Title
Essays in Urban and Labor Economics

Author
Asquith, Brian James

Publication Date
2017

UNIVERSITY OF CALIFORNIA,
IRVINE

Essays in Urban and Labor Economics

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Economics

by

Brian J. Asquith

Dissertation Committee:
Professor Jan Brueckner, Chair
Professor Marianne Bitler
Professor David Neumark

2017

DEDICATION

To my wife, Esther Volchek, who now knows more about rent control, abortion, Social Security, and clustering standard errors than she ever wanted to. This is also your blood, sweat, and tears.

And also to my three laptops, including the two have gone on to that big, Microcenter in the sky. This is for you, boys.

Contents

List of Figures

List of Tables

ACKNOWLEDGMENTS

I would like to thank Jan, Marianne, and David for all of their help and input. I hope you enjoyed supervising my Ph.D. as much as I enjoyed writing it. I'd also like to thank the shadow member of my dissertation committee, Nanneh Chehras, without whom I wouldn't have finished. I'd also like thank my friends in the foxhole: Jose Luis Luna Alpizar, Ian Finn, Henry Korman, and Mark Kagan. I'd like to thank Eileen Raney, Samantha Spallone, and Yael Katz for their excellent research assistance and support.

Lastly, I'd like to thank my parents and my brother because I never would have gotten here without you all. My family has been my strength from the beginning and I never forget it.

This dissertation was made possible by generous grants from the National Institutes on Aging (P01 AG029409), the Center for Economics and Public Policy, the UC Irvine Data Science Initiative, and the UC Irvine Economics Department.

CURRICULUM VITAE

Brian J. Asquith

Biography

- Date and Place of Birth: August 27, 1986, Cincinnati, Ohio, Citizenship: USA

Employment

- National Bureau of Economic Research, Cambridge, MA 2017-2018
  Post-Doctoral Fellow in the Economics of an Aging Workforce
- The W. E. Upjohn Institute for Employment Research, Kalamazoo, MI 2018-onwards
  Economist

Education

- Ph.D. Economics, University of California, Irvine, 2017
  Committee: Jan Brueckner (Chair), Marianne Bitler, David Neumark
- M.A. Economics, University of California, Irvine 2013
- B.A. Mathematics, Economics, Brandeis University, *cum laude* 2009

Publications

- Asquith, Brian 2016. "Beggar Thy Neighbor, Beggar Thy Neighborhood", Forthcoming in *Papers in Regional Science*

Working Papers

- "Rent Control and Evictions: Evidence from San Francisco", Job Market Paper
  - *2017, Winner of the Tiebout Prize in Regional Science for Best Graduate Student Paper, Western Regional Science Association*
- "Grandchildren and Grandparents' Labor Force Attachment"

- "The Long-Run Effects of Minimum Wages and Other Anti-Poverty Policies on Disadvantaged Neighborhoods", with David Neumark and Brittany Bass.

Works in Progress

- "Labor Market Networks and Recovery from the Great Recession" with Judy Hellerstein, Mark Kutzbach, and David Neumark.
- "A Theory of Love and War: The Vietnam Draft, Assortive Mating, and Crime in the 1980's and 1990's" with Nanneh Chehras, Ian Finn, and Jose Luis Luna Alpizar.
- "Unfortunate Sons: Conscription and Assortive Mating"
- "An Analysis of the Quality of the Neighborhood Change Database's Weights"
- "U.S. Job Flows and the 'China Shock'" with Sanjana Goswami, David Neumark, and Antonio Rodriguez Lopez.

Awards and Research Grants

- NBER Post-Doctoral Fellowship on the Economics of an Aging Workforce 2017 – 2018
- Associate Dean's Fellowship, UCI — Spring 2017
- 31st Annual Tiebout Prize for Best Graduate Student Paper, Western Regional Science Association — 2017
- Seed Grant, National Institute on Aging P01 AG029409, "Grandchildren and Grandparents' Labor Force Attachment", $17,000 — 2015–2016
- Data Science Graduate Summer Fellowship, Data Science Initiative, UCI, $6,000 2015
- Borowski Student Policy Research Initiative, Center for Economics and Public Policy, UCI, $2,000 — 2016
- Economics Summer Research Fellowship, $5,025 UCI — 2014, 2015, 2016
- Associated Graduate Students Travel Grant, $800 UCI — 2014, 2016
- Merit Fellowship in Economics, UCI — 2011 – 2016
- Social Science Tuition Fellowship, UCI — 2011 – 2016

ABSTRACT OF THE DISSERTATION

Essays in Urban and Labor Economics

By

Brian J. Asquith

Doctor of Philosophy in Economics

University of California, Irvine, 2017

Professor Jan Brueckner, Chair

This dissertation answers three questions in urban and labor economics. Chapter 1 investigates under what conditions wealthy, high-skilled landowners wind up worse off when they use housing regulations to make their communities harder for poorer, low-skill workers to move into. If there is labor complementarity in production between high-skill and low-skill workers, low-skill workers will have lower productivity when they cannot advantage themselves of neighborhood effects, which in turn will hurt high-skill workers. Chapter 2 asks whether rent control, a strong form of housing regulations, incentivizes landlords to evict their tenants when the unregulated rents rise. Exploiting a known price shock to different residential buildings throughout San Francisco over an 11 year period, I find that a 2% increase in prices leads to a 1% increase in the monthly probability of an eviction from a controlled unit. The analysis also highlights the perverse incentives created by rent control, that medium-term (3-5 year) market withdrawals increase when prices increase. Chapter 3 studies whether grandchildren change their grandparents' labor force participation, and then researches whether the fall and rise in labor force participation between 1970 and 2009 can be ascribed to the rise and fall in grandparenthood due largely to the Baby Boom. I find that being a grandparent lowers labor force attachment for both grandfathers and grandmothers on both the intensive and extensive margin, but that grandchildren play only a limited role in the trends in older men's labor force participation.

Chapter 1

Beggar Thy Neighbor, Beggar Thy Neighborhood

The effects of restrictions on housing and land consumption are well-understood and researched in the economics literature. Similarly, the role of neighborhood and peer effects on educational outcomes is a well-studied phenomenon. Yet, there is a lacuna in the literature on the interaction between housing consumption restrictions and neighborhood effects. Namely, do certain housing restrictions interfere with neighborhood effects that benefit the poor? And do housing restrictions yield unintended consequences for the wealthy and poor alike?

Recognizing that artificially creating housing shortages in a neighborhood can reduce the opportunities for workers to learn from higher-skilled peers is key to understanding how laws effectively limiting the housing supply may influence local labor markets. We might therefore expect that economically segregated cities will have a lower rate of skill diffusion from high-income, high-skill individuals to others. In light of these considerations, restrictive housing policies can have negative externalities affecting socio-economic mobility and the welfare of all individuals across income levels.

Labor market externalities generated by community housing policies, and who bears the associated costs of these externalities, will be addressed. This paper models a city with two neighborhoods that has both a labor market and a housing market. The skill level of each resident is exogenous but is affected by neighborhood effects, so that the low-skill individuals' productivity is partially determined by the skill diffusion from high-skill individuals living in the same neighborhood. Further, high-skill and low-skill labor are complements in production, meaning that both types benefit from the increased efficiency of the other. Individuals consume land in the housing market, but neighborhoods can enact a local minimum land consumption requirement. Because impacting the labor market is not the typical intent of zoning policy, changes in labor market outcomes caused by the consumption restriction are regarded as unintended consequences of the new regulation. Collectively, this model's features allow us to answer the questions posed above.

The minimum land consumption requirement allows the rich to acquire more land by reducing competition from poor residents, who are forced to move to the other neighborhood. [1] Changing the demographics of the neighborhood in this way is often the intent of land-use restrictions, as explained by Hughes and Turnbull [64]. As a result, residents who are priced out will be unable to take advantage of the positive productivity spillovers from the rich, which they enjoyed from living in the same neighborhood. These positive spillovers are well-documented in the literature.[2] For example, peer effects can increase student GPA (Sacerdote [103], and moving from a poor to a wealthier neighborhood improves educational outcomes (Katz, Kling, and Leibman [71], Aaronson [1], and Fauth, Leventhal, and Brooks-

[1]The minimum land consumption requirement is analogous to a real-world zoning policy called minimum lot size zoning. Brueckner [26] gives a theoretical treatment of the implications for the housing market of implementing minimum land consumption restrictions. Glaeser and Ward [50] empirically show that minimum lot size zoning causes a fall in permits for new housing construction and an ambiguous effect on prices. Glaeser and Ward's analysis does show that communities that are not sufficiently dense do not maximize their aggregate land values, and thus also not their social welfare.

[2]See Dietz [40] for a discussion of the types of neighborhood effects and their mechanisms, including a review of how positive externalities generated by the wealthy or the educated emerge.

Gunn [45]). These and other papers strongly imply that the poor gain by living among the rich, and policies reinforcing income segregation make the poor less productive. Thus, the effects of the land-use restriction imposed in the model extend beyond the housing market.

The rich also benefit from a having a more efficient low-skill labor force, with labor complementarity ensuring that the rich have a stake in the poor's productivity.[3] Eeckhout, Pinheiro, and Schmidheiny [43] find strong evidence in U.S. data for "extreme-skill" complementarity, where high-skill individuals' productivity is boosted by the presence of low-skill workers. Therefore, the poor's productivity loss would be of concern to the rich, with the loss reducing their own incomes through the complementarity effect. This outcome could be one of the unintended consequences discussed above that could undermine the case for imposing the restriction in the first place.

The paper's approach advances our understanding of housing-market and labor-market interactions by going beyond previous theoretical work on income segregation via a combination of housing regulations and peer effects. Some papers, such as Fernandez and Rogerson [46] and Calabrese, Epple, and Romano [30], employ a minimum housing consumption requirement, which is also incorporated into this paper, to demonstrate that a primary outcome of these regulations is to increase income segregation. This paper will go beyond both of these studies by adding a citywide labor market that all residents participate in, but will do so at the expense of simplifying their setup. Fernandez and Rogerson assign a continuous distri-

[3] Selecting where to live is an endogenous choice, and the subset of the poor who choose to live among the rich is of great interest. They would be the ones most likely to move up the income ladder, and the presence of land-use restrictions would only make that ascent harder. Some authors have shown that the qualities of families who choose to live in certain neighborhoods account for a large amount of the variation in outcomes. Evans, Oates, and Schwab [44] show that using simultaneous equation models to account for neighborhood self-selection completely removes the neighborhood effects on teen pregnancy and high school dropout rates. Similarly, Oreopolous [96] uses an administrative dataset from Toronto that allows him to track children randomly assigned to different housing projects until they were more than 30 years old and found that family differences were much stronger than neighborhood differences in explaining labor market outcomes. For the purposes of this paper, these concerns are largely irrelevant. If it is only the motivated poor who take advantage of the neighborhood effects, this outcome complements the story this paper is presenting by showcasing how even the most "motivated" among the poor are frustrated in their attempt to improve their skill level.

bution of incomes to their city residents, while this paper considers discrete income groups, an approach that helps to better identify outcomes on specific worker types. The restricted community-wide housing consumption level is chosen endogenously in both Fernandez and Rogerson and Calabrese et al., but in this model, the finding from Calabrese et al. that the pivotal voter will impose binding zoning restrictions is incorporated directly as an exogenous change, so as to sharpen the emphasis on labor market outcomes.

