies. In tracts with a high proportion of rental units, rent control is associated with a \$1,300 increase in average tract-level income. Lastly, I report suggestive evidence that rent control has small long-term benefits for children at the bottom of the parent income distribution.

**Keywords:** Rent Control, Housing, Urban, Inequality, Anti Poverty

---

\*Department of Economics, University of Michigan. Contact: [mbgross@umich.edu](mailto:mbgross@umich.edu)

†I would like to thank Mike Mueller-Smith, Sara Heller, Brian Jacob and Charlie Brown for their time, effort and support throughout the writing process. I would also like to thank the attendees of the University of Michigan Labor Seminar for helpful comments at different stages of the research. The GM of GMM have been a constant source of support and friendship throughout my dissertation, and I would not have been able to do this without them.

# 1 Introduction

There is an ongoing national conversation about inequality of opportunity and the perception that economic mobility has become increasingly difficult for those born at the bottom of the income and wealth distribution. Research by Chetty *et al.* (2017) confirms that economic mobility in the United States decreased for birth cohorts between 1940 and 1980, and suggests that the socioeconomic status of parents is a growing predictor of a child's life outcomes. Researchers have identified early childhood as an especially important development period when well-timed interventions can mitigate some of the gaps in achievement between children of different income or wealth levels. Because of the cumulative nature of childhood development, relatively small childhood investments can have large long-run impacts, particularly when targeting children from marginalized populations (Cunha *et al.*, 2006). As a result, there is a large body of literature studying and measuring the effects of these interventions.<sup>1</sup>

Rent control, defined as government regulation of allowable rent, is a policy tool that is intended to transfer resources from landlords to renters. The goal of the policy is to make it easier for low-income tenants and families to remain in their current housing. In most cities with rent control, the policies were enacted in response to low rental vacancy rates, rising rental prices and the fear that without regulations, many tenants would face a heightened threat of eviction and homelessness. In the ideal scenario, rent control operates as a transfer from landlords to renters, both in the form of below-market rents and insurance against rent increases in the future.

This paper is organized around a central research question: how does rent control affect the long-term outcomes of children? In addition to the main question, I also seek to understand how this long-term impact varies by family income. Does rent control provide benefits to children growing up at the bottom of the income distribution?

While the short-term effect of rent control on impacted renters is likely positive, the long-term effects on children have never been studied. The literature on childhood interventions suggests a number of channels through which the child beneficiaries of rent control could have improved long-run outcomes as a result of growing up in rent controlled housing. Rent control can be thought of as a form of government mandated housing assistance, which has been shown in other contexts to increase earnings and decrease incarceration rates of impacted children (Andersson *et al.*, 2016). The benefits of rent control also include an income effect component, allowing families to shift expenditures from housing to other goods

---

<sup>1</sup>See Almond *et al.* (2018) for a recent review of the literature.

such as improved healthcare or education. Hoynes *et al.* (2016) shows that cash assistance to families with young children can improve child health. Lastly, the added housing security associated with rent control has the potential to reduce parental stress and the number of housing moves a child makes during childhood. Research shows that a mother's exposure to distressing news, such as a threat of eviction, can impact a newborn's birth weight (which is itself correlated with long-term health and achievement outcomes), while frequent childhood moves are known to be harmful to a child's academic performance and long-term development (Carlson, 2015; Wood *et al.*, 1993; South *et al.*, 2007).

Research also shows that one's childhood neighborhood has a causal exposure effect on long-run labor market and social outcomes, suggesting that the direction of the effect of rent control on outcomes may depend on the neighborhood (Chetty and Hendren, 2018a,b; Chyn, 2018, for example). If rent control leads children to have longer tenant durations in neighborhoods that provide negative exposure effects, then rent control could even lead to declines in long-term outcomes when compared with similar children who grew up in non-rent controlled cities.

Despite the potential effects of rent control on tenants, the economics literature has mainly focused on quantifying rent control's effect on the housing market and negative spillovers on neighborhoods. Rent control is often associated with a decrease in the quality of controlled housing (Gyourko and Linneman, 1990; Sims, 2007, for example) and a misallocation of tenants and apartments (Glaeser and Luttmer, 2003; Krol and Svorny, 2005). More recently, Autor *et al.* (2014, 2017) show that rent control has potential negative spillover effects, not only on the value of rent controlled units, but also on the value of neighboring properties that are not rent controlled. One of the large drivers of this negative capitalization is increased crime in areas close to rent controlled units suggesting that rent control suppresses gentrification. Diamond *et al.* (2019) shows that rent control in San Francisco leads landlords to remove their units from the rental market, thereby decreasing the supply of rental housing and ultimately leading to a more segregated and unequal housing landscape. Asquith (2018) confirms that landlords are more likely to convert rentals to owner occupied housing as the local price of housing increases.

While the costs of rent control are well-established, the benefits are much harder to quantify. Traditionally, economists used hedonic price regressions to quantify the benefit of rent control to renters by estimating the rent that would prevail in the absence of controls (Gyourko and Linneman, 1989). The difference between the controlled rent and the estimated market-rate rent is a measure of the compensating variation of the policy; however, a static measure of rent control benefits ignores the potential long-term benefits that the policy

confers. From a policy perspective, these long-term benefits are fundamental to determining whether rent control passes a cost-benefit analysis.

To answer the main research questions, I use a matching method to compare tract-level outcomes of areas that received rent control to counterfactual tracts that did not. This corresponds to the average treatment effect on the treated (ATT) of rent control on long-term tract outcomes. I measure these long-term outcomes using the publicly available Opportunity Insights data described in Chetty *et al.* (2018). This data set is constructed by linking children born between 1978 and 1983 to their childhood census tracts using restricted federal tax data, allowing the authors to measure the average exposure effects of neighborhoods on long-term outcomes such as economic mobility, employment, marriage, teen pregnancy and incarceration. For each tract and outcome variable, the data reports the unconditional mean over all children from the analysis cohort that lived in the tract between ages 6 and 23. As an example, the data on economic mobility reports the tract-level probability that a child will reach the top 20% of the income distribution in 2015. The realization of this variable gives the proportion of all children linked to a tract who have income in the top 20% of their birth cohort.

From 1970 to 1985, a number of municipalities enacted rent control legislation in response to rising inflation and low rental vacancy rates. I identify 116 cities in California, Massachusetts and New Jersey that codify new rent control laws during this time period. For each rent controlled city, I determine the census tracts that comprise the city, enabling me to map rent control laws to the outcome data set which is measured at the tract level. Rent control laws are passed by cities according to an unknown function of economic, housing, demographic, political and other local characteristics. Many of the predictors of rent control are also likely to be correlated with the long-term outcomes of children, implying that the difference in average outcomes between places with and without rent control is a biased estimate for the effect of rent control on long-term outcomes. Unfortunately, the direction of the bias is not immediately clear based on observable traits. While tracts that receive rent control have higher unemployment rates, minority population and single parent rates, for example, they also have higher average income, college attendance and property values.

To recover a causal estimate of rent control on the average long-term outcomes measured at the tract-level, I utilize a Mahalanobis distance nearest neighbor matching procedure to pair each treated census tract with a similar comparison tract that did not receive rent control. The Mahalanobis distance is a common metric used to measure distance between two points based on underlying covariate values. I estimate the Mahalanobis distance between tracts using data from multiple sources including the 1970 decennial census, the 1972

Census of Governments and county-level voting preferences in the 1968 presidential election. This data allows me to account for observable differences between census tracts and cities that enact rent control. The census data is reported at the tract-level and includes controls for demographic, housing, income and other characteristics that are predictive of receipt of rent control and correlated with the potential outcomes of children in a given tract. The Census of Governments data includes detailed municipal-level information on government expenditures and revenue. Under the strong ignorability assumption that conditioning on the observed covariates removes all confounding variation in the assignment of rent control, I am able to interpret the estimated average treatment effect as a causal parameter. To minimize imbalance on observed covariates in the matched sample, I implement a caliper match to prune treated observations that do not have comparison tracts with similar underlying covariate values. After pruning slightly more than half of all rent controlled tracts, I show evidence that the matching strategy does an adequate job of balancing the covariates for the remaining rent controlled tracts. When estimating the average effects of rent control on treated tracts, I also include bias adjustments for all covariates as proposed by Abadie and Imbens (2011).

My baseline estimates show that rent control leads to a 3.6% increase in the average time that children spend in a given census tract, which implies that rent control laws achieve their primary policy goal of allowing families to stay in their housing for longer. I also show that rent control leads to slight decreases in average tract-level economic mobility. Rent control is associated with a 5.4 and 3.5% decrease in the average probability of reaching the top 20% of the family and individual income distributions, respectively. The estimates show that rent control has a negligible effect on teen pregnancy, incarceration and employment. The standard errors on these baseline estimates are relatively large, leading to large confidence intervals. For the economic mobility outcomes, I can rule out effects of rent control that are larger than 10% and smaller than -20%.

The sample of tracts used to estimate the baseline results includes tracts with a low proportion of rental housing that are unlikely to have large direct effects resulting from rent control. I also generate matching estimators while limiting the sample to tracts where rental units represent at least 30% of all housing units. This removes approximately one half of the non-rent controlled tracts and 25% of the treated tracts from the sample. I find that when limiting the sample to high rental tracts, rent control increases the average time that a child spends in a given tract by 12%. This provides even stronger evidence that rent control leads families to remain in rental housing, since the tract-level effect is magnified in areas where we expect there to be more rent control.

Using the high rental sample, I also show that rent control increases the average probability of reaching the top 20% of the family (individual) income distribution by 5.9% (3.9%). Rent control also increases the tract-level average employment rate by 2.7%. Alternatively, rent control has a minimal effect on the average teen pregnancy rate while increasing the tract-level probability of being incarcerated during the 2010 census by 16.8%.

Assessing the results from the baseline and high rent samples, there is suggestive evidence that rent control does improve the average tract-level economic mobility in areas with a high percentage of renters, while also negatively impacting the economic mobility of non-rent controlled children living in cities with rent control. This result is consistent with the literature on the impact of government transfers on child outcomes as well as the literature showing that rent control is associated with negative spillovers on non-controlled housing. Furthermore, I use data from the 1980 to 2000 censuses to show that rent control leads to a decrease in the tract-level percentage of college attendance and an increase in the tract-level poverty rate and unemployment rate. According to Chetty *et al.* (2018), each of these demographic variables is associated with declines in the long-term outcomes of children.

An important shortcoming of the Opportunity Insights data is the fact that it is reported at the census tract level. Tract-level averages will include many children who did not grow up with rent control when aggregating over all children in a tract, making it more difficult to credibly measure small treatment effects. Assuming 20% of all children in a tract receive rent control, and rent control improves the average probability of reaching the top income quintile by 10% from a baseline of 0.1, the average tract level economic mobility rate would be  $(0.8 \times 0.1) + (0.2 \times 0.11) = 0.102 \approx 0.1$ . In this hypothetical example, the tract-level outcome is approximately unaffected by the existence of rent control despite the large benefit it provides to children who live in rent controlled units.

I attempt to alleviate this data concern by calculating the baseline nearest neighbor matching estimators while weighting the treated tracts by the number of children in each tract that are used to calculate the Opportunity Insights outcomes. This does not solve the measurement issues raised by using tract-level averages, though it does scale the treatment effect estimates by the number of children in each tract. In these weighted estimates, tracts with more children will play a larger role in determining the average tract-level effects of rent control. After implementing this weighting, the baseline results are unchanged. This implies that the average effect of rent control on a child growing up in a rent controlled tract is roughly equivalent to the average outcome of a child growing up in a similar tract.

Lastly, I utilize the Opportunity Insights data on predicted outcomes for children at the 25th and 75th percentiles of parental income to determine how rent control affects children

at the bottom and top of the parent income distribution. I generate estimates using both the baseline sample and the high-rent sample. In the baseline sample, rent control has a negative effect on the predicted economic mobility of children with parents at the 25th percentile of income distribution. By contrast, rent control has a minimal effect on the economic mobility of children with parents at the 75th percentile of the income distribution. In the high-rent sample, rent control leads to small and statistically significant improvements in the predicted economic mobility for children at the 25th percentile of the parent income distribution. For children at the 75th percentile of the family income distribution, rent control has a significant positive effect on predicted rates of reaching the top income quintile as adults. The results from the high rent sample suggest that rent control helps individuals at the bottom of the income distribution, though the effects are substantially stronger for children growing up at the top of the parent income distribution. These results are consistent with previous findings showing that rent control is poorly targeted to lower income families; however, I view the results from this exercise as merely suggestive and warranting future study with individual-level data to better grasp the heterogeneity of the effect of rent control on future outcomes by income levels.

This research adds to the economics literature on rent control by tracking the outcomes of people that are affected by the policy and quantifying the long-term benefits. This article is the first to estimate these benefits in a causal framework and will be of immediate interest to policymakers deciding whether rent control policies pass a cost benefit analysis. While the results from the high-rent sample suggest that there are positive long-term benefits for children growing up with rent control, future research can build on this work by generating more precise estimates of these effects and potentially utilizing alternative data sources to leverage individual-level variation in the assignment of rent control.

## 2 Relevant Literature

Rent control is a commonly studied topic in the economics literature going back to Grampp (1950) who argued strongly in favor of removing rent regulations to help avoid housing shortages and improve economic efficiency. This view is consistent with the implications of a simple supply and demand model which predicts that rent control leads to over-consumption and deadweight loss. For many years, a lack of natural experiments and suitable data prevented economists from estimating well-identified causal effects of rent control. As a result, there is a significant body of theoretical work exploring the implications of various rent control regimes (Fallis and Smith, 1984; McFarlane, 2003; Suen, 1989, for

example). In addition, the standard economic model's clear predictions of efficiency costs due to price controls may have led some economists to think that empirical research on this topic would be superfluous (Gyourko and Linneman, 1990).

There are a number of papers that attempt to quantify the costs of rent control on housing quality (Moon and Stotsky, 1993; Gyourko and Linneman, 1990; Sims, 2007, for example), generally showing that rent controlled units are maintained at a lower quality than they would be in the absence of price controls. Other research shows how rent control negatively affects housing prices for the controlled (Autor *et al.*, 2014) and uncontrolled stock (Fallis and Smith, 1984; Early, 2000). Lastly, there is a body of literature that attempts to characterize and quantify the costs of rent control that result from inefficiently long tenant durations (Krol and Svorny, 2005; Ault and Saba, 1990; Ault *et al.*, 1994) and the misallocation of tenants and apartments (Glaeser and Luttmer, 2003).

Measuring the benefits of rent control to renters can be challenging without longitudinal data. There is ample evidence that rent control is associated with increased tenant durations which implies that renters with rent control receive some benefit from the policy (Olsen, 1972; Ault *et al.*, 1994; Nagy, 1995; Munch and Svarer, 2002, for example). One common method used to estimate the size of the benefit in the absence of exogenous variation is to measure the difference between controlled rent and the predicted rent that would occur in the absence of controls. This can be done using the two-step method proposed by Gyourko and Linneman (1989) to estimate hedonic rent regressions of the uncontrolled rental stock on housing characteristics. These regressions are then used to predict what the rent would be at controlled units, conditional on observable characteristics. The difference between the predicted and actual rent can be thought of as the compensating variation or monetary value of rent control to the renter. In the second step, one can regress the compensating variation on tenant characteristics to determine how the benefits of rent control are distributed to different groups. Other papers that use this methodology include Gyourko and Linneman (1990), Ault and Saba (1990), Munch and Svarer (2002), and Early (2000). In general, these papers find that the benefits are not particularly well-targeted to the lower-income groups that price controls are intended to help.

In 1995, Massachusetts voters banned rent control in a closely contested statewide referendum, providing economists with a natural experiment to measure the effect of the end of rent control. Sims (2007) was the first to utilize this policy variation and found that rent control in Boston was associated with both decreases in the price of housing as well as housing quality. While his results indicate that rent control had no impact on the construction of new housing, he presents evidence that rent control decreased the value of neighboring,



unaffected housing stock, though by a relatively small amount. Autor *et al.* (2014) focus on Cambridge, Massachusetts and study both the direct effect on home values of rent decontrol, and the effects on housing values of homes that were never regulated. They find that rent control suppresses the value of controlled homes, and also find substantial neighborhood effects implying that rent controlled units also suppress the value of nearby unregulated units. The end of rent control in Cambridge caused nearly \$2 billion in housing value appreciation. Using a similar methodology in a follow-up paper, Autor *et al.* (2017) utilize detailed crime data from 1992 to 2005 to measure the effect that the end of rent control had on local crime rates. They conclude that areas with more rent controlled housing prior to decontrol saw larger decreases in crime rates than otherwise similar areas. This implies that the end of rent control had a significant effect on decreasing crime rates in Cambridge and accounts for 15% of the home value appreciation as a result of rent control found by Autor *et al.* (2014).

Building on the research using the 1995 Massachusetts rent decontrol natural experiment, Diamond *et al.* (2019) estimate a well-identified causal effect of rent control on tenants, landlords and inequality using a 1994 change in the San Francisco rent control regime. Prior to 1994, all buildings built before 1980 were subject to rent control except those that contained four or fewer units. In 1994, the small building exemption was removed such that all rental buildings with four or fewer units were now subject to rent control. Buildings with four or fewer units built after 1980 continued to be exempt from the rent control ordinance, providing a natural control group. Using a novel linkage, they collect address histories for San Francisco residents in addition to building and landlord information. The authors find that receipt of rent control led to a 15% increase in the duration of rental stays; however, they also find that rent control incentivizes landlords to remove units from the market, thereby decreasing the number of rental units. Unlike previous studies, they conclude that the benefits of rent control were well targeted to minorities, but that rent controlled units were more likely to be in neighborhoods with lower amenities (where the benefits of rent control are lower). Lastly, because landlords were more likely to remove rent controlled units in neighborhoods with more amenities, the authors argue that rent control has accelerated gentrification, inequality and rental prices in San Francisco. Asquith (2018) shows a similar result that San Francisco landlords react to increasing land values by removing tenants through no-fault evictions, allowing them to convert to non-rental uses such as condominiums.

The paper by Diamond *et al.* (2019) is the only research that attempts to track renters over time to measure the effect of rent control on the mobility and location decisions of tenants; however, we still do not know how rent control affects important long-term economic, labor

market and social outcomes for tenants. Measuring these outcomes is key to a more complete understanding of the costs and benefits of rent control. I fill a critical gap in the literature to date with evidence from a new intergenerational perspective on the consequences of rent control.

### 3 Rent Control Sites and Institutional Background

The rent control policies that I utilize in this paper were passed in the 1970s and early 1980s, and are considered part of the second generation of rent control in the United States; however, rent control policies in the U.S. date back to the end of World War I when housing shortages in a number of cities caused states to restrict evictions of soldiers and workers involved with the war effort.<sup>2</sup> During World War II, the Federal Emergency Price Controls Act (EPCA) subjected many aspects of the economy to price regulation including rental housing. These federal controls continued in modified form until 1952, though some areas and units that had been initially controlled by the EPCA were decontrolled before the end of federal rent regulation. The 1947 Housing and Rent Act gave states additional authority to either extend rent regulations or decontrol rents on their own. By 1948, 10 states had some form of rent control legislation, though by the mid 1950s, New York was the only remaining state with rent controlled housing (Lett, 1976).

In the 1970s, high levels of inflation and low rental vacancy rates in cities around the country led renter advocates to push for new laws to regulate the level and growth of rents. As a result, a number of states and municipalities began to implement new rent control regimes. States that added rent control during this wave include California, New Jersey, Maryland, Massachusetts, the District of Columbia and Alaska. I focus on the laws passed in cities in California, Massachusetts and New Jersey.<sup>3</sup> <sup>4</sup> In most of these places, the justification for passing rent control was that low rental vacancy rates coupled with large rent increases constituted an emergency that incentivized “rent gouging” and placed many residents at risk of eviction (Lett, 1976). Furthermore, evicted residents were more likely

---

<sup>2</sup>See Fogelson (2013) for a detailed historical account of New York City’s experience with rent control in the post World War I period.

<sup>3</sup>In Maryland, Takoma Park enacted rent control in 1981. Lett (1976) claims that a number of other counties implemented rent control in the early 1970s, though I have been unable to independently verify these laws. As a result, I drop Maryland from the analysis sample to ensure that I do not have measurement error in the treatment group.

<sup>4</sup>Washington D.C. implemented rent control in 1975 in the middle of a major, unrelated demographic shift. The population in Washington D.C. dropped 15% between 1970 and 1980 and declined from 800 thousand residents in 1950 to 570 thousand in 2000. These unrelated changes might add additional noise to the measured effects of rent control leading me to drop Washington D.C. from the analysis.

to end up homeless given inadequate supplies of rental housing. Many other cities and states considered implementing rent control during this time but had proposals fail to garner sufficient support.<sup>5</sup> In Maine, the state passed legislation enabling localities to implement rent control though none ended up being passed (Lett, 1976).

Rent control laws regulate the legal terms of rental agreements as well as the rent that a landlord can charge a tenant. In practice, there are many different ways that governments implement rent regulation. In the most restrictive cases, governments determine an exact price for rental housing or place a freeze on rents to prevent them from increasing for any reason.<sup>6</sup> In other cases, governments may place limits on the maximum possible rent increase, either through arbitrarily defined price ceilings or by tying rent increases to inflation. In these cases, rent control is only binding if the landlord would be able to raise rents above the government imposed limit in a competitive market. Another common aspect of rent control legislation is vacancy decontrol, which determines the rent that landlords are allowed to charge the next tenant after the previous tenant voluntarily vacates the rental. Depending on the law, landlords may be allowed to raise rents to market rates under full vacancy decontrol while in other cases, landlords may only be allowed to raise rents by a fixed percentage. Lastly, rent control legislation often places additional limits on evictions though the implementation of eviction restrictions varies widely by location.

In New Jersey, most rent control laws are based on the legislation enacted in 1972 by the municipal government of Fort Lee, which was the first city in New Jersey to implement rent control. Landlords quickly challenged the legality of the legislation, but in 1973, the State Supreme Court of New Jersey ruled that local governments were allowed to regulate local rent. After the court decision affirming the legality of local rent control, many New Jersey municipal governments followed Fort Lee's example and instituted their own regulations. By 1976, nearly 100 cities and townships had rent control laws. Although the laws in New Jersey are not identical, many are based on the original law from Fort Lee (Lett, 1976). In general, the laws set base rents at current (as of the date of enactment) levels and then tied allowable rent increases to inflation. The laws generally exempted small-scale landlords (usually owners of buildings with fewer than three rentals) from the law. Also, rental units constructed after rent control enactment were often exempt from the legislation to incentivize new housing, and landlords were given permission to raise rents by more than inflation in the event that they invested in capital improvements or if operating costs increased.

---

<sup>5</sup>For example, municipal rent control proposals in Colorado, Pennsylvania and Wisconsin were all considered and ultimately not approved during the early 1970s.

<sup>6</sup>For example, see Washington D.C.'s temporary rent freeze, Regulation 74 - 13 passed in 1974.

In Massachusetts, the state passed rent control enabling legislation in 1969 which allowed certain cities to pass rent control laws. Following this legislation, Boston, Somerville, Cambridge, Brookline and Lynn passed rent control laws which went into effect in 1970. Lynn and Somerville repealed rent control in 1974 and 1979 respectively. Under the laws, base rents were set at current levels and rents were allowed to increase to return a reasonable net operating income. Rent increases were also allowed for capital improvements and changes in operating expenses. New buildings and units in owner occupied houses were exempt, while all other extant rentals were subject to the law (Lett, 1976). In 1995, the voters of Massachusetts narrowly approved a referendum which made it illegal for cities to enact rent control legislation. Although Boston and Brookline loosened rent control restrictions prior to 1995, both still had a substantial percentage of units subject to control (Autor *et al.*, 2014). Cambridge still had a heavily controlled housing market at the time of the ballot initiative in 1995.

In California, Berkeley was among the first cities to pass rent control legislation in 1972; however, this law was ruled unconstitutional by the California Supreme Court. Starting in 1979, a number of large cities began passing rent control laws including San Francisco, Los Angeles and Oakland. By 1985, 12 cities in California had implemented rent control legislation. Rent control laws varied by the city; however, all cities were forced to adhere to both the Ellis Act passed in 1985 and the Costa-Hawkins act passed in 1995. The former allowed landlords to evict tenants if they wished to remove their rental housing from the market. This was passed by the state legislature in response to a State Supreme Court ruling which stated that cities could prevent landlords from evicting tenants even when the landlord wanted to occupy the house. This forced cities with strong restrictions to allow landlords the ability to exit the rental market, though the administration of this law varied by city. The Costa-Hawkins act forced cities to allow for vacancy decontrol after a renter leaves a rent controlled unit. This allowed base rents to rise to reflect market conditions after a tenant leaves, regardless of the rent that the previous tenant paid. In terms of exemptions, most new construction and small-scale rental buildings (1-3 units) were not subject to rent regulation.

In summary, rent control laws rolled out across California, Massachusetts and New Jersey in the 1970s and early 1980s. While these policies were not identical, each placed new restrictions on landlords that could have long-run consequences for tenants, especially children. Next I discuss the data I use to estimate these effects before presenting the key results.

## 4 Data

In the states that began passing rent control legislation during 1970s, the power to implement rent control devolved to local political units, meaning that rent control existed in some cities but not others. This was particularly true in California, Massachusetts and New Jersey, which are the three states I focus on to estimate the effect that rent control has on long-term outcomes of children. I follow Krol and Svorny (2005) in using Lett (1976) to collect information on local rent control laws passed in the 1970s, particularly in New Jersey and Massachusetts. This book includes a comprehensive list of cities that passed rent control by 1976, covering nearly all of the New Jersey and Massachusetts cities that added rent control. I supplement this resource with internet searches of legislative histories for large municipalities, particularly in California, to determine which cities added rent control legislation in the years following 1976.

Throughout the paper, I measure rent control as a binary variable and do not distinguish between rent control policies in Massachusetts, California and New Jersey. Though there are differences in the regulations across municipalities and states, the number of different cities with unique regulations makes it difficult to account for policy variation. Future research should attempt to study specific dimensions of rent control and the heterogeneity of treatment effects by rent control policy type. Despite this, there are reasons to believe that policies within states are fairly similar to each other. Most laws at this time were in reaction to low vacancy rates which allowed landlords to raise rents quickly. In New Jersey, rent control laws are enacted around the same time and are based on a law passed by the municipal government of Fort Lee. In Massachusetts, despite some differences between Boston, Brookline and Cambridge, all three cities had a substantial number of controlled units until 1995 when all rent control laws were invalidated by the statewide ballot initiative. Lastly, in California, the laws across cities had similar exemptions and rent increase mechanisms and were subject to statewide legislation that standardized vacancy decontrol and landlord exit.

I utilize geographic and shapefile data provided by the Census Bureau to identify census tracts that were subject to rent control. First, I merge a Census shapefile of incorporated cities from the 1990 census with a shapefile of the 2010 census tracts to create an overlap layer. Using this intersection, I determine the 2010 census tracts that comprise every city in the country. I then merge in the list of cities passing rent control between 1970 and 1985 to generate a binary variable for rent control status for each census tract in the United States. In New Jersey, there are a number of municipalities that passed rent control that do not appear in the list of census places. For these remaining locations, I use the more detailed

county maps provided at [www2.census.gov](http://www2.census.gov) to manually identify the tracts associated with each rent controlled city. This leaves me with a database of census tracts for California, Massachusetts and New Jersey along with a binary variable indicating whether the tract had rent control established between 1970 and 1985.

Since rent control is implemented by local elected representatives, the decision to enact these policies is likely dependent on underlying observable and unobservable city characteristics. In other words, rent control is not assigned randomly throughout the country, so comparing outcomes of rent controlled and uncontrolled cities is likely to be a biased measure of the effect of rent control. Instead, I utilize data from the 1970 Decennial Census as balancing covariates in a matching framework to estimate the effect of rent control on long-term outcomes. The Census data is pulled from the SocialExplorer website, which aggregates individual responses from the 1970 Census up to the census tract level. Census tract borders change over time, so I use the 1970 Census data that is reported at the 2010 census tract level to maintain a consistent measure of geography.<sup>7</sup>

The matching procedure also includes city-level data on municipal spending and revenues. The municipal tax and revenue data comes from the Government Finance Database described by Pierson *et al.* (2015).<sup>8</sup> The database compiles information from the Census of Governments beginning in 1967. In years ending in either a 2 or 7, the U.S. Census Bureau collects information on the finances of every incorporated government in the United States. Unfortunately, this full census did not begin until 1972, so I use the data collected from the 1972 census to ensure that I have maximum coverage of all cities in my sample.

The enactment of rent control is a local political decision. As a result, it is necessary to control for local political views when comparing places that did or did not have rent control. To account for local political differences, I use data on the county-level partisan vote shares for the presidential election of 1968. This data is collected by Clubb *et al.* (2006) and distributed by the ICPSR.

The data for long-term outcomes is described in Chetty *et al.* (2018) and is available for public download on the Opportunity Insights website.<sup>9</sup> This data combines multiple sources

---

<sup>7</sup>SocialExplorer uses the area interpolation method described in Logan *et al.* (2014) to convert 1970-1990 tracts to the 2010 tracts. Area interpolation assigns populations from one area to another based on area overlap and does not account for the distribution of population density within tracts. This is a potential source of error, particularly in tracts with changing borders and unequal distribution of population. To the best of my knowledge, no research has theorized the direction of the expected bias.

<sup>8</sup>The data is publicly available at <https://willamette.edu/mba/research-impact/public-datasets/>; however, I downloaded the data through the Inter-university Consortium for Political and Social Research (ICPSR) website (study number 37641): <https://www.icpsr.umich.edu/web/pages/ICPSR/index.html>.

<sup>9</sup><https://opportunityinsights.org>

of restricted government data to measure a series of financial, social, educational and other outcomes for children born between 1978 and 1983. Using federal income tax returns from 1989 to 2000, the authors identify all children who are listed as tax dependents and were born between 1978 and 1983. Next, they utilize the Census Bureau’s Protected Identification Key (PIK) to link these children to the 2000 and 2010 Decennial Census waves, 2000-2015 American Community Surveys and IRS income tax returns from 1989-2015.<sup>10</sup> The sample selected is representative of all children in the 1978 to 1983 birth cohort that were born in the United States or authorized immigrants and whose parents were either born in the U.S. or authorized immigrants.<sup>11</sup>

Once the sample is selected, the authors map children to the census tracts they grow up in through their age 23 year. A child born in 1983 can be linked to a particular tract through 12 distinct years of tax returns (1989, 1994-1995 and 1998-2006) between ages 6 and 23. Children born in 1978 are only linked to 7 years of tract data (1989, 1994-1995 and 1998-2001) between ages 11 and 23. For each 2010 census tract, the authors report the unconditional mean outcome value for all children linked to the tract. Children that appear in multiple tracts due to childhood moves are weighted to represent the relative time spent in each tract. As an example, a child born in 1983 who is linked to tract A in 6 years of tax returns and tract B for the other 6, would receive 0.5 weight in both tracts A and B when calculating tract level outcomes.