This paper also expands on other work that explores the interplay between labor market outcomes, peer effects, and income segregation. Benabou [19] examines how high-skill and low-skill individuals choose to segregate into separate communities in the presence of neighborhood effects. Benabou includes, as this paper does, labor complementarity between high- and low-skill individuals, and finds that a negative externality of segregation is that it reduces the productivity of the city as a whole. However, Benabou's model did not include a housing market, and so while in this paper skills will not be endogenously chosen, it expands upon Benabou's work by introducing a policy tool that induces income segregation by neighborhood. Creating a model containing both a labor market and a housing market thus offers a clear innovation over previous models.

By adding more emphasis than other authors to the role of neighborhood effects in the interplay between labor and housing markets, certain other simplifications have to be made in order to arrive at intelligible conclusions. The main simplification is that unlike in Benabou, the high-skill are confined to one neighborhood at the outset, effectively incorporating directly into the model Benabou's finding that income segregation is a stable equilibrium. This immobility makes the model more tractable, but also has advantages in plausibility and improving the reader's intuition about how land regulations, labor markets, and neighborhood effects might interact.

This simplification also reflects the very strong likelihood that neighborhood effects and peer effects are probably less effective if high-skill workers are too dispersed. In Chetty and Hendren [33], the authors use a quasi-experimental method to show that for every 1 percentile increase in the average permanent resident's income, a child who moves into that neighborhood improves his or her expected adulthood income percentile position by 0.03-0.04 for every year of residence. Taking these results as a baseline, if a child moves from a neighborhood where the average income percentile is the 10th into one where the percentile is the 50th at the age of 1, he or she is expected to move up 28.8 percentiles. However, if a child moves from a 10th percentile neighborhood to an 80th percentile neighborhood, the child's expected adulthood income moves up 50 percentiles, a substantive difference. The latter neighborhood is more likely to be the result of income segregation by the rich, and concordantly, has more powerful neighborhood effects. Thus, concentration of the rich is an advantageous and parsimonious assumption.

The literature therefore strongly suggests that there is a link between a neighborhood's housing consumption patterns and the production externalities generated by high-skilled individuals, while leaving unanswered what the interaction of these forces might be. This model will contribute to our understanding of this connection by using a two-neighborhood model with the rich confined to one neighborhood but the poor mobile, with the analysis deriving the equilibrium allocation of poor residents across the two neighborhoods. The rich then impose a minimum lot size restriction, and the ramifications of this policy are explored for each class individually, and for society as a whole.

Using numerical methods, the paper shows that in the presence of labor complementarities and productivity spillovers (neighborhood effects), imposing a minimum lot size restriction may lead to welfare losses for all segments of the population. In the absence of complementarity and with low spillover effects, the wealthy may benefit from a minimum lot size restriction, although society as a whole sees a net loss. The magnitude of the social loss varies with the magnitude of the complementarity and neighborhood effects and the level of the minimum lot size requirement, with a small, isolated rich population increasing that group's incentive to impose a restriction, but only if complementarity and productivity spillovers are low.

1.1 Model

The city contains two neighborhoods. This could be thought of as being akin to a city divided into two by a river. The population of the city has two types of workers, innately high-skilled and innately low-skilled workers. These workers act as a collective labor supply for the city, and have identical preferences.

All labor participates in a city-wide labor market, being employed by a firm using the production function

$$F = \gamma E_H + \beta E_L + \delta(E_H E_L). \tag{1.1}$$

where E_H and E_L are high-skilled and low-skilled efficiency units of labor in the entire city, respectively, δ is a measure of labor complementarity, and $\gamma > \beta$. Labor complementarity is experienced by all workers across both neighborhoods. The total numbers of people of each skill type are given as

$$N_L = N_{L1} + N_{L2} \text{ and } N_H = N_{H1} + N_{H2}, \tag{1.2}$$

where N_{ij} is the number of people of type $i = \{L, H\}$ in neighborhood $j = \{1, 2\}$. Each rich individual provides one efficiency unit, so that $E_H = N_H$, while $E_L = E_{L1} + E_{L2}$, where E_{Lj} denotes the efficiency units in the supply of low-skill people in neighborhood j.

All the high-skill workers live in one neighborhood and are thus immobile. Throughout the rest of this paper, the variables denoting low-skill individual's choices will be indexed by neighborhood but those for the high-skill will not (with their residence in neighborhood 2 implicit).

The low-skill workers enjoy a positive externality from living in the same neighborhood as high-skill individuals, gaining efficiency units in their supply of labor from exposure to high-skill workers in their neighborhood. This relationship is given as

$$E_{L2} = \sigma N_{L2},\ E_{L1} = N_{L1},\ E_L = N_{L1} + \sigma N_{L2} = N_L + (\sigma - 1)N_{L2} \tag{1.3}$$

where σ is a constant greater than 1.

The production function can then be rewritten as

$$F = \gamma N_H + \beta[N_L + (\sigma - 1)N_{L2}] + \delta N_H(N_L + (\sigma - 1)N_{L2}), \tag{1.4}$$

From the production function, the wages for each group are thus

$$w_H \equiv \frac{\partial F}{\partial N_H} = \gamma + \delta(N_{L1} + \sigma N_{L2}), \tag{1.5}$$

$$y_{L1} \equiv \frac{\partial F}{\partial N_{L1}} = \beta + \delta N_H, \tag{1.6}$$

and

$$y_{L2} \equiv \frac{\partial F}{\partial N_{L2}} = \sigma(\beta + \delta N_H). \tag{1.7}$$

For consistency of the model, σ is constrained to ensure that $w_H > y_{L2}$, which requires $\sigma < w_H/y_{L1}$. But because N_L is assigned to be significantly greater than N_H, this constraint will be satisfied.

Following Brueckner and Lai [27], the city has resident landlords who collect rental income in addition to wage income. The high-skilled residents are the landlords in this city, so that each high-skill individual collects $(r_1 + r_2)/N_H$ where r_j is the land rent in neighborhood j. Their combined income from wages and rents is thus

$$y_H = w_H + \frac{r_1 + r_2}{N_H}. \tag{1.8}$$

The above condition means that in addition to $w_H > y_{L2}$, $y_H > y_{L2}$. Having high-skill income be a function of both wages and rents is a key part of the model's structure, and it matches a well-known stylized fact that many wealthy people only get a portion of their income from compensation.[4] The presence of this rental increase allows for a richer exploration of the trade-offs between increasing high-skill land consumption and potential losses from changes to income when low-skill residents switch neighborhoods.

To solve for the optimal consumption and settlement patterns, first assume that the land areas of the neighborhoods are fixed and equal, each being normalized to unity. The land

[4] For example, Rosenberg [100] shows that compensation income drops down to only 34% of total income for earners in the top 0.1% of the distribution.

constraints can therefore be expressed as

$$1 = N_{L1} q_{L1} \tag{1.9}$$

and

$$1 = N_{L2} q_{L2} + N_H q_H, \tag{1.10}$$

where q_i is the land consumption of an individual of type $i = H$, $L1$, $L2$.

The other good consumed by all individuals is a composite consumption good, c_i. Preferences are Cobb-Douglas, given by

$$U_i = c_i^{\alpha} q_i^{1-\alpha}, \tag{1.11}$$

and utility is maximized subject to the budget constraint $y_i = c_i + r_j q_i$, $i = H$, $L1$, $L2$, $j = 1$, 2. Low-skill individuals are freely mobile between the two neighborhoods, and they will move until utility is equalized between the two neighborhoods, so that

$$U_{L1} = U_{L2}. \tag{1.12}$$

The marginal products here will more than exhaust the total product, but if the firm is a monopolist, it could still earn a positive profit.

1.2 Solving for the Equilibrium

Given the above structure, the equilibrium of the model can be computed. First, the optimal consumption bundle for each skill type is found and the equilibrium neighborhood population levels are determined. A single asterisk will indicate equilibrium values.

The first step is to find the optimal consumption bundles as a function y_i and r_j, which equal

$$q_i^* = \frac{(1-\alpha)y_i}{r_j}, \tag{1.13}$$

$$c_i^* = \alpha y_i. \tag{1.14}$$

Land rent, r_j, can be computed by using the neighborhood land constraints. Substituting (13) into (9) and (10) and rearranging yields

$$r_1 = (1-\alpha)N_{L1}y_{L1} = (1-\alpha)(N_L - N_{L2})y_{L1}, \tag{1.15}$$

$$r_2 = (1-\alpha)(N_{L2}y_{L2} + N_H y_H). \tag{1.16}$$

The optimal land consumption bundles can then be expressed as

$$q_{L1} = \frac{1}{N_{L1}}, \tag{1.17}$$

$$q_{L2} = \frac{y_{L2}}{N_{L2}y_{L2} + N_H y_H}, \tag{1.18}$$

and

$$q_H = \frac{y_H}{N_{L2}y_{L2} + N_H y_H}. \tag{1.19}$$

It is clear from inspection that $q_H > q_{L2}$, because $y_H > y_{L2}$ is assumed to hold. Substituting (15) and (16) back into (8) gives

$$y_H = \gamma + \delta(N_L + (\sigma - 1)N_{L2}) + \frac{(1-\alpha)N_{L1}y_{L1} + (1-\alpha)(N_{L2}y_{L2} + N_H y_H)}{N_H}. \tag{1.20}$$

Equation (20) can be rearranged to isolate y_H, yielding

$$y_H = \frac{N_H\left(\gamma + \delta(N_L + (\sigma-1)N_{L2})\right) + (1-\alpha)\left((N_L - N_{L2})y_{L1} + N_{L2}y_{L2}\right)}{\alpha N_H}. \tag{1.21}$$

N_{L1} and N_{L2} can be determined by using the equal utility constraint (12) and the low-skill population constraint given in (2). Substituting into (12) using (14), (17), and (18), equal utilities require

$$(\alpha y_{L1})^{\alpha}(1/N_{L1})^{1-\alpha} = (\alpha y_{L2})^{\alpha}\left(\frac{y_{L2}}{N_{L2}y_{L2}+N_H y_H}\right)^{1-\alpha}. \tag{1.22}$$

Equation (22) can be rearranged to yield

$$N_{L1} = N_L - N_{L2} = \frac{\left(N_{L2}y_{L2} + \left(\gamma + \sigma_2 \delta N_{L2}\right) N_H\right) \sigma_2^{\frac{\alpha}{\alpha-1}}}{\alpha y_{L2} - \sigma_2^{\frac{\alpha}{\alpha-1}} \left(\delta N_H + (1-\alpha) y_{L1}\right)}. \tag{1.23}$$

Substituting $y_{L2} = \sigma y_{L1}$ and solving (23) for N_{L2} yields N_{L2}^{*}:

$$N_{L2}^{*} = \frac{N_L \left(\alpha \sigma^{\frac{1}{1-\alpha}} y_{L1} - \left(\delta N_H + (1-\alpha) y_{L1}\right)\right) - \gamma N_H}{\alpha \sigma^{\frac{1}{1-\alpha}} y_{L1} + (\sigma - 1 + \alpha) y_{L1} + \delta \left(\sigma - 1\right) N_H}. \tag{1.24}$$

1.3 Comparative Statics

The change in neighborhood 2's low-skill population after an increase in γ is clear. Since the denominator of (24) is positive,

$$\frac{\partial N_{L2}^{*}}{\partial \gamma} < 0, \tag{1.25}$$

follows from inspection. Thus, as high-skill productivity increases more low-skill people move into neighborhood 1.

The change in neighborhood 2's low-skill population as the productivity parameter β increases is harder to determine, as the derivative is very involved. However, it can be shown that $\frac{\partial N_{L2}}{\partial \beta} > 0$ holds, showing that low-skill workers move into neighborhood 2 as their productivity increases.

For the low-skill workers, it is clear then that they move into the mixed neighborhood as their productivity rises relative to that of the high-skill workers and leave as it falls. For the high-skill people, y_H rises with N_{L2}, as shown in the derivative

$$\frac{\partial y_H}{\partial N_{L2}} = \frac{\delta(\sigma-1)}{\alpha} + \frac{(1-\alpha)(\sigma-1)y_{L1}}{\alpha N_H} > 0. \tag{1.26}$$

The first term in the derivative is the increase in income due to complementarity effects in the high-skill wage, while the second term is due to the increase in rental income. It is then clear that as low-skill individuals leave, high-skill individuals suffer a loss in income due to lower wages and lower rents, but their land consumption, q_H, increases as low-skill individuals move out. This outcome can be seen in the derivative of q_H, where from (19)

$$\frac{\partial q_H^*}{\partial N_{L2}} = (N_{L2}y_{L2}+N_H y_H)^{-1}\frac{\partial y_H}{\partial N_{L2}} - y_H(N_{L2}y_{L2}+N_H y_H)^{-2}\Big(\frac{\partial y_H}{\partial N_{L2}}N_H + y_{L2}\Big) < 0. \tag{1.27}$$

The inequality can be signed after substituting (26) and rearranging. Clearly, if the gains from increased land consumption are greater than the loss in consumption of the numeraire good, c_H, due to lower y_H, then high-skill residents stand to benefit from inducing low-skill individuals to leave.

1.4 Introducing A Housing Restriction

As discussed in the introduction and shown in the previous section, the rich may have an incentive to restrict housing access in their neighborhoods by preventing too many low-skill individuals from moving in and occupying more land themselves. To understand when the rich might do so and what consequences this action will have, suppose that the high-skill workers can somehow impose a minimum lot size requirement, which states that all inhabitants of neighborhood 2 must consume at least q_{min} worth of land.

For the purposes of this paper, this constraint will be treated as exogenously imposed.[5] Not endogenizing the choice of a housing restriction and thereby permitting a policy choice that might result in a welfare loss for the choosers captures a situation where policymakers are ignorant or misinformed about the consequences of their decisions. For example, the high-skill workers may think that there is no labor complementarity with low-skill workers ($\delta = 0$). They would then support a housing policy commensurate with that assumption, and may only discover later they were wrong.

Assume that there therefore exists a q_{min}, the minimum amount of land that must be consumed, such that

$$q_{L2}^* < q_{min} \leq q_H^*. \tag{1.28}$$

[5] Fernandez and Rogerson () point out that viewing zoning requirements as exogenous is sensible because in many communities, the oldest housing requirements are effectively independent of present-day community characteristics and are not easily changed.

The only group initially affected are then low-skill people in neighborhood 2, whose land consumption must now equal q_{min} even though a smaller value is preferred.

With the q_{min} provision, the land consumption constraint (10) becomes

$$1 = N_H q_H + N_{L2} q_{min}. \tag{1.29}$$

Substituting $q_H = \frac{(1-\alpha)y_H}{r_2}$, land rent in neighborhood 2 can be written as

$$r_2 = \frac{N_H(1-\alpha)y_H}{(1 - N_{L2}q_{min})}. \tag{1.30}$$

The low-skill budget constraint then implies

$$c_{L2} = y_{L2} - \frac{N_H(1-\alpha)y_H}{1 - N_{L2}q_{min}} q_{min}. \tag{1.31}$$

The income for the high-skill residents changes to reflect the new rent function for neighborhood 2. Substituting (30) and (15) into (8) gives

$$y_H = (\gamma + \delta\left(N_L + (\sigma - 1)N_{L2}\right)) + \frac{(1-\alpha)y_H}{1 - N_{L2}q_{min}} + \frac{(1-\alpha)N_{L1}y_{L1}}{N_H}. \tag{1.32}$$

Rearranging (32) to solve for y_H yields

$$y_H = \left(\frac{1 - N_{L2}q_{min}}{\alpha - N_{L2}q_{min}}\right)\left(\frac{\left(\gamma + \delta\left(N_L + (\sigma - 1)N_{L2}\right)\right)N_H + (1-\alpha)(N_L - N_{L2})y_{L1}}{N_H}\right). \quad (1.33)$$

Equation (33) can be substituted into (30) to yield a more informative expression for r_2:

$$r_2 = (1-\alpha)\left(\frac{\left(\gamma + \delta\left(N_L + (\sigma - 1)N_{L2}\right)\right)N_H + (1-\alpha)(N_L - N_{L2})y_{L1}}{\alpha - N_{L2}q_{min}}\right). \quad (1.34)$$

From (33) and (34), positivity of y_H and r_2 requires

$$\alpha - N_{L2}q_{min} > 0, \quad (1.35)$$

a condition that will be assumed to hold (recalling that $\alpha < 1$).

The imposition of q_{min} changes N_{L1} and N_{L2}. To find these effects, first rewrite the equal utility condition (12) as

$$(\alpha y_{L1})^{\alpha}(1/N_{L1})^{1-\alpha} = (y_{L2} - r_2 q_{min})^{\alpha}\left(q_{min}\right)^{1-\alpha}. \quad (1.36)$$

Rearranging yields

$$N_{L1} = \left(\frac{y_{L2} - r_2 q_{min}}{\alpha y_{L1}}\right)^{\frac{\alpha}{\alpha-1}} \frac{1}{q_{min}}. \tag{1.37}$$

Now after substituting y_{L2}, and y_{L1} into (37) along with r_2 from (34), the population constraint (23) then yields

$$N_L = \left(\frac{y_{L2}(\alpha - N_{L2}q_{min}) - (1-\alpha)\left((\gamma + \delta\left(N_L + (\sigma-1)N_{L2}\right)\right)N_H + (1-\alpha)\left((N_L - N_{L2})y_{L1}\right)\right)q_{min}}{\alpha y_{L1}(\alpha - N_{L2}q_{min})}\right)^{\frac{\alpha}{\alpha-1}} \frac{1}{q_{min}} + N_{L2}. \tag{1.38}$$

There is no closed-form solution for N_{L2} from (38) unless $\alpha = 1/2$. However, (38) is used below in numerical analyses of the model. The solution to (38), denoted N_{L2}^{**}, is derived numerically and then substituted into (33) to get y_H^{**} and into (31) and (34) to get c_{L2}^{**} and r_2^{**}, respectively.

Despite the lack of a closed-form solution for N_{L2}, some important results can still be derived. Consider the effect on rent of an increase in q_{min}, starting with $q_{min} = q_{L2}^*$, while holding N_{L2} fixed at N_{L2}^*. From (34), the relevant derivative is

$$\left.\frac{\partial r_2}{\partial q_{min}}\right|_{N_{L2}=N_{L2}^*} = \frac{r_2^* N_{L2}^*}{\alpha - N_{L2}^* q_{min}} > 0. \tag{1.39}$$

Therefore r_2 rises above r_2^* as q_{min} increases above q_{L2}^*, holding N_{L2}^* fixed. The effect of this change on low-skill utility is found by differentiating

$$U_{L2} = (y_{L2} - r_2 q_{min})^{\alpha} \, q_{min}^{1-\alpha}, \tag{1.40}$$

again holding N_{L2} fixed and starting with $q_{min} = q^*_{L2}$ and $r_2 = r^*_2$. Since q^*_{L2} maximizes U_{L2} when $r_2 = r^*_2$, increasing q_{min} above q^*_{L2} lowers low-skill utility. When the assumption that $r^{**}_2 = r^*_2$ is relaxed, their utility falls further through the indirect effect operating through the increase in r_2 (recalling (39)). Thus, from (39) and (40), it can be seen that

$$\left.\frac{\partial U^{**}_{L2}}{\partial q_{min}}\right|_{N^{**}_{L2}=N^*_{L2}} < 0. \tag{1.41}$$

The decline in U_{L2} as q_{min} rises disrupts the equal utility condition, and equalizing utilities again between the neighborhoods requires a decline in N_{L2}. The decline in N_{L2} reduces r_2 from (34), which raises U_{L2}, and the population shift to neighborhood 1 reduces U_{L1}, as can be seen from (36). Ultimately, utilities are reequalized so that $U_{L1} = U_{L2}$. Thus,

$$\frac{\partial N^{**}_{L2}}{\partial q_{min}} < 0. \tag{1.42}$$

One implication of these findings is that the *overall* impact of imposing q_{min} on high-skill utility is ambiguous. After q_{min} is imposed, the high-skill workers' utility function becomes

$$U_H^{**} = (\alpha y_H^{**})^\alpha \left(\frac{1 - N_{L2}^{**} q_{min}}{N_H}\right)^{1-\alpha}, \tag{1.43}$$

which is found by substituting (14) and (19) into (11). To see how U_H^{**} changes with q_{min}, the derivate is computed:

$$\frac{\partial U_H^{**}}{\partial q_{min}} = \alpha \frac{\partial y_H^{**}}{\partial q_{min}} \left(\frac{\alpha y_H^{**}}{q_H^{**}}\right)^{\alpha-1} - \frac{(1-\alpha)}{N_H} \left(\frac{\alpha y_H^{**}}{q_H^{**}}\right)^\alpha \left(N_{L2}^{**} + \frac{\partial N_{L2}^{**}}{\partial q_{min}} q_{min}\right). \tag{1.44}$$

Since N_{L2} shrinks as q_{min} grows, the overall impact on land-consumption and utility is unclear. The next section discusses in greater detail what kinds of impacts the high-skill group can anticipate from an increase in q_{min}.

1.5 The Impact of q_{min} on the High Income Group

This section explores what happens to the high-skill workers' utility as q_{min} increases. The derivative of U_H with respect to q_{min} is analytically ambiguous (as shown in (44)), so to illustrate what high-skill residents can expect from an increase in q_{min}, the behavior of U_H is simulated in Figures 1-4. These simulations show how increasing q_{min} changes U_H under different assumptions on the level of labor complementarity (δ), productivity spillovers (σ), and the skill dispersion in the city (N_L/N_H).[6] Variation in δ and σ will illustrate the impact of labor market externalities from housing restrictions on high-skill utility, and variation in

[6]Increasing the ratio of γ to β produces an effect similar to changing the ratio of N_L to N_H, so for economy of presentation, only the difference in response to the N_L/N_H ratio is shown.

skill dispersion will show how income inequality changes the incentive to use housing market restrictions when labor market outcomes are fully incorporated.