I focus on six census tract outcomes reported in the Opportunity Insights data: 1. the probability of reaching the top income quintile<sup>12</sup>; 2. the average income percentile<sup>13</sup>; 3. the probability of having a teen birth, 4; the probability of being in a correctional facility at the time of the 2010 Decennial Census, 5; the probability of having positive W2 earnings in 2015; 6. the probability of living in a low poverty neighborhood as an adult. These variables allow me to measure how rent control impacts long-term economic mobility and other important social outcomes.

In addition to the unconditional mean outcomes, Chetty *et al.* also report average predicted outcomes at the tract level for children at 5 different levels of the parent income

---

<sup>10</sup>The PIK is created using a probabilistic matching algorithm that is based on an individual’s Social Security Number, as well as name, date of birth and address. The PIK can be used to follow an individual across a number of Census Bureau, IRS and other governmental data sets.

<sup>11</sup>External validity is a common concern when using data limited to those who file tax returns. According to Chetty *et al.*, the sample used to create the public Opportunity Insights data is representative of the overall population covered by the American Community Survey and the Current Population Survey.

<sup>12</sup>The income quintile is measured relative to all other people born in the same year to account for rising expected earnings with age.

<sup>13</sup>Income percentiles are reported based on either the distribution of family income or individual income.

distribution. These fitted values are generated from a regression of individual outcomes on parental income level at the tract level. The regressions are estimated using all children linked to a particular tract (weighted for the number of linked years). I utilize the estimates at the 25th and 75th percentile of the parent income distribution to investigate whether rent control has a differential effect on children at the bottom or top of the distribution; however, it is important to emphasize that these estimates represent fitted values of a regression and may not reflect true outcomes of children. For example, a very wealthy tract may have relatively few parents at the bottom of the national income distribution. Running a regression of child outcomes on parental income percentile might predict that children with parents at the 25th percentile have positive outcomes in this hypothetical tract; however, this prediction is based entirely on a projection of children at the top of the parent income distribution. I only focus on predicted values from the 25th and 75th percentiles, ignoring estimates from the 1st, 50th and 100th percentiles. See appendix A for a slightly more technical description of the Opportunity Insight data.

## 5 Methodology

Rent control legislation is not randomly assigned to cities, but is instead implemented according to an unknown function of local housing, demographic and political characteristics. Tables 1 and 2 show summary statistics broken down by rent control status at the tract and municipal level respectively. In both tables, the treated group are the tracts and cities located in California, Massachusetts and New Jersey that receive rent control between 1970 and 1985, while the control groups represent tracts and cities from the remainder of the country that are located in incorporated cities with a population greater than 5,000 people and that are represented in the 1972 Census of Governments. I also remove observations from New York, Maryland and Washington D.C. due to a mix of timing issues (New York), rent control measurement ambiguity (Maryland), and likely confounding variation (Washington D.C.).

Out of 31,261 total census tracts, 2,444 are treated with rent control. These 2,444 treated tracts comprise the 99 cities that received rent control between 1970 and 1985, have populations over 5,000 and responded to the 1972 Census of Governments.<sup>14</sup> In California, 26% of the population lived in cities with rent control. In Massachusetts, 15% of the population lived in a rent controlled city while in New Jersey, over 48% of the state population lived in

---

<sup>14</sup>There are 16 small cities that enacted rent control between 1970 and 1985 but are not included in the Census of Governments. These are mostly located in New Jersey.



cities with rent control.

From Table 1, it is clear that tracts that receive rent control are fundamentally different than those that do not. The difference in average value between treatment and control groups is statistically significant for most of the covariates suggesting that rent control is not distributed quasi-randomly. Instead, places with rent control have a higher single parent rate, minority population and higher unemployment rate, which are all variables that are negatively correlated with long-term outcomes for children. On the other hand, tracts with rent control have higher college attendance rates, average incomes and home values which are correlated with improved child outcomes.

In Table 2, I report summary statistics at the city level comparing cities that enacted rent control to those that never implemented a rent control law. Not surprisingly, cities with rent control had lower vacancy rates and a higher percentage of rentals as a share of total units. These cities with rent control were more likely to be in counties that had higher share of votes for Hubert Humphrey, the Democratic candidate, in the 1968 presidential election. In addition, the cities with rent control have higher municipal revenue per capita, and spend a higher fraction of total expenditures on education, police and welfare compared to non-rent controlled cities. The results from Tables 1 and 2 show that assignment of rent control is correlated with observable characteristics that are also likely correlated with the long-term outcomes of children.

Each row of Table 3 represents an outcome variable of interest. We can see that the average fraction of years spent in a rent controlled tract is only slightly (0.4 percentage points) higher than the non-controlled tracts; however, on average, tracts with rent control are much more likely to report a higher probability of children living with their parents as adults and staying in the same tract or commuting zone as an adult. Rows 5 through 12 show that tracts with rent control report greater average economic mobility, lower teen pregnancy and incarceration rates and higher rates of employment and living in tracts with low poverty rates as an adult. In general, the naive treatment effects suggest that rent control improves long-term outcomes; however, the differences in pre-rent control characteristics, particularly at the tract-level, imply that a simple difference in outcomes between treated and control tracts is likely to be a biased measure of the effect of rent control. Unfortunately, it is not immediately clear which direction this bias shifts the naive estimates given the countervailing effects of the individual variable imbalances.

In many observational studies, treatment is not assigned randomly but is instead assigned according to some (unknown) function of observable and unobservable characteristics. I lean on Rubin's model of causal inference to formalize the analysis and provide theoretical

justification for the estimand of interest (Holland, 1986). In my setting, there is a rent control treatment  $T \in \{0, 1\}$  that is assigned to the population of census tracts of size  $N$ . I hypothesize that rent control treatment  $T$  has a causal effect on long-term average outcomes at the tract level, denoted  $Y(T)$ . The causal effect of  $T$  on  $Y_i$  for tract  $i$  can be measured as  $Y_i(T = 1) - Y_i(T = 0)$ . Aggregating up to the full sample of census tracts, the average treatment effect (ATE) can be written as  $\frac{1}{N} \sum_{i=1}^N [Y_i(1) - Y_i(0)]$ . Alternatively, the average treatment effect on the treated (ATT) which measures the effect of a treatment only on the treated tracts, is written as  $\frac{1}{N_1} \sum_{i=1}^{N_1} ([Y_i(1) - Y_i(0)] | T = 1)$ . Since we cannot observe the treated and untreated potential outcome at the same time, measurement of the causal effect of interest is reduced to a missing data problem. Unless otherwise noted, I estimate the average treatment effect on the treated throughout my analysis.

To bypass the missing data issue, I utilize techniques that balance the treatment and control groups on observable characteristics to find a suitable counterfactual and allow for improved estimates of the effect of rent control on outcomes. Rosenbaum and Rubin (1983) were among the first to formalize the framework for achieving causal estimates in observational studies by conditioning on a vector of control variables to facilitate matching. The assumptions required to identify the average treatment effect on the treated of rent control can be written as:<sup>15</sup>

$$(Y_0) \perp T | X, pr(T = 1 | X = x) < 1$$

where  $X$  is a vector of covariates and  $pr(T = 1 | X = x)$  is the probability that a tract is rent controlled conditional on any realization of the covariate vector. The second part of the assumption states that the covariates cannot perfectly predict treatment. Under this strong ignorability assumption, Rosenbaum and Rubin show that treatment effects can be recovered by matching observations of different treatment levels with the same value of the conditioning function based on  $X$ . In the original case, Rosenbaum and Rubin use an estimated propensity score; however, any function of  $X$  can be used. The intuition behind this result is that controlling for the covariates  $X$  is sufficient to make the treatment assignment random. For this to be true, there must not be any unobserved variables that predict treatment and are correlated with the outcome after controlling for  $X$ . The assumption that controlling for  $X$  removes all confounding variation is quite strong and hard to prove.

I use the Mahalanobis distance metric (MDM) to determine the “closest” counterfactual

---

<sup>15</sup>Note that the identification assumption is slightly less restrictive than the one needed to identify the ATE (Abadie and Imbens, 2006).

match for each rent controlled tract. The Mahalanobis distance for any two tracts  $i, j$  is calculated as:

$$MD = \sqrt{(X_i - X_j)^T S^{-1} (X_i - X_j)}$$

where  $X$  is a vector of covariates and  $S$  is the covariance matrix of  $X$ . The intuition behind the Mahalanobis metric is that it calculates the distance between two points in a way that is independent of the scale of each component of  $X$ . Recent work by King and Nielsen (2019) suggests that matching on the MDM is preferable to matching on estimated propensity scores, since the matched pairs come closer to mimicking a fully blocked experiment compared to propensity score matches which mimic a fully randomized experiment.<sup>16</sup> Fully blocked experiments are more efficient and should have substantially less noise in the estimated treatment effects. I match tracts with replacement, which allows a control tract to serve as the counterfactual for multiple observations.

Throughout the paper, I use 38 main covariates to calculate the MDM as well as test for balance in the subsequent matched or weighted samples. In Chetty *et al.* (2018), the authors report tract level variables that correlate with long-term outcomes. In particular, they show that education, poverty rates, single parenthood, income, unemployment, minority population and proxies for social capital are all correlated with economic mobility at the tract level. Therefore, it is crucially important to include these variables when calculating the distance between tracts and to ensure that they are balanced in the post-match sample. In addition to the variables highlighted by Chetty *et al.*, I also include various housing and population variables that are likely to predict the imposition of rent control at the municipal level. These variables include per capita municipal revenue, expenditures and county level data on voting behavior from the 1968 presidential election. The municipal and county level data allows me to control for city-level differences that are not accounted for at the tract level and that could be correlated with the outcomes of interest. The full set of covariates are the same tract and city-level variables that are listed in Tables 1 and 2.

As Ho *et al.* (2007) suggests, the main goal of a matching strategy is to reduce covariate imbalance across treatment status (with the hope that the covariates remove all confounding variation). Despite the large body of research on propensity score and distance metric matching, there is limited consensus on the best way to estimate matching metrics (Hainmueller, 2012). This is particularly true in the context of estimating propensity scores, though still

---

<sup>16</sup>In a fully randomized experiment, treatment is assigned at random across a given population. In a fully blocked experiment, the sample is stratified based on pre-treatment characteristics, and treatment is assigned randomly within strata.

relevant for the MDM when deciding which variables to include and whether to use higher order terms when assessing the distance between two observations. Given the large number of tract and city-level covariates I control for, I do not include any higher order terms when calculating the Mahalanobis distance.<sup>17</sup>

Another common strategy in the matching literature is to limit matches to observations that are within a given distance caliper or radius. Conducting a nearest neighbor match without a caliper will find the closest match for each treated observation. As one decreases the size of the matching caliper, the matches that are farthest in measured Mahalanobis distance are pruned, leaving the better matches with closer covariate realizations. This process of determining the match caliper is another way of describing the bias - variance trade off common to many empirical approaches. In addition, as treated observations are pruned, the estimated treatment effect from the reduced sample may not be relevant for the target sample. As a result, it is up to the researcher to determine the caliper that minimizes covariate imbalance while also maintaining a sufficient sample of observations to estimate treatment effects.

I use an iterative process to determine the optimal caliper by instituting different caliper values and checking the number of treated observations that are pruned and the resulting covariate imbalance. In all iterations of the model, I check covariate balance by comparing the standardized mean difference between the treated and control group before and after implementing the matching procedure. The standardized mean difference is given by  $\frac{\bar{Y}_1 - \bar{Y}_0}{\sqrt{\frac{V_1 + V_0}{2}}}$ , and is thus measured in units of the pooled standard deviation of the treatment and control means (Austin, 2009). Its use allows us to standardize the measure of divergence regardless of the units of each covariate. In addition, it satisfies the conditions of a good balance check statistic suggested by Imai *et al.* (2008) in that it is a characteristic of the sample (and not a hypothetical population) and that the value is unaffected by sample size. It is also important to note that the iterative caliper selection process is estimated without estimating treatment effects to avoid data mining particular results while searching for the optimal match radius.

Once I have finalized the caliper selection, I estimate the average treatment effect on the treated for all outcome variables by averaging the difference between the treated and matched counterfactual across all treated tracts. I also account for the bias that results when matching on continuous covariates (Abadie and Imbens, 2006) by including post matching

---

<sup>17</sup>There is some older research showing that the Mahalanobis distance metric performs poorly with many covariates (Gu and Rosenbaum, 1993, for example); however, in my context, the Mahalanobis distance provides the best matches resulting in the lowest residual imbalance on observables despite the large number of covariates.

regression adjustment on all covariates. This is equivalent to the bias-adjusted matching estimators proposed in Abadie and Imbens (2011). I report clustered standard errors at the city level using the influence function proposed by Jann (2019). Note that this standard error is likely conservative, compared to the consistent standard errors proposed by Abadie and Imbens (2006) which are harder to calculate while including clusters. In addition Abadie and Imbens (2008) show that bootstrapped standard errors are not consistent for matching estimators despite their use in the literature.

The ideal post-matching sample will have a standardized mean difference of 0 though researchers sometimes use simple thresholds such as 0.1 or 0.25 to signify when a covariate has a large enough difference to warrant returning to the matching procedure (Stuart *et al.*, 2013; Rubin, 2001; Cohen, 1977; Normand *et al.*, 2001; Austin, 2009). Ultimately it is up to the practitioner to identify whether the imbalanced covariate could plausibly lead to biased estimation of treatment effects. Slight imbalances of covariates that are highly correlated with the outcome will be more problematic than larger differences in a variable which has a weak correlation with the outcome of interest.

I also measure the ratio of the mean variance for the treated and control groups for each covariate. Austin (2009) shows that propensity scores that are balanced on means still may be incorrectly specified. Covariates that are fully balanced should have a variance ratio of 1, though the literature gives little direction when it comes to determining what variance ratio threshold implies an imbalanced covariate. I view variance imbalances as subordinate to imbalances of the mean when determining the match caliper.

Table 4 show the standardized mean differences (SMD) and variance ratios for the raw and matched data after selecting the optimal caliper. Given the large number of covariates, the resulting imbalances are expected though larger than ideal. The caliper drops slightly more than half of the original treated tracts, leaving a sample of 1,174 tracts which will be used to estimate the relevant ATTs. Out of 38 covariates, 18 have SMD of less than the desired 0.1, while 13 have SMD values of between 0.1 and 0.25. This leaves 7 covariates that have post matching SMD values of greater than 0.25. The post-matching variance ratios show that the majority of covariates have variance ratios within the [0.7, 1.3] interval; however, there are still 13 covariates with variance ratios outside of this interval, further suggesting some residual imbalance. Given the resulting imbalance, the post matching regression adjustment is especially important when estimating treatment effects.

Across the distribution of included treated tracts, there are a total of 549 unique control tracts that are selected as matches. This means that each counterfactual observation is selected as the closest neighbor match for an average of 2.1 treated observations. The

maximum number of times a control observation is matched is 24 while the median is 1. The caliper selection process dropped 1,290 treated tracts including most of the tracts from rent controlled cities in Massachusetts. Of the remaining treated tracts, approximately 35.4% are located in New Jersey (416 tracts), 63.1% are located in California (741 tracts) and 1.4% are located in Massachusetts (17 tracts). These tracts represent 88 out of the potential 99 cities that are included in the final analysis sample.

The observation pruning suggests that the estimated average treatment effect on the treated is not equivalent to the overall sample average treatment effect on the treated. I view the change in underlying sample to be a worthy tradeoff to achieve better covariate balance. It is also worth noting that matching models with more permissive calipers yield very similar results to this more restrictive matching procedure.

Lastly, I also estimate a matching model that limits treated and control tracts to those that have a high proportion of rental units. This removes over one half of the total tracts in the original sample, including 25% of the treated tracts. I use the same baseline covariates and matching caliper to construct these matches, which are used to determine the effect of rent control on a sample of tracts that are especially likely to be affected by the policy. The balance results from this matching procedure are reported in Table A.1. Of the 840 tracts included in the high rent sample, 72% are located in California and 28% are located in New Jersey.

## 5.1 Additional Methodological Assumptions

The publicly available outcome data has shortcomings that require me to make a strong assumption before interpreting my results. The sample of children included in each tract's outcome is determined by links made starting in 1989, when the analysis cohort is between 6 and 11 years old. Figure 1 shows a timeline of the relevant dates when states began implementing rent control, as well as the dates when children were eligible to be linked to a given census tract. Given the timing, I must assume that rent control did not cause any selective immigration to or from census tracts in the years after it is implemented and before 1989. Under this assumption, the people living in rent controlled tracts in 1989 should be roughly equivalent to those living in the matched counterfactual tracts in 1989. This assumption is stronger for the rent controlled cities in New Jersey and Massachusetts than it is in California due to the longer duration between rent control implementation and address linkage.

Prior work suggests that violations of this assumption could bias the estimated treatment

effects in either direction. Research by Autor *et al.* (2014, 2017) shows that the removal of rent control in Cambridge, Massachusetts led to significant property appreciation and decreased crime, suggesting that rent control is associated with reverse gentrification. If this is true, rent controlled areas may be less likely to receive an influx of new residents than similar areas without rent control. Under this scenario, the estimated ATT of rent control on long-term outcomes would likely be biased downward. Alternatively, research by Diamond *et al.* (2019) suggests differences in the short and long run effects of rent control on neighborhood immigration. While rent control allows some lower income individuals the ability to stay in their homes for longer, it also incentivizes landlords to exit the rental market. In the long run, they suggest that rent control increases gentrification by substituting rental housing for more expensive private uses. Under this latter scenario, we might expect the children living in rent controlled neighborhoods in 1989 and later to be relatively wealthier than the children comprising the counterfactual tracts. This would likely create an upward bias on the estimated effect of rent control on long-term outcomes, since the outcome sample would include wealthier people that select into areas with rent control.

In Figure 2, I test this immigration assumption by looking at the effect that rent control has on the tract level proportions of immigrants from other counties, states and countries. These variables are based on the census question asking respondents over 5 years old where they lived 5 years ago and are aggregated up to the census tract level. For each year of the decennial census, I estimate the average treatment effect on the treated on these immigration outcomes using the baseline Mahalanobis distance matching estimator. The top left panel shows the effect of rent control on the tract level percentage of people who migrated from a different census tract. The estimated ATT is negative in 1980 and positive in 1990, suggesting that rent control does not have a consistent effect on cross county migration. The top right panel shows that rent controlled tracts have fewer residents migrate from other states than the counterfactual tracts. The bottom panel shows that rent controlled tracts have a slightly higher proportion of international immigrants than non-controlled tracts in 1980 and 1990, though this effect reverses signs in 2000. These findings are very similar when using the high-rent sample and are shown in Figure A.1. In general, the balance of the evidence from this figure suggests that the Autor *et al.* prediction is better supported by the evidence, that rent control should lead to fewer new immigrants. Given this fact, I expect any violation of the selective immigration assumption to bias downwards the estimated benefits of rent control.

## 6 Results

I first show evidence that rent control affects the location decisions of the households included in the Opportunity Insights sample. The central finding is that rent control allows families to stay in their housing for longer. Specifically, the first column of Table 5 shows that rent control leads to a 3.6% increase in the tract-level average duration of tenancy compared to counterfactual tracts. Each child included in the Opportunity Insights data can be linked to a particular tract in up to 12 distinct tax years, but these years are not continuous. As a result, a 3.6% increase in the tract-level average of the fraction of years that a child is linked understates the increased tenant durations resulting from rent control. As an illustration of this fact, a child that lives in a given tract from 1989 to 1995 would only be linked to that tract in 3 separate years (1989, 1994, 1995) despite living there for a total of 6 years. I interpret the result as evidence that rent control increases rental durations without focusing on the exact magnitude of the increase.

The second column reports that rent control leads to a 14.6% increase in the tract-level probability that a child will live with their parents as an adult. The third column shows that rent control leads to a 7.1% increase in the average probability that a child will live in their childhood tract as an adult. The final column shows that rent control has minimal effect on the tract-level probability that a child will live in the same commuting zone as an adult. The main takeaway is that rent control allows families to stay in their housing for longer which is exactly what we would expect if rent control laws were binding and provided benefits to families in the form of lower than market rents.

The baseline long-term outcome results are reported in Table 6. The columns represent different outcome variables while the top row represents the estimated ATT of rent control on the relevant tract-level average outcome. This row is measured in percentage point differences. Standard errors are reported in parenthesis beneath the ATT estimates.<sup>18</sup> The first two columns report the tract-level probability that a child will reach the top family or individual income quintile in 2015. The third and fourth columns report the tract-level average income percentile in the family or individual income distribution in 2015. The fifth column reports the tract-level probability that someone linked to a given track will have a teenage birth (females only). The sixth and seventh columns report the tract-level probability that someone linked to a given tract will be in jail or prison in April, 2010 or will

<sup>18</sup>As a reminder, the standard errors in Table 6 are calculated using the influence function method proposed by Jann (2019) and account for clustering at the city level. These standard errors are likely to be overly conservative compared to the consistent matching standard errors proposed by Abadie and Imbens (2006) which are more difficult to calculate including clusters.



have positive earnings in 2015. Lastly, the eighth column reports the tract-level probability that someone linked to a given tract will live in a low poverty neighborhood as an adult, defined as a tract with a poverty rate of less than 10% according to 2000 Census data. The baseline row reports the average counterfactual value for the relevant outcome while the % $\Delta$  row calculates the percent change the ATT row represents from the baseline value.

Columns 1 and 2 show that rent control has a -5.4 and -3.5% effect on the tract-level probability that a child will reach the top family and individual income quintile respectively. Column 3 shows that rent control has a small (-1.3%) effect on the tract-level average family income percentile and a negligible effect on the average individual income percentile. Columns 5 through 7 show that rent control has a minimal effect on the tract-level probability that a child will have a teenage pregnancy, be incarcerated during the 2010 census or report positive employment earnings in 2015. Lastly, column 8 shows that rent control is associated with a 5.2% reduction in the tract-level probability that children will live in low poverty neighborhoods as an adult.

The unit of observation for the results in Table 6 is the tract level. Although tracts are constructed to have roughly the same population, it is possible that tracts with more children may have different estimated treatment effects. Table 7 reports the same matching ATT results but weights each tract by the number of children that are used to calculate the outcome variable.<sup>19</sup> While this does not allow for an individual-level interpretation of the results, it does scale the estimated ATT to better reflect differences in the size of tracts and, more importantly, the number of children used to measure the outcome variable. Columns 1 and 2 suggests that rent control leads to a -4.2 and -2.6% decrease in the probability that a child reaches the top family and individual income quintile respectively. Columns 4 and 5 show that the weighted treatment effect of rent control on the average income percentile is negligible. Column 5 shows that rent control leads to a 1.7% decrease in the average probability that a female child will experience a teenage pregnancy. Column 6 shows that rent control is associated with a 9.1% decrease in the probability that a child will be incarcerated during the 2010 decennial census. Column 7 shows that rent control has a negligible effect on the average probability that a child has positive earnings in 2015. Lastly, column 8 shows that rent control is associated with a 2.4% decrease in the average probability that a child lives in a low poverty neighborhood as an adult. On balance, these results are fairly similar to the unweighted baseline results.

In the last part of my baseline analysis, I limit the sample of tracts to include only those tracts that have a high proportion of renters, defined as any tract with greater than a 30%

---

<sup>19</sup>The Opportunity Insights data reports the underlying sample size for each outcome variable.

share of rentals of the total housing units. This sample modification drops slightly more than half of the non-rent controlled tracts and one quarter of all rent controlled tracts. The theory underpinning this adjustment is that tracts with a high percentage of rentals are more likely to exhibit effects directly related to rent control. Indeed, Table 8 shows that rent control leads to an estimated 12% increase in the fraction of years spent in a given tract and a 13.7% increase in the probability of living in the same tract as an adult. These results provide further proof that the rent control laws I study allow families to stay in their homes for longer.

In Table 9, I show the effect of rent control on long-term outcomes for children growing up in tracts with a high proportion of rentals. Column 1 shows that rent control leads to a 5.9 and 3.9 percentage increase in the average tract-level probability of reaching the top 20% of the family and individual income distribution. Rent control is also associated with a 2.7 and 3.3% increase in the tract-level average family and individual income percentile. Based on the crosswalk converting income percentiles to income in 2015 dollars, the increase in average individual income attributed to rent control is between \$850 and \$1,709. The results in columns 5 and 6 shows that rent control has minimal impact on the average teen-pregnancy rate, and leads to a 16.8% increase in the tract-level probability of being in jail on April 1, 2010. Column 7 shows that rent control leads to a 2.7% increase in the tract-level employment average. Lastly, rent control leads to a 4.2% decrease in the probability of living in a low poverty neighborhood as an adult. One thing to note is that the sample of tracts used for the high-rental estimates is fairly small. By definition, these tracts are not representative of the full sample of rent controlled tracts. Despite this, the change in estimated effects between the full sample results and the high rental results shows that rent control leads to greater economic and employment outcomes in areas where there is a higher probability of residents living in rent controlled housing.

These results strongly suggest that rent control leads to long-term benefits for children that grow up in rent controlled housing. The difference in estimated effects between the baseline and high rent sample also suggest that rent control has spillover effects that may harm the long-term outcomes of children that live in close proximity to rent controlled housing. In the next subsection, I try to determine how these long-term benefits are distributed across the income distribution.

## 6.1 Rent Control at Different Points of the Income Distribution

In Table 10, I show the estimated average treatment effect on the treated tracts of rent control on the predicted outcomes at the 25th and 75th percentile of the parent income distribution. For these results, the outcome variable is a predicted value of a regression performed for each census tract in the data, regressing the relevant outcome on parent income rank. The 25th (75th) percentile outcome is thus the fitted value from this regression at the relevant parent income level. This means that the predicted outcomes are a projection based on the outcomes of all children that are linked to a given tract.

Panel A of the table shows the effect of rent control on predicted tract-level economic mobility for children at the 25th percentile of parent income. In general, the estimated tract-level effects of rent control on economic mobility for children at the 25th income percentile are similar to those reported in Table 6. Column 1 shows that rent control is associated with a 4.9% decrease in the predicted tract-level probability of reaching the top family income quintile, but this estimate has a standard error over 2 times larger than the point estimate. Columns 3 and 4 show that rent control has little effect on the tract-level predicted family and individual income percentile. Column 5 shows that rent control is associated with a 4.5% decrease in the tract-level probability of having a teen pregnancy. Column 7 shows that rent control is associated with a 1.9% increase in the predicted tract-level probability of being employed as an adult. Lastly, column 8 shows that rent control decreases the average probability of living in a low poverty neighborhood by 3.2%.

Panel B of the table shows the effect of rent control on predicted economic mobility for children at the 75th percentile of parent income. The reported effects of rent control on economic mobility and average income are quite small. Rent control is associated with a slight increase in the average probability of having a teen pregnancy, and a large (though statistically insignificant) decrease in the average probability of being incarcerated.

In Table 11, I show the results of this same exercise but limiting the sample to the high rental tracts. In panel A, I show that rent control has a small and slightly positive effect on the average predicted economic mobility of children at the 25th percentile of parent income. This contrasts with the effect shown in panel B for children at the 75th percentile of the parent income distribution, where rent control is associated with a 13.2 and 8.2% increase in the predicted probability of reaching the top family and individual income quintile. The estimates in columns 1-4 are all significant at the 5% significance level, further suggesting that rent control leads to better economic outcomes for children at higher income levels.

Comparing the effect of rent control on average predicted teen pregnancy between panels

A and B, the estimates indicate that rent control leads to a 5% decrease in average teen pregnancy for lower income children, and a 15.5% increase in teen pregnancy for children towards the top of the parent income distribution. The result in panel B is surprising, especially considering that rent control is associated with large improvements in economic mobility, and we might expect the estimates to move in the same direction.

The results reported in Table 11 suggest that rent control is not particularly well-targeted to individuals at the bottom of the income distribution; however, since the outcomes are projections, it is entirely possible that the results are driven by people towards the top of the parent income distribution. As such, these results should be interpreted with care and certainly do not provide conclusive proof of the effect that rent control has on children at the bottom or top of the income distribution.

## 6.2 Rent Control Effects on Housing and Demographics

In this section, I use census data from the 1980-2000 surveys to see how rent control affects the demographic composition of tracts, as well as the housing market. For each census year between 1980 and 2000, I estimate the average treatment effect on the treated tracts of rent control on a host of demographic and housing variables. Part of the goal of this section is to examine possible mechanisms to explain the baseline ATT results. Prior research has shown that rent control can suppress gentrification in a neighborhood or city by making it easier for tenants to remain in their current housing. This might explain why rent control has a slight negative effect on economic mobility when estimating the baseline ATT model. People who do not receive rent control but live in a city with rent control may receive a negative spillover from the policy. I use the same nearest neighbor matching methodology with bias adjustments to generate these estimates. As a result, the analysis sample is the same, meaning that the results are based on 1,174 rent controlled tracts spread over California, Massachusetts and New Jersey. I also estimate these effects using the high rental sample, though the results are largely the same. These supplemental figures are reported in appendix B.

Figure 3 displays estimated average treatment effects on the treated for housing market conditions. The first panel on the left column shows the average tract-level effect of rent control on the probability that a resident has been living in their current house for at least 5 years. The estimated effect of rent control on the probability of living in one's house for more than 5 years is statistically significant in all three census years reported. This indicates that rent control increases tenancy duration in tracts with rent control and provides corroborating

evidence that rent control allows renters to stay in housing for longer. The top panel on the right shows that rent control leads to a small but negative effect on the average tract-level vacancy rate, though these estimates are not statistically significant.

The middle panels show the effect that rent control has on the tract-level proportion of people that have long commutes, and the tract-level percentage of rentals as a share of total housing units. The panel on the left provides weak evidence that rent control leads to decreased commutes of over 1 hour. This contrasts with Krol and Svorny (2005) who find that rent control in New Jersey lead to spatial mismatch between where people live and work further resulting in longer average commutes. The panel on the right shows that compared to counterfactual tracts, rent control has a negative effect on the proportion of rentals in 1980, and zero effect on the proportion of total units that are rentals in 1990 and 2000. The panel on the bottom shows that rent control has a statistically significant (though small) effect on the average total number of units in a tract. Taken together, this partially corroborates evidence reported in Diamond *et al.* (2019) and Sims (2007) that rent control leads landlords to remove rental units from the rental market.