Figures 1-2 assume that there are four times as many low-skill as high-skill people and Figures 3-4 assume that there are six times as many low-skill as high-skill people. Figures 1-2 and 3-4 can be thought of as comparing cities where the elite are a small group versus more broadly based. Likewise, the results in Figures 1 and 3 show the results in a "low"-σ scenario and Figures 2 and 4 show a "high"- σ scenario. Both values of σ are calibrated to ensure that $N_{L2}^* > 0$ when $q_{min} = q_{L2}^*$. In the graphs, the two values are given in increments above a baseline σ, denoted σ_0, which is about 1.49.[7] In each graph, lines for four different values of δ ($\delta = \{0, 1/5, 1/3, 1/2\}$) are included to demonstrate how labor complementarity impacts the outcome, and q_{min} varies from $q_{min} = q_{L2}^*$ to $q_{min} = q_H^*$ for given values of the parameters.

For the cases shown in Figures 1-4, incorporating labor market outcomes does not overturn the finding by others in the literature that high-income residents are better off after a housing restriction is imposed, although including labor market effects does reduce their potential utility gains. All low-skill residents exit neighborhood 2 well before q_{min} reaches q_H^*, which can be seen in the eventual flattening of the U_H lines in each graph. By contrast, high-skill residents improve their welfare when q_{min} is raised high enough that all low-skill residents are induced to move, but the utility increases are small in all cases.[8] Perfect income segregation yields lower incomes and utility for low-skill residents while not greatly improving high-skill resident's utility. As δ and σ rise, the high-skill resident's potential utility gains from perfect income segregation diminish. In fact, when $\delta = 1/2$ and $\sigma \approx 0.5 + \sigma_0$, perfect income segregation across the two neighborhoods leaves the high-skill residents about as well off as

[7]More specifically, the initial value of σ is chosen to make the numerator of (24) positive, such that $\sigma > \left(\frac{\gamma \frac{N_H}{N_L} + \delta N_H + (1-\alpha) y_{L1}}{\alpha y_{L1}}\right)^{1-\alpha}$, which for $\alpha = 1/2$, $\delta = 1$, $6N_H = N_L = 600, \frac{\gamma}{\beta} = 4$ is ≈ 1.4916, although the magnitude varies slightly with δ.

[8]Even in the no complementarity scenarios modeled, the high-skill residents never see a utility increase of more than 5%.

they are under the mixed-income equilibrium. Even so, the high-skill residents still improve their utility in the other scenarios if q_{min} is raised modestly.

Greater skill dispersion causes the gains from perfect income segregation to be larger. Utility increases more with q_{min} in Figures 3-4 than in Figures 1-2. This is because increasing N_L increases the competition for land in the mixed-income neighborhood, so the high-skill residents have more to gain from the exit of low-skill residents. Figures 5-8 show what happens to the high-skill residents' consumption of land and the numeraire good (directly proportional to y_H via (14)) as q_{min} increases.[9] Increasing q_{min} creates a trade-off for the high-skill residents of gaining more land consumption at the expense of less income and hence numeraire consumption. The results here are clear: as long as the utility gains from more land consumption outweigh income losses, then high-skill residents are better off from a housing restriction.

The results in Figures 1-8 suggest that the traditional view that imposing a minimum land-consumption requirement yields a higher utility level for high-skill (*i.e.*, high-income) residents is borne out. At best, factoring in labor market outcomes leaves the high-skill residents essentially indifferent to living in a income segregated city, but in most cases, they are better off. However, Figures 1-4 actually understate the role of both σ and δ because as δ increases, the wages of both types of worker automatically adjust upwards. This creates the level effect in δ that can be seen in all four graphs. Removing the automatic adjustment can be done by setting

$$\gamma = \gamma_0 - \delta(N_L + (\sigma - 1)N^*_{L2}) \tag{1.45}$$

[9]Land consumption is scaled up by a factor of $10e^5$ for clarity of comparison

and

$$\beta = \beta_0 - \delta N_H, \tag{1.46}$$

so that y_{L1} and y_{L2} are initialized at the same income level regardless of the value of δ. Likewise, y_H^* will now be the same across all values of δ. As q_{min} increases, y_{L1} and y_{L2} will remain the same (as they do when the level effect is not removed), but all values of w_H will fall from the same base level, γ_0.

Figures 9-12 show the changes in the high-skill worker's utility as q_{min} increases with the level effect in income removed. All four graphs are set on the same scale for ease of comparison. The findings from Figures 1-4 are effectively reversed: now, the high-skill experience utility losses from perfect income segregation except in the cases where complementarity and spillovers are low.[10] Labor market outcomes in fact make a substantive difference in whether q_{min} increases or decreases high-skill individuals' utility. Even low levels of complementarity and productivity spillovers yield utility losses, suggesting that high-skill workers' outcomes are indeed sensitive to small changes improvements in these parameters.

Comparing Figures 9-10 with 11-12, it is clear that there is a wider "spread" in outcomes as inequality increases. When complementarity is low or nonexistent, and the productivity spillover is low, high-skill residents in a more unequal society achieve greater utility gains from increasing q_{min}. This is a more nuanced finding than stereotypical "rich get richer" stories. If there is no complementarity, then greater inequality does increase the positive impact of q_{min} for the high-skill residents. But increasing q_{min} when there is complementarity leaves them more worse off the more unequal is the skill dispersion is in the city. Clearly, models that do

[10] While not shown here, as γ_0 rises in relation to β_0, the change in U_H becomes even more sensitive to the value of δ.

not take into account labor complementarity and spillover effects could greatly *understate* the welfare impacts of a housing consumption restriction.

Overall, these graphs demonstrate that if neighborhood effects and labor complementarity are taken into account, then the high-skill residents can be acting against their own interests by mandating a minimum lot size restriction. Even a small q_{min} with low levels of complementarity and modest productivity spillovers can cause them to experience a net decline in utility. This affirms the intuition of the model: the decision to impose a minimum lot size restriction may not fully anticipate for all of the repercussions to both the high- and low-skill workers. When the repercussions are accounted for, the high-skill residents should only impose a housing restriction if neighborhood effects are weak and complementarity is trivial.

1.6 Overall Welfare Effect

If the decision to raise q_{min} is made, what will be its effects on social welfare? Defining U_{all} as our overall social welfare, where

$$U_{all} \equiv U_H N_H + U_{L2} N_{L2} + U_{L1}(N_L - N_{L2}), \tag{1.47}$$

the social welfare will respond to a change in q_{min} across all three groups, but signing the change is difficult since we do not know the relative magnitudes of the terms below:

$$\frac{\partial U_{all}}{\partial q_{min}} = \frac{\partial U_H}{\partial q_{min}} N_H + \left(\frac{\partial U_{L2}}{\partial q_{min}} N_{L2} + U_{L2} \frac{\partial N_{L2}}{\partial q_{min}} \right) + \left(\frac{\partial U_{L1}}{\partial q_{min}} N_{L1} + U_{L1} \frac{\partial N_{L1}}{\partial q_{min}} \right)$$

This expression can be simplified since $U_{L1} = U_{L2}$, $\frac{\partial N_{L2}}{\partial q_{min}} = -\frac{\partial N_{L1}}{\partial q_{min}}$, and $\frac{\partial U_{L2}}{\partial q_{min}} = \frac{\partial U_{L1}}{\partial q_{min}}$:

$$\frac{\partial U_{all}}{\partial q_{min}} = \frac{\partial U_H}{\partial q_{min}} N_H + \left(\frac{\partial U_{L2}}{\partial q_{min}} N_{L2} - U_{L1} \frac{\partial N_{L1}}{\partial q_{min}} \right) + \left(\frac{\partial U_{L2}}{\partial q_{min}} N_{L1} + U_{L1} \frac{\partial N_{L1}}{\partial q_{min}} \right)$$

$$= \frac{\partial U_H}{\partial q_{min}} N_H + \frac{\partial U_{L2}}{\partial q_{min}} N_L \tag{1.48}$$

The second term is unambiguously negative, but from Figures 1-4 and 9-12 we know the first term is positive in the case of low values of δ and σ, but can become negative as δ and σ rise and the change in q_{min} exceeds a small increase. Figures 13-16 demonstrate that even in the case of no complementarity, social welfare still declines, and as mathematical intuition would suggest, declines more steeply as δ increases.

An other important consequence of this model is that production in the city declines as q_{min} rises:

$$\frac{\partial F}{\partial q_{min}} = \beta\sigma_2 \frac{\partial N_{L2}}{\partial q_{min}} + \beta \frac{\partial N_{L1}}{\partial q_{min}} + \delta N_H \frac{\partial N_{L2}}{\partial q_{min}} < 0, \tag{1.49}$$

which is true since $\frac{\partial N_{L2}}{\partial q_{min}} = -\frac{\partial N_{L1}}{\partial q_{min}}$ and $\beta\sigma_2 > \beta$.

1.7 Conclusion

Imposing an arbitrary minimum on land consumption (the proxy for housing in this model) leads to an unequivocal welfare loss for the low-skilled group and for society as a whole, but the welfare change for the high-skilled group is dependent on the level of labor complementarity and productivity spillovers. By inducing the low-skill residents to relocate out of the mixed-skilled neighborhood, society becomes overall less productive.

This result suggests several answers to the questions posed in the introduction. The first is that except for small values of δ and σ, imposing even small increases in q_{min} leads to a utility loss for high-skilled workers. Productivity spillovers amplify the impact of labor complementarity, such that increasing either can switch expelling all of the low-skill individuals out of the neighborhood from a net utility gain to a net utility loss for the high-skill residents. The second is that this outcome has interesting implications for income inequality, suggesting that societies where wealth is concentrated in a small elite who are isolated from the poor stand to gain the most from reinforcing inequality through government intervention. However, if spillovers and complementarity are substantial, income elites in more unequal societies can in fact be left worse off from housing market restrictions. The third is that restricting access to housing leads to a loss in productivity, as the low-skill individuals are pushed into places where they cannot enjoy any neighborhood or peer effects.

In cases where the rich impose the restriction that results in a welfare loss to them, the interpretation would be that they are not aware of the impact of the minimum lot size restriction beyond gaining more land. In some cases, very limited increases in q_{min} can be welfare enhancing for the high-skill, high-income residents but substantial increases in the presence of even limited complementarity and productivity spillover effects can result in a net welfare loss for them. If q_{min} were endogenously chosen in the presence of perfect information, it is clear from the results that there is a welfare-maximizing level of $q_{min} > q^*_{L2}$, but that this utility-maximizing level of q_{min} is very sensitive to the values of δ and σ so would be sensitive to even minor underestimations of the parameters.

To see if this result bears out empirically, a test for negative productivity effects of housing restrictions identified by the model would ask whether cities with many areas of restricted housing access have lower productivity than cities with relatively freer housing markets. Another test derived from this model would be to see if cities that have small numbers of wealthy elites are more likely to have housing restrictions.

This paper represents an crucial first step towards improving the understanding of the intersection between housing policy, worker productivity, and neighborhood and peer effects. Additional changes can be made to clarify how these issues interlock, but even this relatively simple model demonstrates that there is a strong connection that warrants further study.

Figure 1.1: Relationship between high-skill utility and the minimum land consumption level (q_{min}) when neighborhood effects are relatively weak ($\sigma = 0.2 + \sigma_0$) and there is a more even skill distribution ($N_L/N_H = 4$).