Figure 4 displays estimated average treatment effects on the treated tracts for demographic variables that could play a role in changing the long-term outcomes of residents. Chetty *et al.* (2018) show that long-term outcomes are correlated with tract level characteristics such as the poverty rate, education, unemployment and rate of single parent families. In the top left panel, I show that rent control is associated with a small increase in the proportion of single parent families. In the top right panel, I show that rent control is associated with a consistent decrease in the tract-level share of college graduates. In the bottom panels, I show that rent control is associated with a slight increase in the tract-level poverty rate and increases in the unemployment rate in the 1990 and 2000 decennial census. Although the measured effects are small, the results suggest that rent control leads to a reversal or hold on neighborhood gentrification. This is mostly consistent with results reported by Autor *et al.* (2014, 2017). In addition, this suggests that the benefits of rent control might be offset by negative spillover effects on long-term outcomes for those that do not live in rent controlled housing.

## 7 Discussion and Limitations

The results from the high rental sample suggest that rent control improves economic mobility and employment outcomes in areas where we should see rent controlled tenants exert a greater effect on the tract-level average outcome. I interpret these latter results to

indicate that rent control provides a long-term benefit to children that grow up in rent-controlled housing, though the exact magnitude of this benefit is difficult to quantify with the data in this study. While few of the estimated ATT effects in the high rent sample are statistically significant, there are reasons to think that I am understating the true benefit of rent control on the long-term outcomes of children. In addition, conservative standard errors are likely to overstate the variance of my estimates, making it more difficult to reject the null hypothesis of 0 effect.

The methodology I use is only capable of estimating the net effect of rent control in a census tract and cannot disentangle the effects on rent-controlled tenants and everyone else. As the figures in section 6.2 show, rent control is associated with a slight demographic shift that leads to lower college attendance, higher unemployment rates and higher poverty rates. Chetty *et al.* (2018) reports that the long-term outcomes of children are negatively correlated with these demographic shifts, implying that rent control could harm children growing up in neighborhoods with rent control but who do not live in rent controlled housing. If this is the case, the estimated benefit of receiving rent control would include both the effect on those who receive rent control and the spillover effects on those children who do not but are affected by the changing neighborhood demographics. This would indicate that the effects I measure understate the true benefit of rent control to children who grow up in rent controlled housing.

Even if we suppose that these spillover effects are negligible, it would still require a large magnitude effect to detect tract-level changes in outcomes as a result of rent control. Suppose there are 100 children in a tract, and the true effect of rent control is that it improves the probability of reaching the top income quintile by 10% (improving the baseline probability from 20 to 22%). If we also assume that half of the children live in rental housing (which is roughly equal to the share of rental housing in rent controlled tracts), and half of these renter children live in rent controlled housing, then 25 children will have improved chances of reaching the top income quintile; however, the tract-level economic mobility probability will be  $((0.75) \times (0.2)) + ((0.25) \times (0.22)) = 0.205$ , which is roughly equal to the baseline probability of 0.2, even though rent control has a large effect on those that receive the benefit.

From the high rental sample, the estimated effect of rent control on the average tract-level individual income is 1.6 percentiles. The difference in annual income between percentile 49.2 and percentile 50.8 is roughly equivalent to \$1,111, measured in 2015 dollars. This represents a 4% increase (from a baseline of \$28,500) that is attributable to rent control. This effect is smaller than the one reported by Andersson *et al.* (2016), who find that an additional year of public housing is associated with a 5% increase in annual earnings. The difference

in the magnitudes of the estimates can be reconciled by the fact that the outcome data is aggregated and includes children that do not receive rent control.

Another related issue that can attenuate the estimated benefits of rent control is the fact that I do not have tract-level measures of rent control prevalence. By treating rent control as a binary variable, I do not account for the fact that different tracts are likely to have different levels of rent control intensity. Given the complicated rules that govern exemptions to rent control laws, it would be very challenging to construct a measure of the true rate of rent control in each tract over all cities in California, Massachusetts and New Jersey that enacted rent control between 1970 and 1985. Despite this, the methodological decision to measure rent control as binary will likely lead to an attenuation bias of the target average treatment effect on the treated. This bias is likely mitigated in the high rental sample analysis, though there is still the probability that there is differential rent control treatment within these tracts leading to measurement error and attenuation bias.

Lastly, the causal interpretation of the matching results rests on the assumption that conditioning the receipt of rent control on the matching covariates removes all confounding variation. I attempt to use a wide variety of tract and city-level covariates when matching to account for as many observable characteristics as possible. The large number of covariates makes it more difficult to balance the full covariate vector after matching, which further leads to possible bias of the estimated ATT on the outcomes of interests. I mitigate this concern by using the bias corrected estimators proposed by Abadie and Imbens (2011); however, this is a suboptimal method for dealing with residual imbalance.

Future research can avoid these problems using a number of different strategies. Most importantly, access to better data that is reported at the individual child level would allow for one to disentangle the direct and spillover long-term effects of rent control. Additionally, studies that include more detail on the level of rent control intensity would have a better chance of avoiding the attenuation bias that is likely to plague my results. Lastly, researchers can utilize random or quasi-random variation in the assignment of rent control to more confidently avoid potential confounding of the effect of interest.

## 8 Conclusion

This paper is the first to provide estimates of the long-term benefits to children of growing up with rent controlled housing. Rent control is a commonly implemented policy that currently exists in 5 states including some of the largest cities in the country. Despite its

ubiquity, previous research on rent control has mainly focused on its impact to housing markets and neighborhoods while ignoring the potential benefits to renters. There has been some attempt to quantify the benefits of rent control by calculating the difference between controlled rent and the rent that would be charged in a competitive market, but these static measures ignore the potential long-term benefits that the policy confers.

I show that rent control leads to decreases in economic mobility and minimal changes in tract-level teen pregnancy rates, incarceration and employment; however, when limiting the sample of tracts to those that have at least a 30% rental share, I show that the estimated effect of rent control on economic mobility, income and employment is positive. These findings are consistent with prior research showing that rent control leads to negative spillovers on surrounding properties, while also providing clear benefits to those who live in controlled housing. To further contextualize these results, I use census data from the 1980 to 2000 census to show that rent control leads to tract-level changes in demographics and housing variables that are consistent with rent control suppressing gentrification.

In addition, I show some evidence that rent control can lead to minor improvements in economic mobility for people at the bottom of the income distribution; however, the economic mobility effects of rent control appear significantly stronger for children growing up at the top of the parent income distribution.

These results represent an important first step in measuring the long-term benefits of rent control for children that live in rent controlled housing, as well as the potential negative spillovers that worsen long-run outcomes for those that live in non-rent controlled housing. This research suggests that the long-term direct benefits are positive, though future research should explore this topic using alternative data to determine whether the estimated effects can be reproduced using individual-level variation instead of aggregated data. Additional inquiry can also help to determine how the direct benefits of rent control compare to the potential negative long-term spillovers which can help policymakers measure the net effect of rent control on residents.



## References

- Abadie, A. and Imbens, G. W. (2006) Large sample properties of matching estimators for average treatment effects, *Econometrica*, **74**, 235–267.
- Abadie, A. and Imbens, G. W. (2008) On the failure of the bootstrap for matching estimators, *Econometrica*, **76**, 1537–1557.
- Abadie, A. and Imbens, G. W. (2011) Bias-corrected matching estimators for average treatment effects, *Journal of Business & Economic Statistics*, **29**, 1–11.
- Almond, D., Currie, J. and Duque, V. (2018) Childhood circumstances and adult outcomes: Act ii, *Journal of Economic Literature*, **56**, 1360–1446.
- Andersson, F., Haltiwanger, J. C., Kutzbach, M. J., Palloni, G. E., Pollakowski, H. O. and Weinberg, D. H. (2016) Childhood housing and adult earnings: A between-siblings analysis of housing vouchers and public housing, Working Paper 22721, National Bureau of Economic Research.
- Asquith, B. J. (2018) Do rent increases reduce the housing supply under rent control? evidence from evictions in san francisco, *Working Paper*.
- Ault, R. and Saba, R. (1990) The economic effects of long-term rent control: The case of new york city, *The Journal of Real Estate Finance and Economics*, **3**, 25–41.
- Ault, R. W., Jackson, J. D. and Saba, R. P. (1994) The effect of long-term rent control on tenant mobility, *Journal of Urban Economics*, **35**, 140 – 158.
- Austin, P. C. (2009) Balance diagnostics for comparing the distribution of baseline covariates between treatment groups in propensity-score matched samples, *Statistics in Medicine*, **28**, 3083–3107.
- Autor, D. H., Palmer, C. J. and Pathak, P. A. (2014) Housing market spillovers: Evidence from the end of rent control in cambridge, massachusetts, *Journal of Political Economy*, **122**, 661–717.
- Autor, D. H., Palmer, C. J. and Pathak, P. A. (2017) Gentrification and the amenity value of crime reductions: Evidence from rent deregulation, Working Paper 23914, National Bureau of Economic Research.
- Carlson, K. (2015) Fear itself: The effects of distressing economic news on birth outcomes, *Journal of Health Economics*, **41**, 117 – 132.
- Chetty, R., Friedman, J., Hendren, N., Jones, M. and Porter, S. (2018) The opportunity atlas: Mapping the childhood roots of social mobility, Working Paper 25147, National Bureau of Economic Research.
- Chetty, R. and Friedman, J. N. (2019) A practical method to reduce privacy loss when disclosing statistics based on small samples, *Journal of Privacy and Confidentiality*, **9**.
- Chetty, R., Grusky, D., Hell, M., Hendren, N., Manduca, R. and Narang, J. (2017) The fading american dream: Trends in absolute income mobility since 1940, *Science*, **356**, 398–406.


- Chetty, R. and Hendren, N. (2018a) The Impacts of Neighborhoods on Intergenerational Mobility I: Childhood Exposure Effects\*, *The Quarterly Journal of Economics*, **133**, 1107–1162.
- Chetty, R. and Hendren, N. (2018b) The Impacts of Neighborhoods on Intergenerational Mobility II: County-Level Estimates\*, *The Quarterly Journal of Economics*, **133**, 1163–1228.
- Chyn, E. (2018) Moved to opportunity: The long-run effects of public housing demolition on children, *American Economic Review*, **108**, 3028–56.
- Clubb, J. M., Flanigan, W. H. and Zingale, N. H. (2006) Electoral data for counties in the united states: Presidential and congressional races, 1840-1972.
- Cohen, J. (1977) Statistical power analysis for the behavioral sciences.
- Cunha, F., Heckman, J. J., Lochner, L. and Masterov, D. V. (2006) Interpreting the evidence on life cycle skill formation, Elsevier, vol. 1 of *Handbook of the Economics of Education*, pp. 697 – 812.
- Diamond, R., McQuade, T. and Qian, F. (2019) The effects of rent control expansion on tenants, landlords, and inequality: Evidence from san francisco, *American Economic Review*, **109**, 3365–94.
- Early, D. W. (2000) Rent control, rental housing supply, and the distribution of tenant benefits, *Journal of Urban Economics*, **48**, 185 – 204.
- Fallis, G. and Smith, L. B. (1984) Uncontrolled prices in a controlled market: The case of rent controls, *The American Economic Review*, **74**, 193–200.
- Fogelson, R. M. (2013) *The Great Rent Wars : New York, 1917-1929*, Yale University Press.
- Glaeser, E. L. and Luttmer, E. F. P. (2003) The misallocation of housing under rent control, *The American Economic Review*, **93**, 1027–1046.
- Grampp, W. D. (1950) Some effects of rent control, *Southern Economic Journal*, **16**, 425–447.
- Gu, X. S. and Rosenbaum, P. R. (1993) Comparison of multivariate matching methods: Structures, distances, and algorithms, *Journal of Computational and Graphical Statistics*, **2**, 405–420.
- Gyourko, J. and Linneman, P. (1989) Equity and efficiency aspects of rent control: An empirical study of new york city, *Journal of Urban Economics*, **26**, 54 – 74.
- Gyourko, J. and Linneman, P. (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.
- Hainmueller, J. (2012) Entropy balancing for causal effects: A multivariate reweighting method to produce balanced samples in observational studies, *Political Analysis*, **20**, 25–46.
- Ho, D. E., Imai, K., King, G. and Stuart, E. A. (2007) Matching as nonparametric preprocessing for reducing model dependence in parametric causal inference, *Political Analysis*, **15**, 199–236.
- Holland, P. W. (1986) Statistics and causal inference, *Journal of the American Statistical Association*, **81**, 945–960.

- Hoynes, H., Schanzenbach, D. W. and Almond, D. (2016) Long-run impacts of childhood access to the safety net, *The American Economic Review*, **106**, 903–934.
- Imai, K., King, G. and Stuart, E. A. (2008) Misunderstandings between experimentalists and observationalists about causal inference, *Journal of the Royal Statistical Society. Series A (Statistics in Society)*, **171**, 481–502.
- Jann, B. (2019) Influence functions for linear regression (with an application to regression adjustment), Working Paper 32, University of Bern.
- King, G. and Nielsen, R. (2019) Why propensity scores should not be used for matching, *Political Analysis*, **27**, 435–454.
- Krol, R. and Svoyny, S. (2005) The effect of rent control on commute times, *Journal of Urban Economics*, **58**, 421 – 436.
- Lett, M. R. (1976) *Rent Control: Concepts, Realities, and Mechanisms*, Center for Urban Policy Research.
- Logan, J., Xu, Z. and Stults, B. (2014) Interpolating u.s. decennial census tract data from as early as 1970 to 2010: A longitudinal tract database, *The Professional Geographer*, **66**.
- McFarlane, A. (2003) Rent stabilization and the long-run supply of housing, *Regional Science and Urban Economics*, **33**, 305 – 333.
- Moon, C.-G. and Stotsky, J. G. (1993) The effect of rent control on housing quality change: A longitudinal analysis, *Journal of Political Economy*, **101**, 1114–1148.
- Munch, J. R. and Svarer, M. (2002) Rent control and tenancy duration, *Journal of Urban Economics*, **52**, 542 – 560.
- Nagy, J. (1995) Increased duration and sample attrition in new york city’s rent controlled sector, *Journal of Urban Economics*, **38**, 127 – 137.
- Normand, S.-L. T., Landrum, M. B., Guadagnoli, E., Ayanian, J. Z., Ryan, T. J., Cleary, P. D. and McNeil, B. J. (2001) Validating recommendations for coronary angiography following acute myocardial infarction in the elderly: A matched analysis using propensity scores, *Journal of Clinical Epidemiology*, **54**, 387 – 398.
- Olsen, E. O. (1972) An econometric analysis of rent control, *Journal of Political Economy*, **80**, 1081–1100.
- Pierson, K., Hand, M. L. and Thompson, F. (2015) The government finance database: A common resource for quantitative research in public financial analysis.
- Rosenbaum, P. R. and Rubin, D. B. (1983) The central role of the propensity score in observational studies for causal effects, *Biometrika*, **70**, 41–55.
- Rubin, D. B. (2001) Using propensity scores to help design observational studies: Application to the tobacco litigation, *Health Services and Outcomes Research Methodology*, **2**, 169–188.
- Sims, D. P. (2007) Out of control: What can we learn from the end of massachusetts rent control?, *Journal of Urban Economics*, **61**, 129 – 151.

- South, S. J., Haynie, D. L. and Bose, S. (2007) Student mobility and school dropout, *Social Science Research*, **36**, 68 – 94.
- Stuart, E. A., Lee, B. K. and Leacy, F. P. (2013) Prognostic score based balance measures can be a useful diagnostic for propensity score methods in comparative effectiveness research, *Journal of Clinical Epidemiology*, **66**, S84 – S90.
- Suen, W. (1989) Rationing and rent dissipation in the presence of heterogeneous individuals, *Journal of Political Economy*, **97**, 1384–1394.
- Wood, D., Halfon, N., Scarlata, D., Newacheck, P. and Nessim, S. (1993) Impact of Family Relocation on Children’s Growth, Development, School Function, and Behavior, *JAMA*, **270**, 1334–1338.

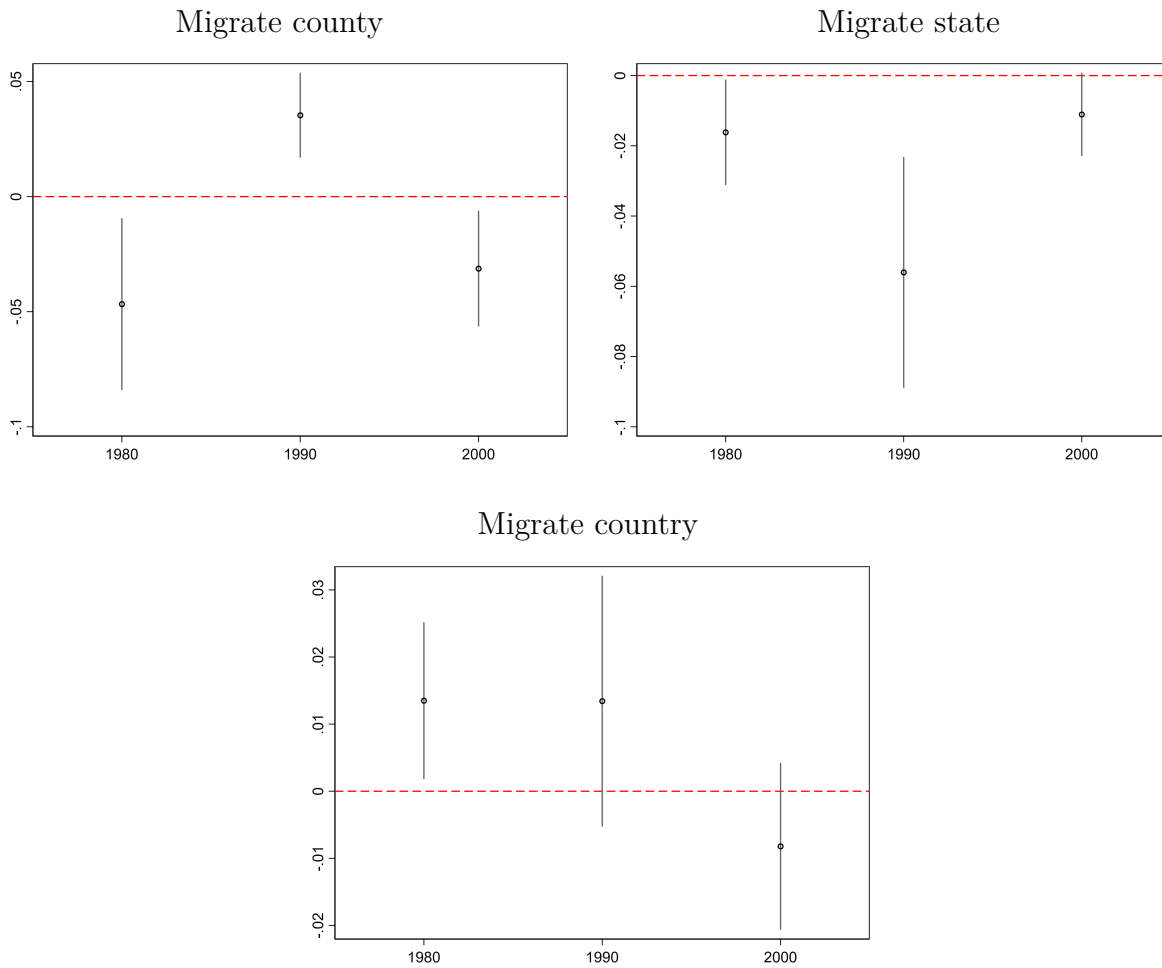
## 9 Figures

Figure 1: Timeline of relevant rent control laws and data construction dates



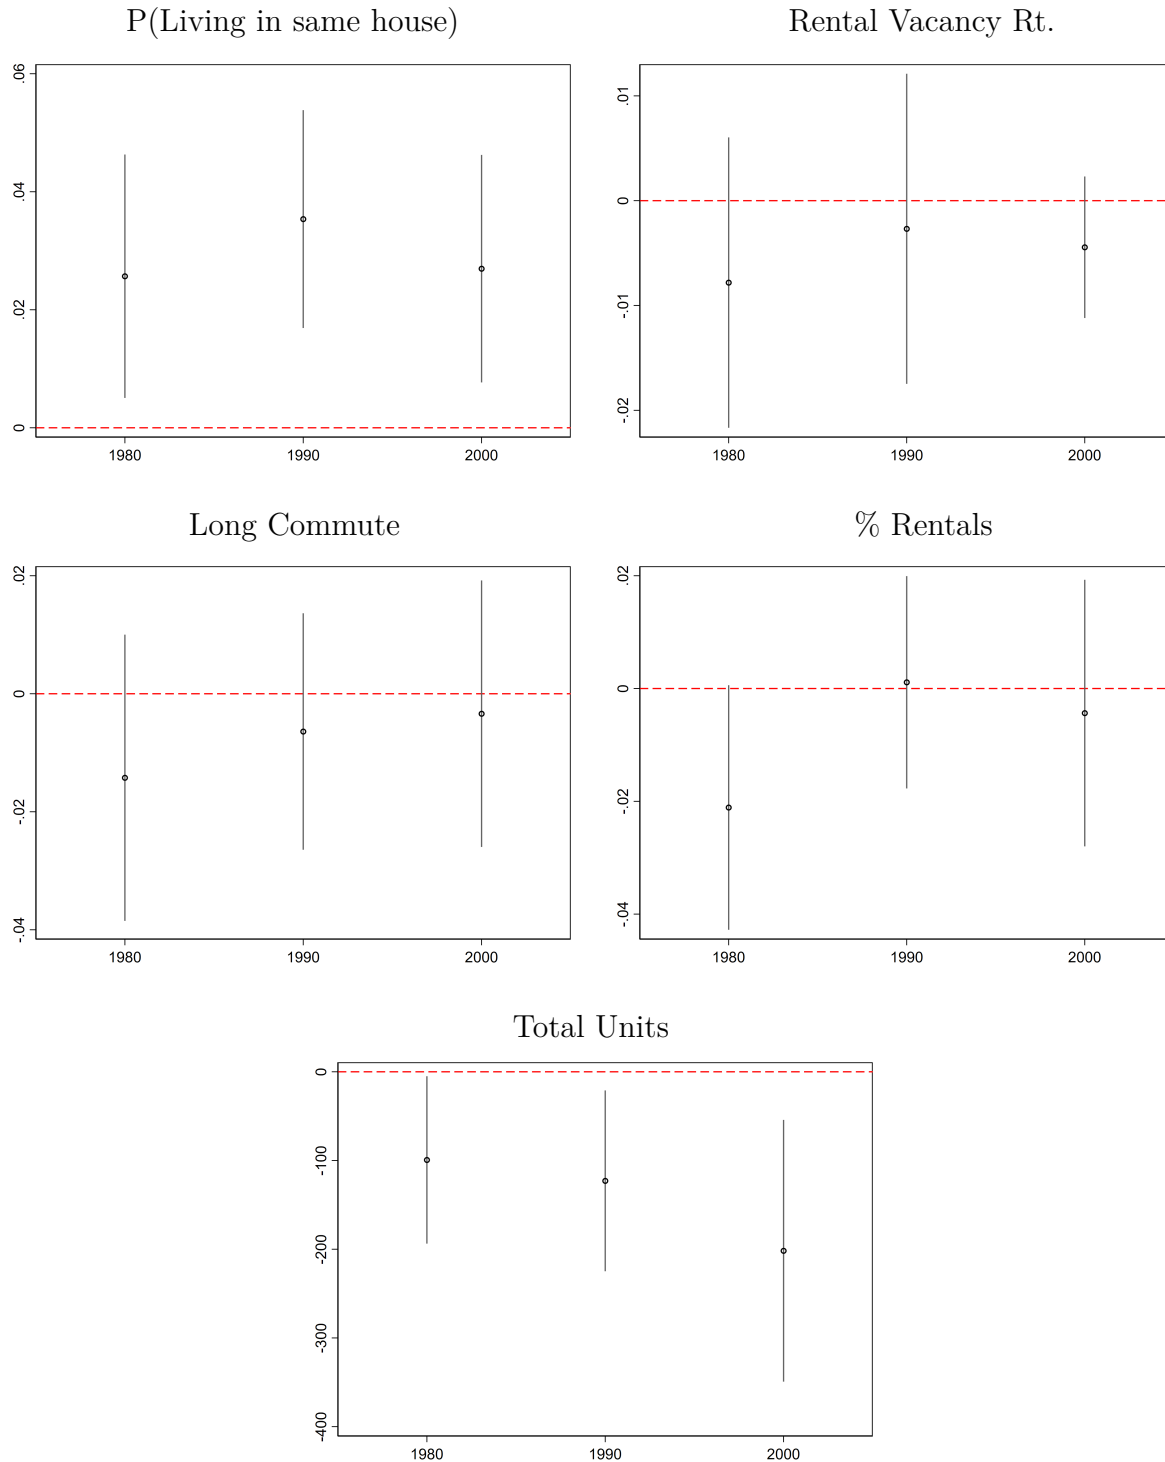
*Note:* Figure shows the years that each state began to implement rent control, as well as the years that are relevant to the construction of the Opportunity Insights data. There is a sizeable gap between the implementation of rent control and the first address link, particularly for New Jersey and Massachusetts. To recover a causal estimate of the impact of rent control, I must assume that there is no selective immigration to areas with rent control between the date of implementation and the address link.

Figure 2: Estimated ATT of rent control on immigration outcomes by year



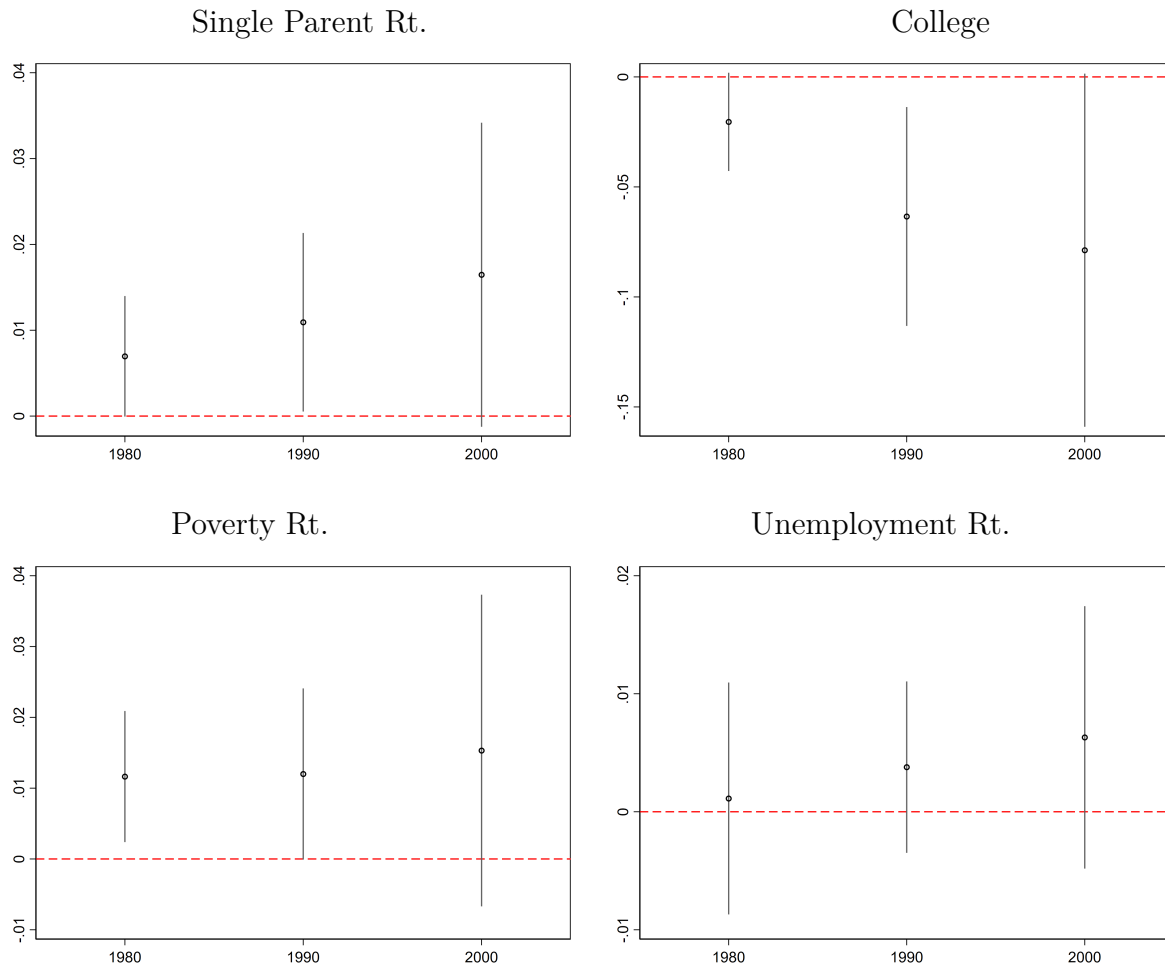
Notes: Figures show the average treatment effect on the treated tracts of rent control on immigration outcomes. Each outcome is the percentage of tract inhabitants that have moved from a given location in the last 5 years. The error bars represent 95% confidence intervals from standard errors that are clustered at the city level.

Figure 3: Estimated ATT of rent control on housing outcomes by year



Notes: Figures show the average treatment effect on the treated tracts of rent control on housing outcomes. The average outcomes are generated using the baseline nearest neighbor match model. The error bars represent 95% confidence intervals from standard errors that are clustered at the city level.

Figure 4: Estimated ATT of rent control on demographic outcomes



Notes: Figures show the average treatment effect on the treated tracts of rent control on employment and demographic outcomes. The average outcomes are generated using the baseline nearest neighbor match model. The error bars represent 95% confidence intervals from standard errors that are clustered at the city level.