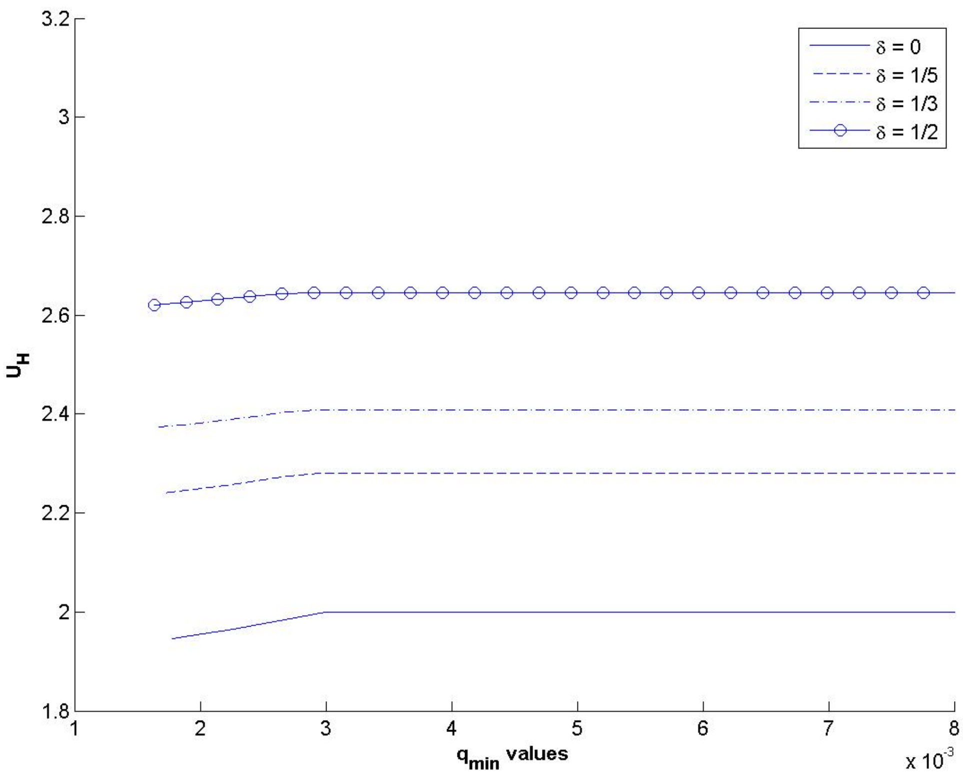

Figure 1.2: Relationship between high-skill utility and the minimum land consumption level (q_{min}) when neighborhood effects are relatively strong ($\sigma = 0.5 + \sigma_0$) and there is a more even skill distribution ($N_L/N_H = 4$).

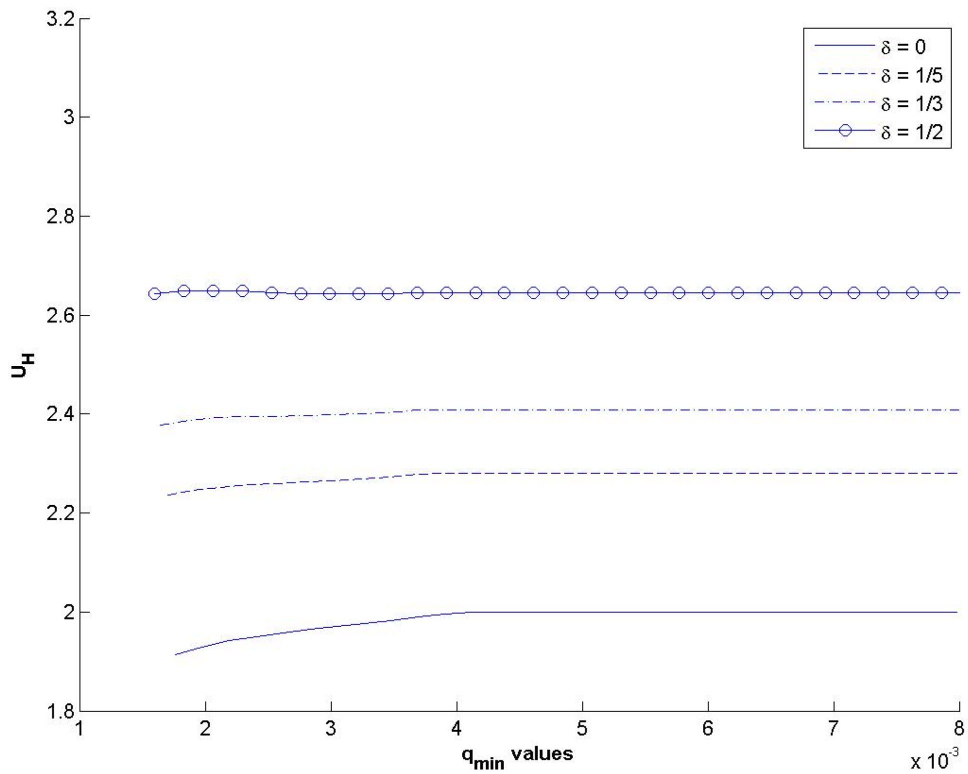

Figure 1.3: Relationship between high-skill utility and the minimum land consumption level (q_{min}) when neighborhood effects are relatively weak ($\sigma = 0.2 + \sigma_0$) and there is a less even skill distribution ($N_L/N_H = 6$).

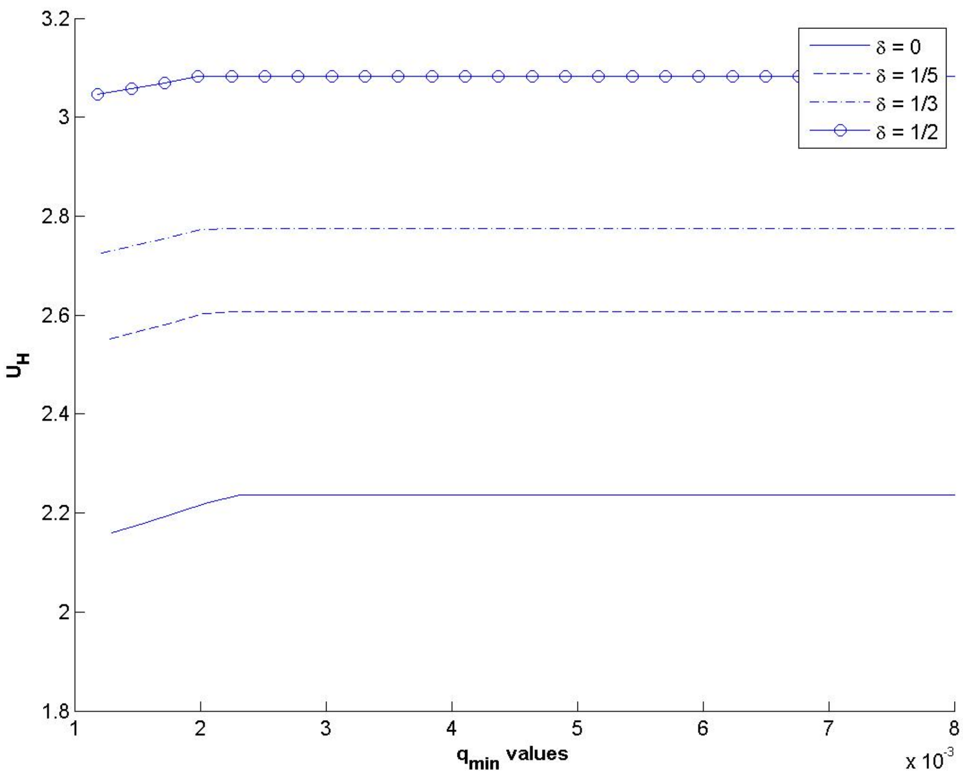

Figure 1.4: Relationship between high-skill utility and the minimum land consumption level (q_{min}) when neighborhood effects are relatively strong ($\sigma = 0.5 + \sigma_0$) and there is a less even skill distribution ($N_L/N_H = 6$).

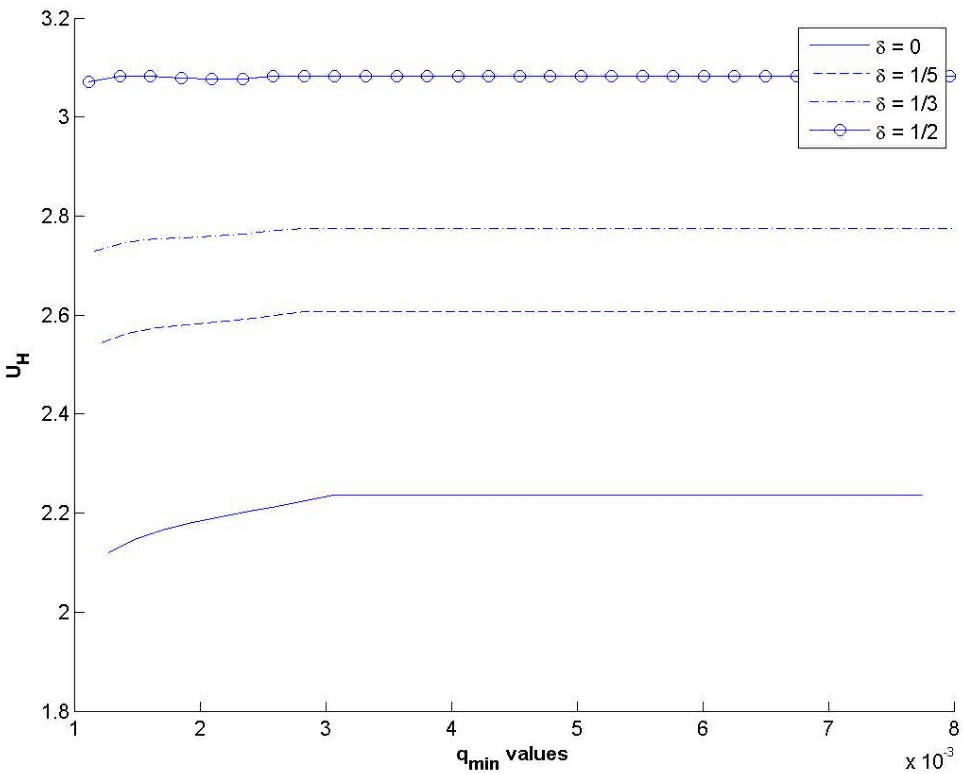

Figure 1.5: Relationship between the two consumer goods, land and the numeraire good, and the minimum land consumption level (q_{min}) when neighborhood effects are relatively weak ($\sigma = 0.2 + \sigma_0$) and there is a more even skill distribution ($N_L/N_H = 4$).

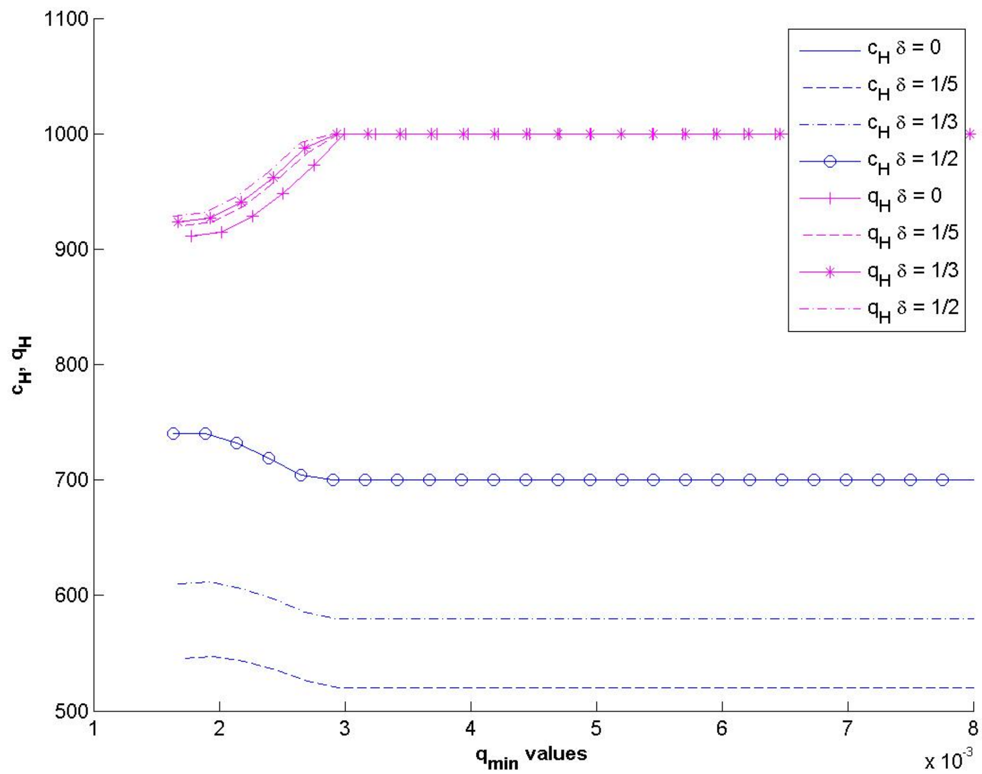

Figure 1.6: Relationship between the two consumer goods, land and the numeraire good, and the minimum land consumption level (q_{min}) when neighborhood effects are relatively strong ($\sigma = 0.5 + \sigma_0$) and there is a more even skill distribution ($N_L/N_H = 4$).

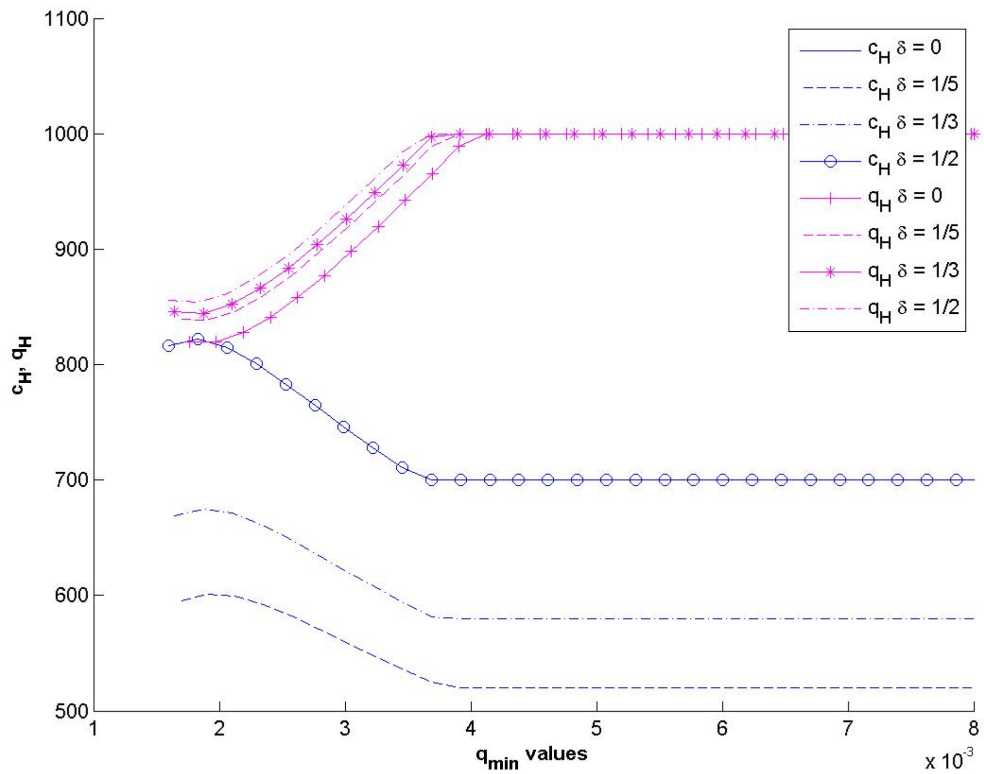

Figure 1.7: Relationship between the two consumer goods, land and the numeraire good, and the minimum land consumption level (q_{min}) when neighborhood effects are relatively weak ($\sigma = 0.2 + \sigma_0$) and there is a less even skill distribution ($N_L/N_H = 6$).

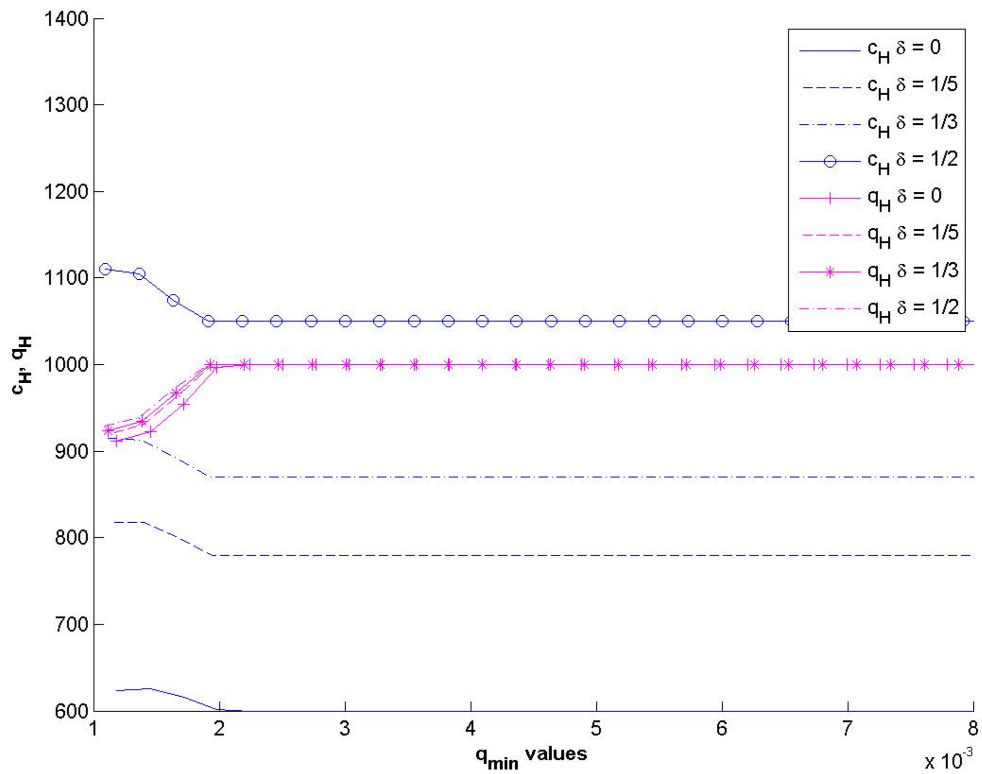

Figure 1.8: Relationship between the two consumer goods, land and the numeraire good, and the minimum land consumption level (q_{min}) when neighborhood effects are relatively strong ($\sigma = 0.5 + \sigma_0$) and there is a less even skill distribution ($N_L/N_H = 6$).

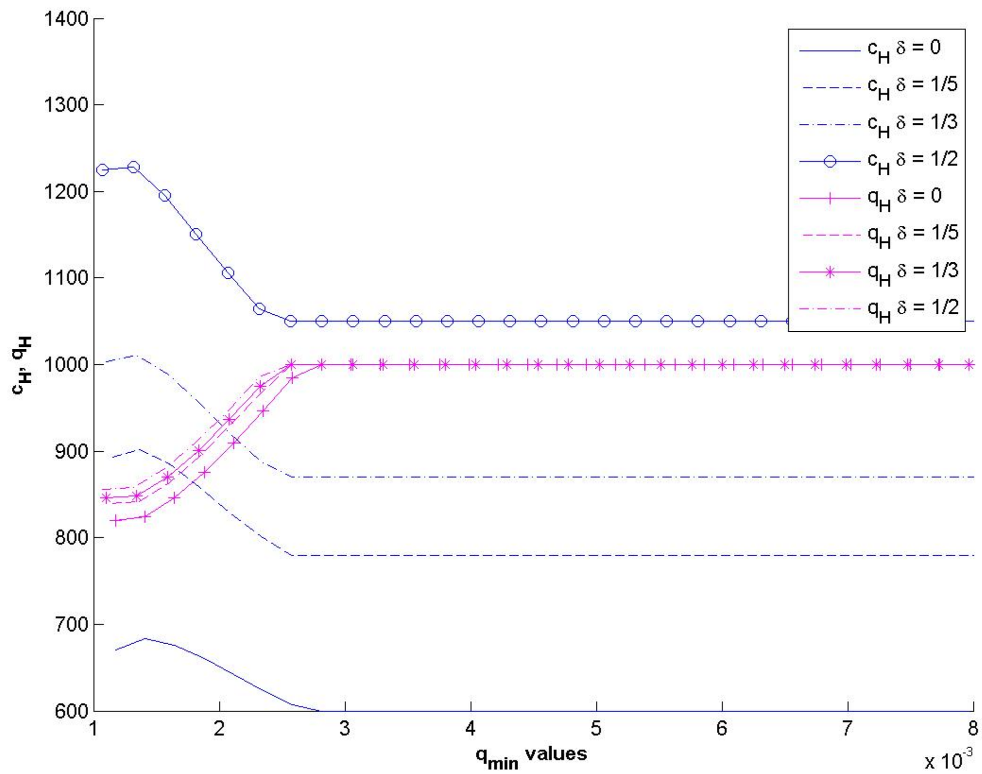

Figure 1.9: Relationship between high-skill utility and the minimum land consumption level (q_{min}) when wages are renormed to be the same across complementarity levels (δ); neighborhood effects are relatively weak ($\sigma = 0.2 + \sigma_0$); and there is a more even skill distribution ($N_L/N_H = 4$).