## 10 Tables

Table 1: T-test of means to compare characteristics of treated and controlled census tracts

	Average		Difference	CA Diff.	MA Diff.	NJ Diff.
	Control	Treat				
Population	3,065.197	3,422.419	357.222***	348.446***	-327.377**	243.072**
Male (%)	0.485	0.478	-0.007***	-0.010***	-0.011***	-0.002*
Pop./sq. mile	4,607.690	13,303.566	8,695.876***	6,557.932***	14,391.206***	11,001.167***
Age median	28.322	30.856	2.534***	2.323***	-0.748	1.762***
white (%)	0.898	0.818	-0.081***	-0.141***	-0.132***	-0.051***
black (%)	0.092	0.144	0.052***	0.113***	0.121***	0.049***
Married (%)	0.634	0.561	-0.074***	-0.082***	-0.125***	-0.032***
Single parent fam. (%)	0.071	0.100	0.029***	0.026***	0.047***	0.016***
Educ. Less than HS (%)	0.437	0.434	-0.002	0.012**	-0.039***	0.073***
Educ. HS (%)	0.320	0.308	-0.012***	-0.025***	-0.012	-0.026***
LFP rate	0.598	0.603	0.005***	0.013***	-0.012**	0.013***
Unemployment rate	0.043	0.060	0.017***	0.007***	0.003*	0.004***
Avg. inc.	4,557.440	4,914.329	356.889***	272.167***	-9.041	-346.089***
Family poverty rt. (%)	0.089	0.093	0.003**	0.016***	0.041***	0.015***
Current addr. 5 years (%)	0.478	0.474	-0.004	0.015***	-0.096***	-0.026***
Housing units	1,026.231	1,237.823	211.592***	232.245***	17.667	154.410***
Rental (% of total units)	0.334	0.532	0.197***	0.142***	0.249***	0.233***
Rental vacancy rate (%)	0.052	0.035	-0.017***	-0.000	0.025***	0.000
Rent vacancy x Rental %	0.020	0.021	0.001*	0.006***	0.018***	0.004***
Avg. rent	127.908	143.727	15.819***	6.621***	18.171***	-3.224
Avg. home value	20,213.463	26,258.793	6,045.330***	3,713.865***	2,200.790**	-343.242
N	28,813	2,444				
N by State:						
California	4,052	1,563				
Massachusetts	538	184				
New Jersey	506	697				
Other States	23,717					

Sample includes tracts in all states except Washington DC, Maryland and New York. These excluded states had cities with rent control and cannot be used as possible control tracts. \* =  $p < 0.1$ , \*\* =  $p < 0.05$ , \*\*\* =  $p < 0.01$ .

Table 2: T-test of means to compare characteristics of treated and controlled cities

	Average		Difference	CA Diff.	MA Diff.	NJ Diff.
	Control	Treat				
City rental (% of total units)	0.298	0.450	0.152***	0.130**	0.311***	0.166***
City rental vacancy rate (%)	0.066	0.023	-0.043***	0.003	-0.004	-0.001
City white (%)	0.914	0.904	-0.010	-0.085*	-0.087	-0.020
City black (%)	0.054	0.069	0.015	0.077*	0.063	0.020
City unemployment rate	0.041	0.038	-0.003	0.001	-0.003	0.002
City avg. rent	118.028	140.819	22.790***	23.756**	50.587	11.518***
County Dem. vote share 1968	41.408	45.876	4.468***	5.802**	3.040	3.296***
County Wallace vote share 1968	14.154	8.005	-6.149***	-0.927**	0.264	-0.766**
City population	33,718.552	88,277.263	54,558.711*	435,493.354	204,199.588	26,450.147***
City family poverty rt. (%)	0.076	0.054	-0.022***	-0.003	0.018	0.006
City revenue per capita	0.173	0.228	0.055***	0.189**	0.214*	0.031*
City tax revenue per capita	0.075	0.146	0.072***	0.086**	0.151**	0.029**
Property tax share of rev.	0.293	0.511	0.218***	0.033	0.022	0.029
Other gov. sources share of rev.	0.142	0.155	0.013	-0.069***	-0.080**	0.008
Educ. share of expenditure	0.032	0.133	0.101***	0.000	-0.115*	0.040
Police share of expenditure	0.155	0.178	0.023***	-0.045***	0.015	0.003
Welfare share of expenditure	0.002	0.010	0.008***	0.027	-0.002	0.003***
N	2,373	99				
N by State:						
California	274	10				
Massachusetts	34	3				
New Jersey	149	86				
Other States	1,916					

Sample includes cities from all states except Washington DC, Maryland and New York. These excluded states had cities with rent control and cannot be used as possible control cities. \* =  $p < 0.1$ , \*\* =  $p < 0.05$ , \*\*\* =  $p < 0.01$ .

Table 3: T-test of means to compare outcomes of treated and controlled census tracts

	Average		Difference	CA Diff.	MA Diff.	NJ Diff.
	Control	Treat				
Fract. years in tract	0.594	0.597	0.004*	0.008***	-0.067***	-0.089***
Live with parents	0.164	0.248	0.083***	0.045***	0.029***	0.037***
Stay tract	0.189	0.235	0.045***	0.027***	-0.002	0.000
Stay comm. zone	0.698	0.741	0.043***	0.041***	-0.005	0.042***
Top 20% fam. inc.	0.194	0.224	0.030***	-0.004	-0.010	-0.064***
Top 20% ind. inc.	0.197	0.254	0.057***	0.011***	0.017**	-0.037***
Percentile fam. inc.	0.491	0.503	0.012***	-0.013***	-0.023***	-0.052***
Percentile ind. inc.	0.498	0.527	0.029***	0.005**	-0.003	-0.029***
Teen birth	0.210	0.164	-0.046***	-0.001	0.002	0.048***
Jail	0.018	0.012	-0.006***	0.000	0.004***	0.003***
Employed	0.766	0.754	-0.012***	-0.009***	-0.025***	-0.016***
Low pov. nbhd.	0.466	0.475	0.009**	-0.058***	-0.110***	-0.112***

Sample includes tracts in all states except Washington DC, Maryland and New York. These excluded states had cities with rent control and cannot be used as possible control tracts. \* =  $p < 0.1$ , \*\* =  $p < 0.05$ , \*\*\* =  $p < 0.01$ .

Table 4: Comparing means and variances of the raw and weighted samples using the two-step nearest neighbor matching algorithm

	Std. Mean Diff.		Var. Ratio	
	Raw	Matched	Raw	Matched
Population	0.185	-0.203	0.655	1.080
Male (%)	-0.226	0.026	1.101	1.259
Pop./sq. mile	0.910	-0.046	3.533	0.928
Age median	0.357	-0.068	1.139	1.220
white (%)	-0.322	-0.080	1.616	0.951
black (%)	0.211	0.041	1.548	0.931
Married (%)	-0.694	-0.053	1.210	1.433
Single parent fam. (%)	0.447	0.191	2.121	1.301
Educ. Less than HS (%)	-0.012	-0.214	0.962	1.075
Educ. HS (%)	-0.157	-0.009	0.695	1.096
LFP rate	0.076	-0.052	0.683	1.175
Unemployment rate	0.583	0.422	1.494	1.311
Avg. inc.	0.216	0.043	1.630	1.238
Family poverty rt. (%)	0.044	0.104	0.894	1.165
Current addr. 5 years (%)	-0.026	-0.362	0.938	1.175
Housing units	0.315	-0.164	0.739	1.107
Rental (% of total units)	0.843	0.116	1.616	1.183
Rental vacancy rate (%)	-0.258	0.066	0.470	1.407
Rent vacancy x Rental %	0.037	0.082	0.762	1.265
Avg. rent	0.367	0.153	1.334	1.124
Avg. home value	0.683	0.266	1.348	1.044
City rental (% of total units)	1.448	0.167	1.399	0.828
City rental vacancy rate (%)	-0.586	-0.167	0.240	0.739
City white (%)	-0.528	0.134	0.960	0.537
City black (%)	0.341	-0.254	0.822	0.416
City unemployment rate	0.912	0.638	0.976	1.659
City avg. rent	0.190	0.041	0.012	0.580
County Dem. vote share 1968	0.875	-0.010	0.536	0.521
County Wallace vote share 1968	-0.904	-0.120	0.021	0.870
City population	1.015	-0.021	4.548	0.681
City family poverty rt. (%)	-0.057	-0.017	0.395	0.880
City revenue per capita	0.882	0.257	2.282	1.454
City tax revenue per capita	0.935	0.048	2.677	1.052
Property tax share of rev.	0.394	-0.221	1.464	1.444
Other gov. sources share of rev.	0.135	-0.131	0.677	1.011
Educ. share of expenditure	0.375	-0.043	2.296	1.128
Police share of expenditure	-0.022	-0.382	0.419	0.687
Welfare share of expenditure	0.293	-0.058	5.377	0.879
N	31,443			
N Treat	1,174			
N unique control	549			

Table shows the standardized differences in means and variances between the raw and weighted sample. The unit of observation is a census tract. The treated tracts are matched to a control tract using a nearest neighbor Mahalanobis distance matching procedure. The Mahalanobis distance metric includes linear terms for tract-level, city-level and county-level characteristics.

Table 5: Mahalanobis distance nearest neighbor match estimates of average treatment effect on the treated of rent control on location as child and adult

	Frac. years in tract	Stay home	Stay Tract	Stay Comm. Zone
ATT	0.021** (0.010)	0.031*** (0.011)	0.016 (0.010)	0.004 (0.025)
Baseline	0.576	0.216	0.219	0.738
% $\Delta$	0.036	0.146	0.071	0.005
N Treat	1,156	1,155	1,156	1,156
N Control	541	539	540	540

The ATT row reports the average treatment effect on the treated of rent control on the average tract outcome. The ATT estimates are generated using a nearest neighbor Mahalanobis distance metric matching estimator. The baseline row represents the average outcome of the matched counterfactual tracts.

\* =  $p < 0.1$ , \*\* =  $p < 0.05$ , \*\*\* =  $p < 0.01$ .

Table 6: Mahalanobis distance nearest neighbor match estimates of average treatment effect on the treated of rent control on long-term outcomes

	Top 20% Fam.	Top 20% Ind.	Percentile Fam. Inc.	Percentile Ind. Inc.	Teen birth	Jail	Employed	Low Pov. nbhd.
ATT	-0.013 (0.020)	-0.009 (0.018)	-0.007 (0.012)	0.000 (0.009)	-0.000 (0.013)	-0.000 (0.002)	0.005 (0.004)	-0.026 (0.022)
Baseline	0.237	0.263	0.510	0.526	0.164	0.012	0.749	0.501
% $\Delta$	-0.054	-0.035	-0.013	0.001	-0.000	-0.001	0.007	-0.052
N Treat	1,172	1,172	1,172	1,172	1,171	1,172	1,172	1,172
N Control	549	549	549	549	549	549	549	549

The ATT row reports the average treatment effect on the treated of rent control on the average tract outcome. The ATT estimates are generated using a nearest neighbor Mahalanobis distance matching estimator that accounts for tract, city and county-level traits. The baseline row represents the average outcome of the matched counterfactual tracts. \* =  $p < 0.1$ , \*\* =  $p < 0.05$ , \*\*\* =  $p < 0.01$ .

Table 7: Mahalanobis distance nearest neighbor match estimates of average treatment effect on the treated of rent control on long-term outcomes, weighted by children linked to a tract

	Top 20% Fam.	Top 20% Ind.	Percentile Fam. Inc.	Percentile Ind. Inc.	Teen birth	Jail	Employed	Low Pov. nbhd.
ATT	-0.009 (0.013)	-0.007 (0.013)	-0.004 (0.009)	0.002 (0.007)	-0.003 (0.016)	-0.001 (0.002)	0.005 (0.005)	-0.012 (0.020)
Baseline	0.219	0.248	0.499	0.518	0.183	0.014	0.752	0.477
% $\Delta$	-0.042	-0.026	-0.008	0.005	-0.017	-0.091	0.006	-0.024
N Treat	1,172	1,172	1,172	1,172	1,171	1,172	1,172	1,172
N Control	549	549	549	549	549	549	549	549

The ATT row reports the average treatment effect on the treated of rent control on the average tract outcome, weighted by the number of children linked to each tract. The ATT estimates are generated using a nearest neighbor Mahalanobis distance matching estimator that accounts for tract, city and county-level traits. The baseline row represents the average outcome of the matched counterfactual tracts.

\* =  $p < 0.1$ , \*\* =  $p < 0.05$ , \*\*\* =  $p < 0.01$ .

Table 8: Mahalanobis distance nearest neighbor match estimates of average treatment effect on the treated of rent control on location as child and adult: high rental tract sample

	Frac. years in tract	Stay home	Stay Tract	Stay Comm. Zone
ATT	0.061** (0.027)	0.047** (0.020)	0.028* (0.017)	0.019 (0.021)
Baseline	0.507	0.204	0.207	0.737
% $\Delta$	0.121	0.230	0.137	0.026
N Treat	827	826	826	826
N Control	372	373	372	372

The ATT row reports the average treatment effect on the treated of rent control on the average tract outcome. The unit of observation is a census tract with at least 30% rental share. The ATT estimates are generated using a nearest neighbor Mahalanobis distance metric matching estimator. The baseline row represents the average outcome of the matched counterfactual tracts.

\* =  $p < 0.1$ , \*\* =  $p < 0.05$ , \*\*\* =  $p < 0.01$ .

Table 9: Mahalanobis distance nearest neighbor match estimates of average treatment effect on the treated of rent control on long-term outcomes: high rental tract sample

	Top 20% Fam.	Top 20% Ind.	Percentile Fam. Inc.	Percentile Ind. Inc.	Teen birth	Jail	Employed	Low Pov. nbhd.
ATT	0.011 (0.012)	0.009 (0.012)	0.012 (0.010)	0.016* (0.008)	-0.001 (0.015)	0.002 (0.002)	0.020* (0.010)	-0.018 (0.021)
Baseline	0.181	0.218	0.466	0.492	0.188	0.012	0.726	0.440
% $\Delta$	0.059	0.039	0.027	0.033	-0.005	0.168	0.027	-0.042
N Treat	840	840	840	840	839	840	840	840
N Control	373	373	373	373	373	373	373	373

The ATT row reports the average treatment effect on the treated of rent control on the average tract outcome. Only tracts with more than 30% rental share are included in the sample. The ATT estimates are generated using a nearest neighbor Mahalanobis distance matching estimator that accounts for tract, city and county-level traits. The baseline row represents the average outcome of the matched counterfactual tracts. \* =  $p < 0.1$ , \*\* =  $p < 0.05$ , \*\*\* =  $p < 0.01$ .



Table 10: Mahalanobis distance nearest neighbor match estimates of average treatment effect on the treated of rent control on long-term predicted outcomes

	Top 20% Fam.	Top 20% Ind.	Percentile Fam. Inc.	Percentile Ind. Inc.	Teen birth	Jail	Employed	Low Pov. nbhd.
<i>Panel A: ATT estimates for children at P25 of parent income distribution</i>								
Estimate	-0.008 (0.018)	-0.006 (0.017)	-0.002 (0.010)	0.003 (0.008)	-0.009 (0.017)	0.001 (0.002)	0.013** (0.007)	-0.015 (0.019)
Baseline	0.172	0.206	0.452	0.480	0.210	0.014	0.714	0.459
% $\Delta$	-0.049	-0.027	-0.003	0.007	-0.045	0.039	0.019	-0.032
<i>Panel B: ATT estimates for children at P75 of parent income distribution</i>								
Estimate	-0.002 (0.012)	0.000 (0.011)	-0.003 (0.008)	0.006 (0.005)	0.003 (0.010)	-0.001 (0.001)	0.006 (0.007)	-0.014 (0.018)
Baseline	0.271	0.303	0.554	0.566	0.108	0.008	0.792	0.537
% $\Delta$	-0.008	0.001	-0.005	0.010	0.024	-0.081	0.007	-0.026
N Treat	1,172	1,172	1,172	1,172	1,171	1,172	1,172	1,172
N Control	549	549	549	549	549	549	549	549

Estimates are reported for the predicted outcomes of children growing up in tracts at the 25th and 75th percentile of parent income. The predicted outcomes are generated by regressing individual outcomes on a transformation of parent income rank and recovering the fitted values at these two points of the parent income rank. The ATT is estimated using the baseline Mahalanobis distance matching estimator.