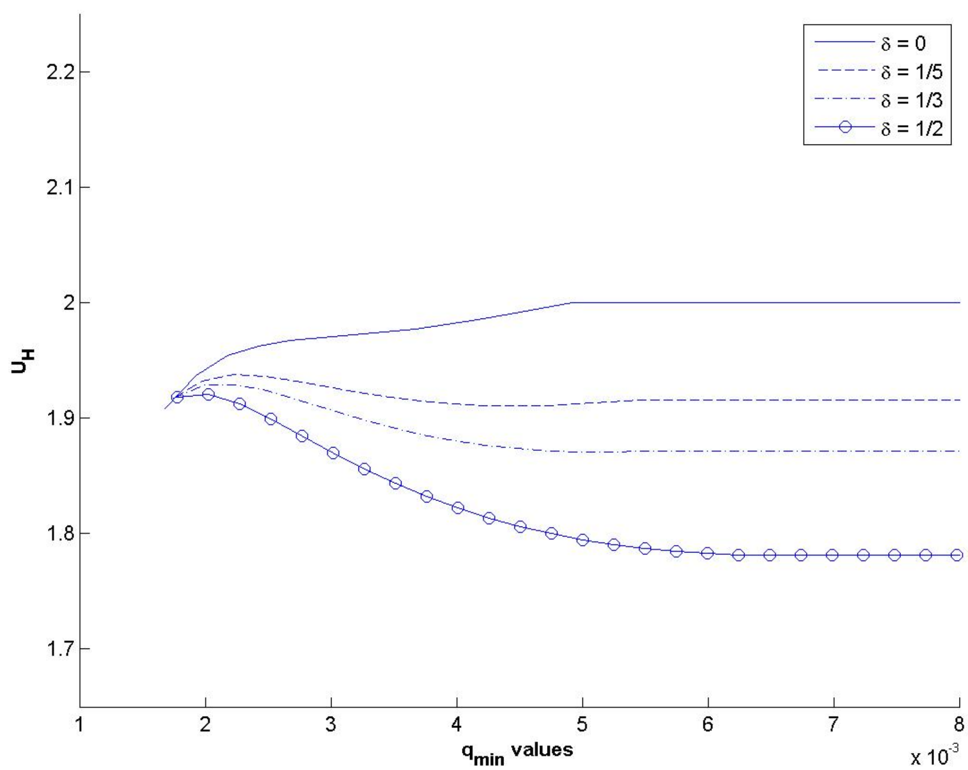

Figure 1.10: Relationship between high-skill utility and the minimum land consumption level (q_{min}) when wages are renormed to be the same across complementarity levels (δ); neighborhood effects are relatively strong ($\sigma = 0.5 + \sigma_0$); and there is a more even skill distribution ($N_L/N_H = 4$).

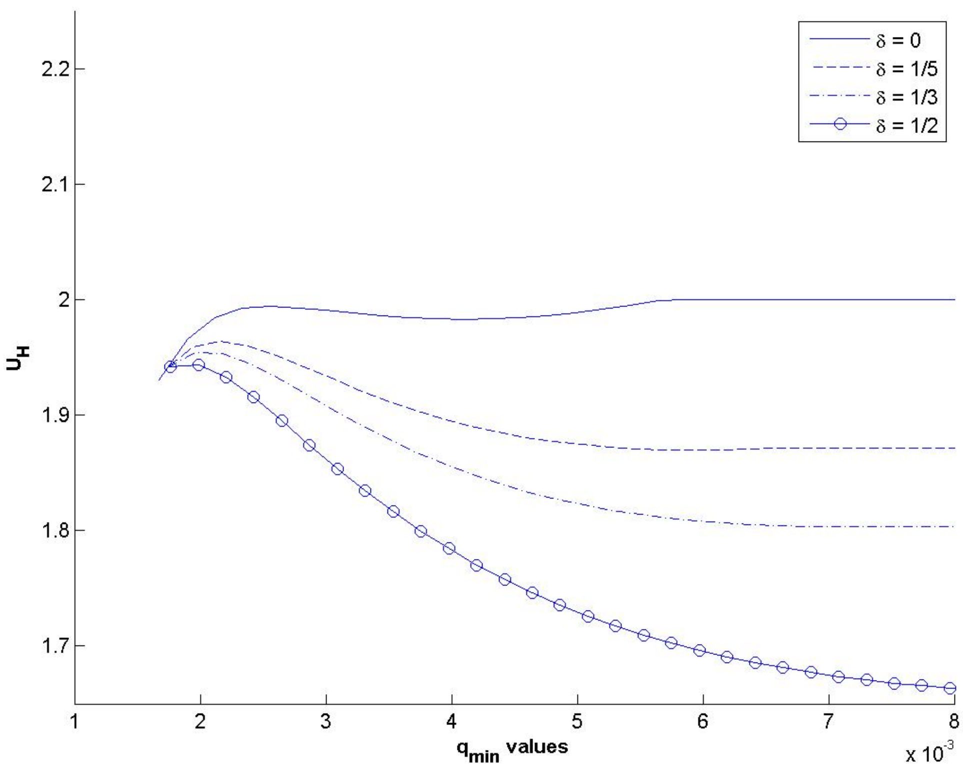

Figure 1.11: Relationship between high-skill utility and the minimum land consumption level (q_{min}) when wages are renormed to be the same across complementarity levels (δ); neighborhood effects are relatively weak ($\sigma = 0.2 + \sigma_0$); and there is a less even skill distribution ($N_L/N_H = 6$).

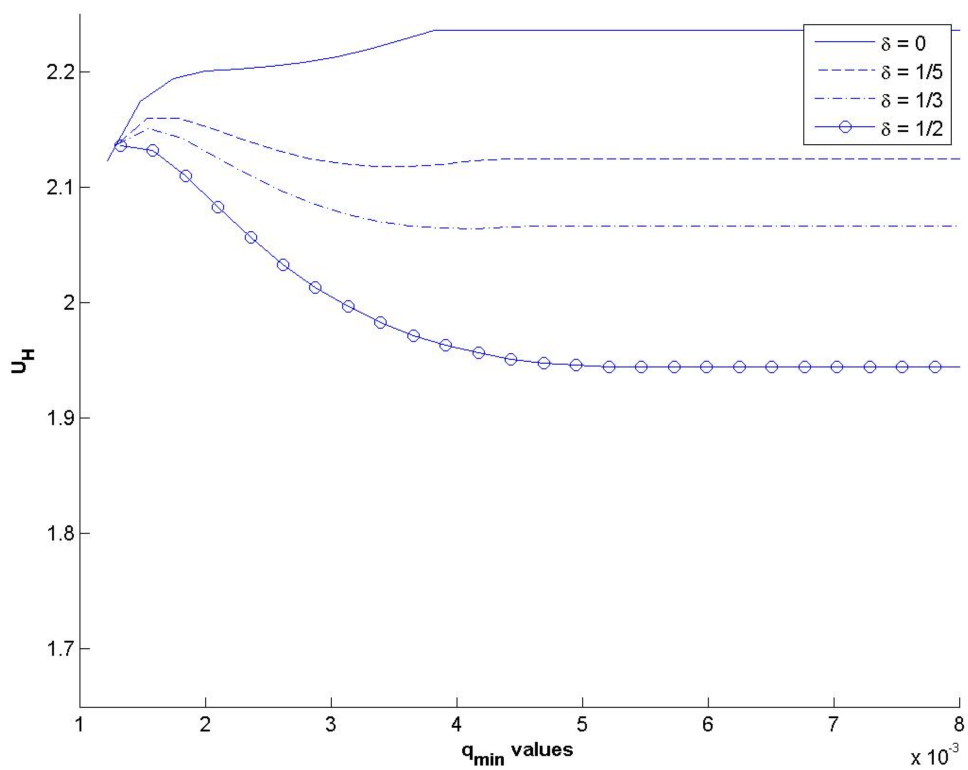

Figure 1.12: Relationship between high-skill utility and the minimum land consumption level (q_{min}) when wages are renormed to be the same across complementarity levels (δ); neighborhood effects are relatively strong ($\sigma = 0.5 + \sigma_0$); and there is a less even skill distribution ($N_L/N_H = 6$).

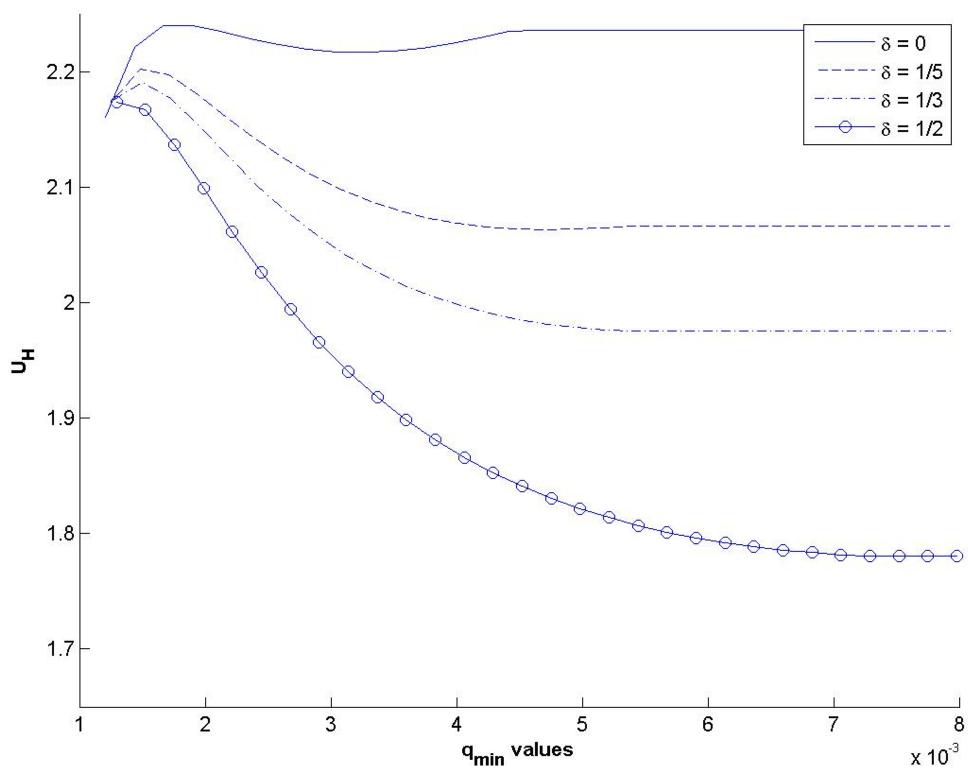

Figure 1.13: Relationship between total welfare and the minimum land consumption level (q_{min}) when wages are renormed to be the same across complementarity levels (δ); neighborhood effects are relatively weak ($\sigma = 0.2 + \sigma_0$); and there is a more even skill distribution ($N_L/N_H = 4$).