\* =  $p < 0.1$ , \*\* =  $p < 0.05$ , \*\*\* =  $p < 0.01$ .

Table 11: Mahalanobis distance nearest neighbor match estimates of average treatment effect on the treated of rent control on long-term predicted outcomes: high rental tract sample

	Top 20% Fam.	Top 20% Ind.	Percentile Fam. Inc.	Percentile Ind. Inc.	Teen birth	Jail	Employed	Low Pov. nbhd.
<i>Panel A: ATT estimates for children at P25 of parent income distribution</i>								
Estimate	0.003 (0.011)	0.004 (0.011)	0.004 (0.009)	0.007 (0.007)	-0.012 (0.016)	0.001 (0.003)	0.012 (0.010)	-0.014 (0.022)
Baseline	0.140	0.180	0.432	0.467	0.231	0.014	0.715	0.408
% $\Delta$	0.019	0.021	0.009	0.015	-0.050	0.094	0.016	-0.034
<i>Panel B: ATT estimates for children at P75 of parent income distribution</i>								
Estimate	0.029** (0.012)	0.022** (0.011)	0.022** (0.010)	0.021** (0.009)	0.016 (0.017)	0.000 (0.002)	0.016* (0.009)	0.006 (0.021)
Baseline	0.222	0.269	0.515	0.543	0.106	0.008	0.778	0.474
% $\Delta$	0.132	0.082	0.043	0.039	0.155	0.004	0.020	0.012
N Treat	840	840	840	840	839	840	840	840
N Control	373	373	373	373	373	373	373	373

Estimates are reported for the predicted outcomes of children growing up in tracts at the 25th and 75th percentile of parent income. Only tracts with more than 30% rental share are included in the sample. The predicted outcomes are generated by regressing individual outcomes on a transformation of parent income rank and recovering the fitted values at these two points of the parent income rank. The ATT is estimated using the baseline Mahalanobis distance matching estimator. \* =  $p < 0.1$ , \*\* =  $p < 0.05$ , \*\*\* =  $p < 0.01$ .

## A Opportunity Insight Data Description

As discussed in the body of the paper, the Opportunity Insights includes unconditional mean outcome estimates for all children linked to a given census tract. In addition, they provide predicted outcomes for children growing up at 5 specific points of the parent income distribution. Using data on the income level of parents, as well as the race, gender and census tract that children grew up in, they run the following regression:

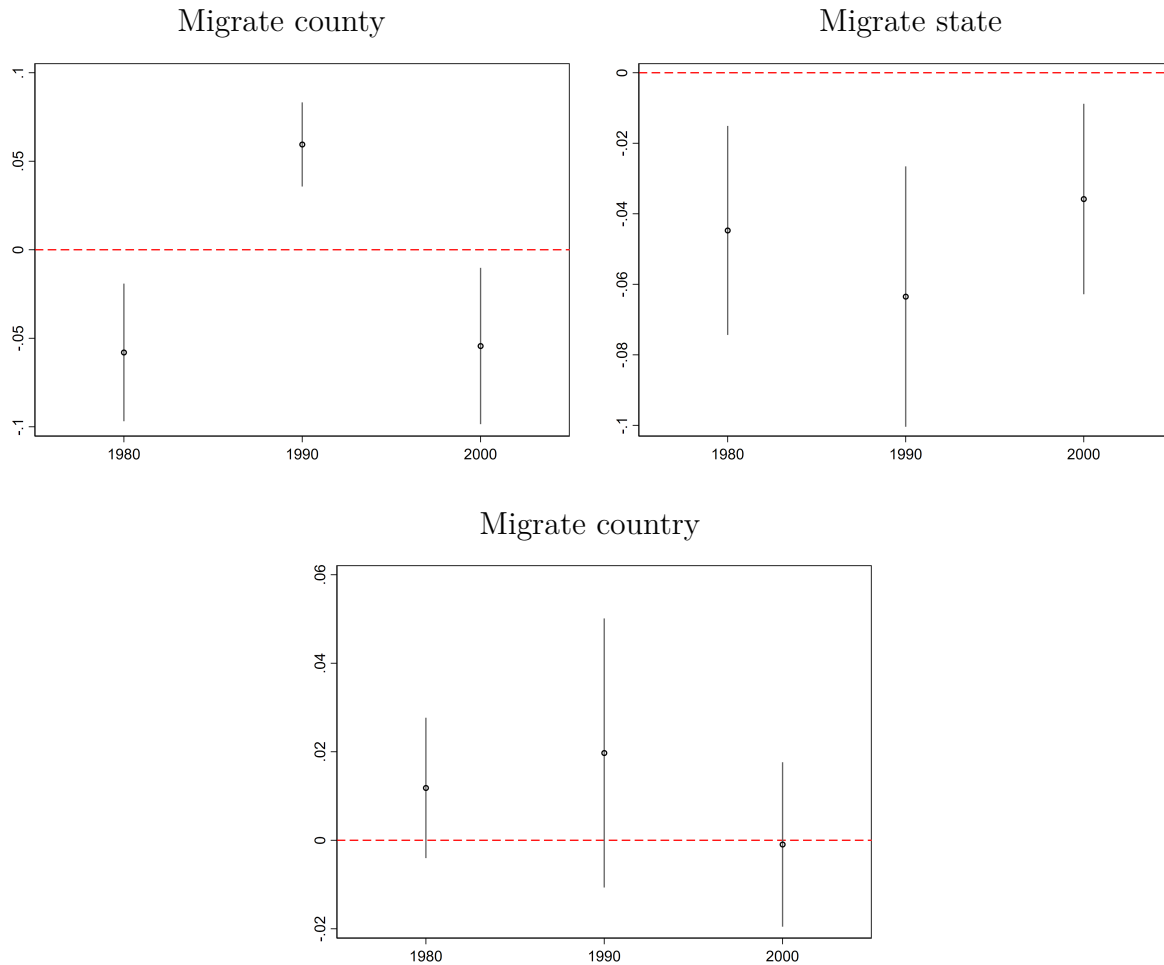
$$y_i = \alpha_{crg} + \beta_{crg} \times f_{rg}(p_i) + \epsilon_i$$

where  $y_i$  is an outcome for child  $i$ ,  $\alpha_{crg}$  is a fixed effect for the combination of census tract ( $c$ ), race ( $r$ ) and gender ( $g$ ).  $f_{rg}(p_i)$  refers to a transformation of parental income rank that is estimated at the national level using local polynomial regression. This is done because the true relationship between parental income and the outcome of children is often non-linear, so the transformed variable allows for the capture of these non-linearities within a simple regression framework. Once the individual regression is estimated, the authors use predicted values at a given level of parental income for each census tract-race-gender combination to create outcome predictions at the census tract level. The outcome data released to the public is the average fitted value over the entire census tract for the given population. Lastly, the authors add a small amount of noise to the estimates to avoid involuntary disclosure before releasing the data to the public.

The noise infusion procedure is detailed in Chetty and Friedman (2019), though the main intuition behind the algorithm is that the amount of noise added is proportional to the sensitivity of the statistic to one observation. For each tract, they determine the observation that has the largest impact on the result by estimating the result in the absence of each observation in the tract. Once they determine the maximum observed sensitivity, they add noise to the final result with a mean zero, normally distributed term that has a standard deviation proportional to the maximum observed sensitivity. Chetty *et al.* (2018) report that the added noise is generally smaller than the sampling variance. This will be especially true in tracts located in cities that have higher populations since the influence of any one observation has a smaller effect on the final output.


## B Appendix Tables and Figures

Figure A.1: Estimated ATT of rent control on immigration outcomes by year: high rental tract sample



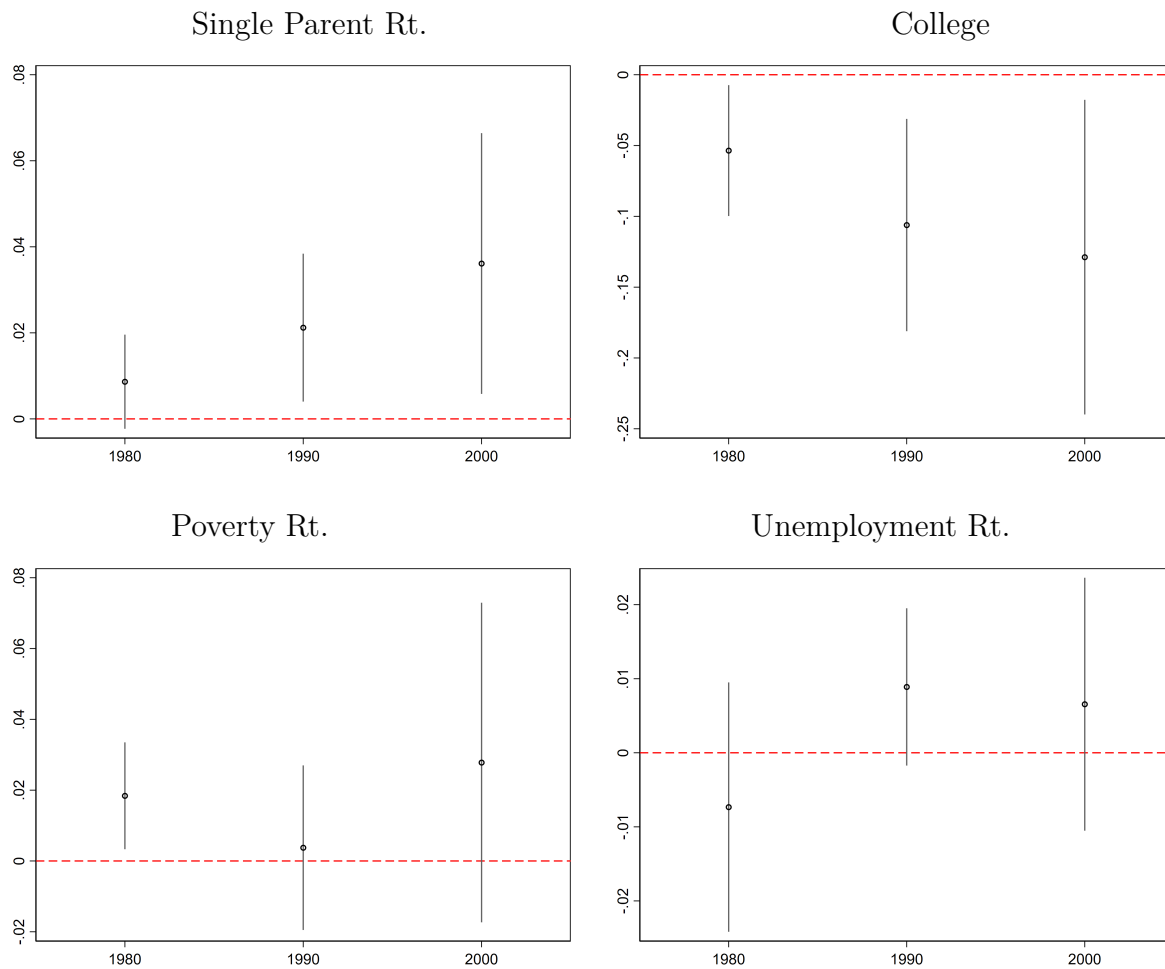
Notes: Figures show the average treatment effect on the treated tracts of rent control on immigration outcomes. The unit of observation is a census tract with at least 30% rental share. Each outcome is the percentage of tract inhabitants that have moved from a given location in the last 5 years. The error bars represent 95% confidence intervals from standard errors that are clustered at the city level.

Figure A.2: Estimated ATT of rent control on housing outcomes by year: high rental tract sample



Notes: Figures show the average treatment effect on the treated tracts of rent control on housing outcomes. The unit of observation is a census tract with at least 30% rental share. The average outcomes are generated using the baseline nearest neighbor match model. The error bars represent 95% confidence intervals from standard errors that are clustered at the city level.

Figure A.3: Estimated ATT of rent control on demographic outcomes



Notes: Figures show the average treatment effect on the treated tracts of rent control on employment and demographic outcomes. The unit of observation is a census tract with at least 30% rental share. The average outcomes are generated using the baseline nearest neighbor match model. The error bars represent 95% confidence intervals from standard errors that are clustered at the city level.

Table A.1: Comparing means and variances of the raw and weighted samples using the two-step nearest neighbor matching algorithm: high rental tract sample

	Std. Mean Diff.		Var. Ratio	
	Raw	Matched	Raw	Matched
Population	-0.033	-0.199	0.567	1.128
Male (%)	-0.101	0.008	1.028	1.225
Pop./sq. mile	0.868	-0.006	2.552	0.882
Age median	0.347	-0.079	1.060	1.150
white (%)	-0.189	-0.104	1.167	1.083
black (%)	0.086	0.064	1.126	1.099
Married (%)	-0.492	-0.135	0.910	1.652
Single parent fam. (%)	0.282	0.220	1.537	1.127
Educ. Less than HS (%)	-0.149	-0.198	0.861	1.340
Educ. HS (%)	0.113	0.050	0.679	1.330
LFP rate	0.145	-0.046	0.696	1.189
Unemployment rate	0.443	0.481	1.245	1.002
Avg. inc.	0.309	0.004	1.269	1.352
Family poverty rt. (%)	-0.142	0.114	0.683	0.941
Current addr. 5 years (%)	-0.019	-0.346	0.844	1.098
Housing units	0.072	-0.160	0.645	1.254
Rental (% of total units)	0.807	0.216	1.128	1.111
Rental vacancy rate (%)	-0.478	0.051	0.358	1.130
Rent vacancy x Rental %	-0.223	0.088	0.572	1.191
Avg. rent	0.478	0.100	0.962	1.305
Avg. home value	0.782	0.315	1.404	1.221
City rental (% of total units)	1.560	0.248	1.249	0.933
City rental vacancy rate (%)	-0.720	-0.298	0.311	0.768
City white (%)	-0.444	0.060	0.816	0.536
City black (%)	0.256	-0.188	0.717	0.390
City unemployment rate	0.842	0.682	0.825	1.412
City avg. rent	0.994	0.165	0.417	0.752
County Dem. vote share 1968	0.792	0.024	0.641	0.232
County Wallace vote share 1968	-0.884	-0.136	0.019	0.735
City population	0.935	0.218	3.001	0.753
City family poverty rt. (%)	-0.136	-0.029	0.387	0.888
City revenue per capita	0.952	0.328	2.169	1.264
City tax revenue per capita	0.959	0.061	2.667	0.760
Property tax share of rev.	0.406	-0.272	1.536	1.308
Other gov. sources share of rev.	0.180	-0.111	0.863	1.354
Educ. share of expenditure	0.426	-0.026	2.233	1.206
Police share of expenditure	-0.121	-0.406	0.377	0.465
Welfare share of expenditure	0.270	-0.077	4.426	0.571
N	15,199			
N Treat	840			
N unique control	373			

Table shows the standardized differences in means and variances between the raw and weighted sample. The unit of observation is a census tract with at least 30% rental share. The treated tracts are matched to a control tract using a nearest neighbor Mahalanobis distance matching procedure. The Mahalanobis distance metric includes linear terms for tract-level, city-level and county-level characteristics.