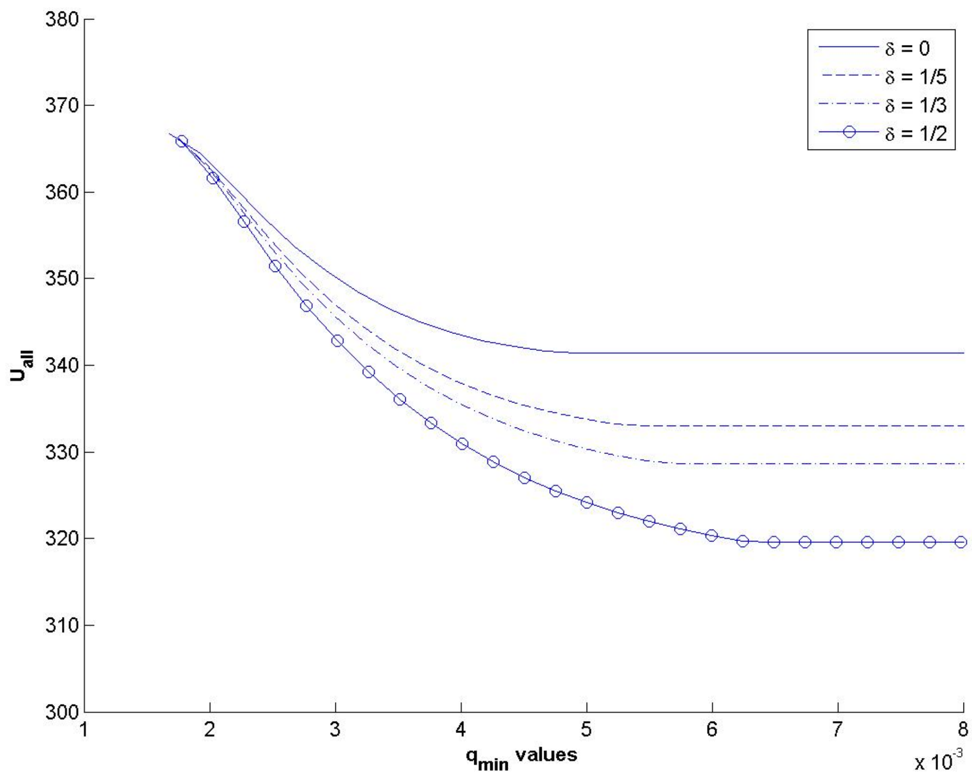

Figure 1.14: Relationship between total welfare and the minimum land consumption level (q_{min}) when wages are renormed to be the same across complementarity levels (δ); neighborhood effects are relatively strong ($\sigma = 0.5 + \sigma_0$); and there is a more even skill distribution ($N_L/N_H = 4$).

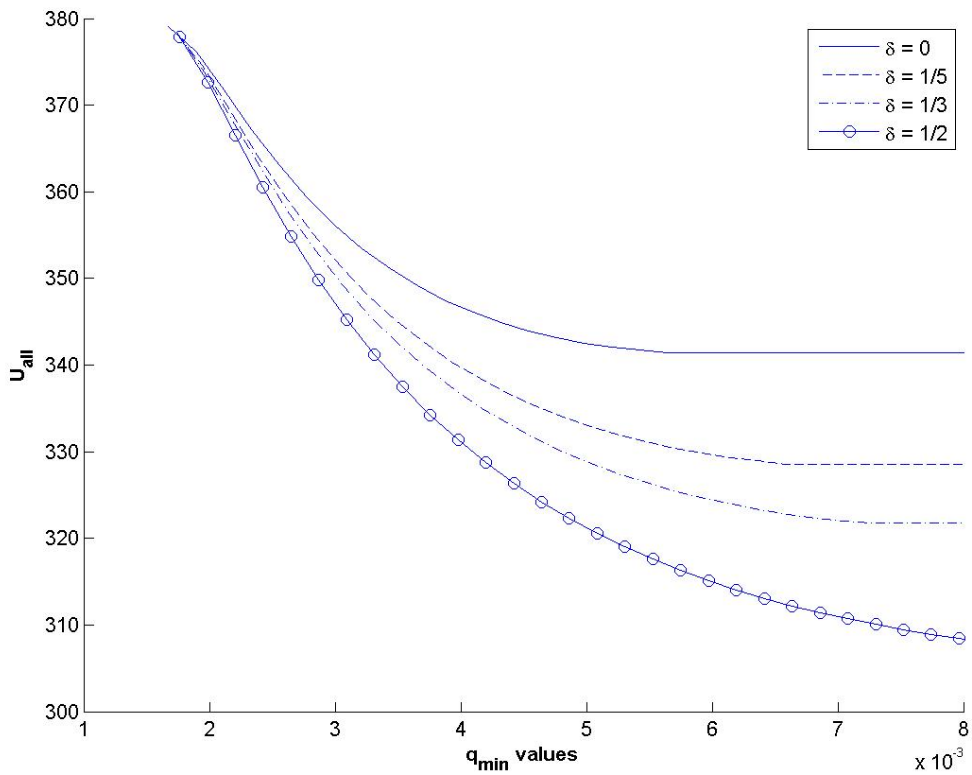

Figure 1.15: Relationship between total welfare and the minimum land consumption level (q_{min}) when wages are renormed to be the same across complementarity levels (δ); neighborhood effects are relatively weak ($\sigma = 0.2 + \sigma_0$); and there is a less even skill distribution ($N_L/N_H = 6$).

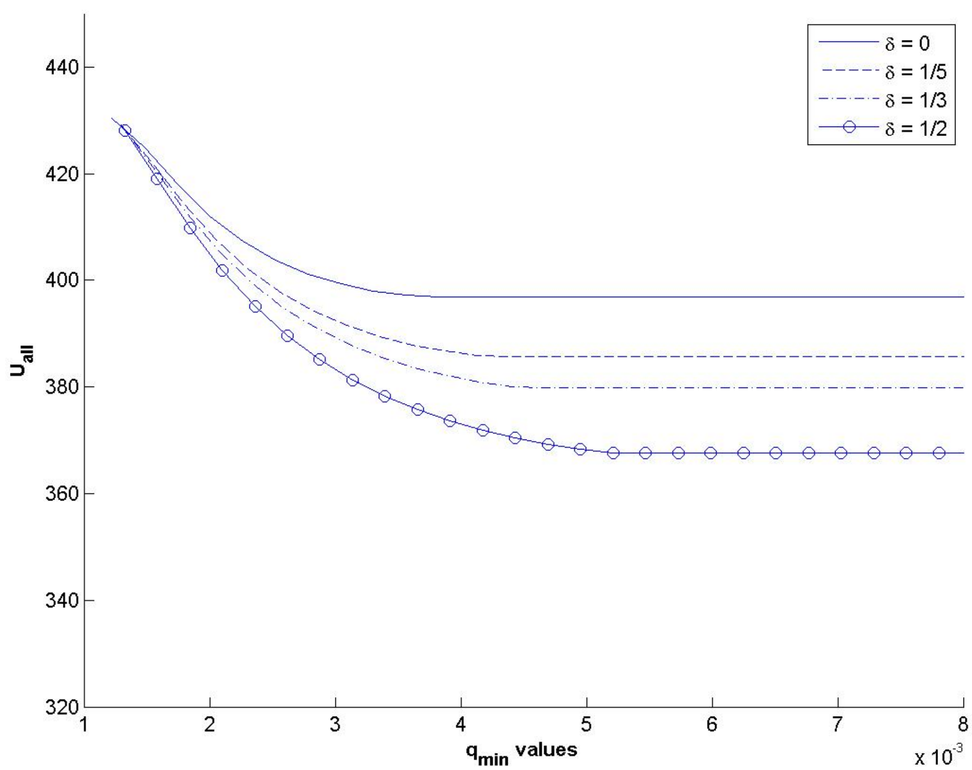

Figure 1.16: Relationship between total welfare and the minimum land consumption level (q_{min}) when wages are renormed to be the same across complementarity levels (δ); neighborhood effects are relatively strong ($\sigma = 0.5 + \sigma_0$); and there is a less even skill distribution ($N_L/N_H = 6$).

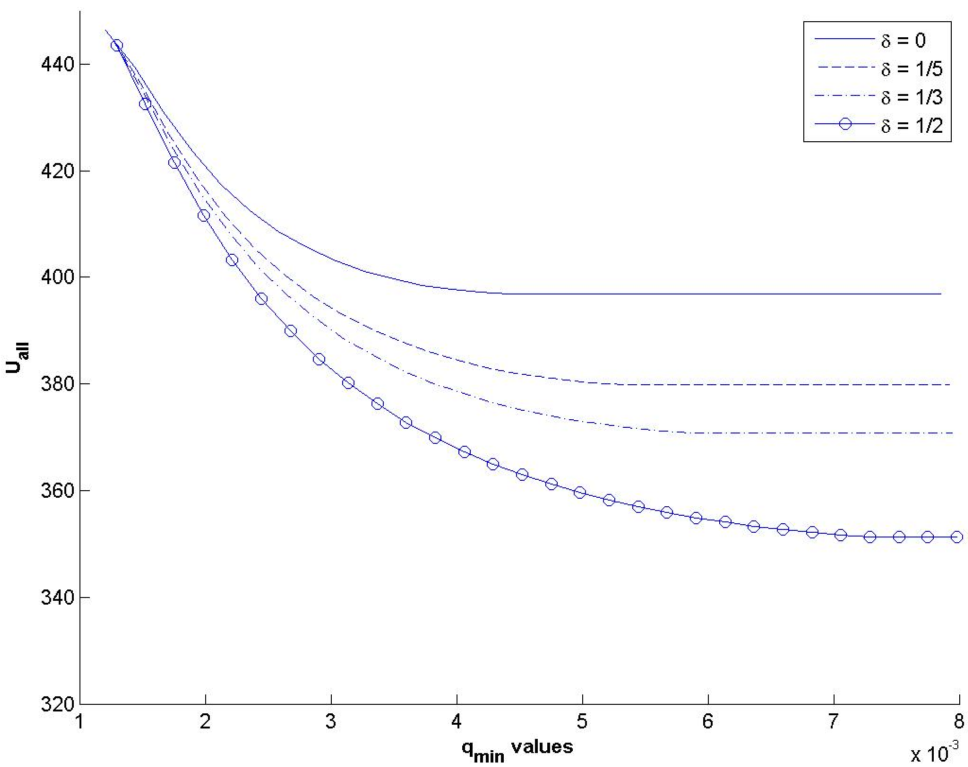

Chapter 2

Rent Control and Evictions: Evidence from San Francisco

Introduction

Forty years after New York City, Los Angeles, Washington D.C., Oakland, San Jose, and San Francisco adopted tenancy rent control, its ability to deliver the promised benefits to tenants is unconvincing. Intended to aid low- and middle-income renters,[1] these same metros are instead plagued by housing shortages and limited housing affordability.[2] This growing affordability crisis in rent controlled cities has been met largely by silence in the literature.[3]

[1]San Francisco's Rent Stabilization and Arbitration Ordinance, SF Administrative Code §37.1(b)(1), claims its purpose is to resolve "..a shortage of decent, safe and sanitary housing in the City and County of San Francisco resulting in a critically low vacancy factor"

[2]According to Census Bureau April-June 2016 figures, all of the major cities' metropolitan areas with rent control had rental vacancy rates at least a point below the national average of 6.7%. These same six cities occupy the top 7 places with the most expensive rents, with Boston (a former rent control city itself) being the seventh. Source: Zumper, Inc. `https://www.zumper.com/blog/2016/04/zumper-national-rent-report-april-2016/`, last accessed October 15, 2016.

[3]Public interest in rent control, on the other hand, has recently reemerged. Seattle and several Bay Area suburbs have seen tenant advocacy groups emerge in favor of rent control, while Alameda, Burlingame, Mountain View, Richmond, and San Mateo in California include controls on the November 2016 ballot.

The dearth of studies is all the more surprising because the controlled rental housing stock is clustered in some of the nation's wealthiest and most productive cities and is a substantial part of each: One-third of rental housing units in San Jose, 85% in Los Angeles, 2/3 in Oakland, 2/3 in Washington DC, roughly 47% in New York City, and 72% in San Francisco.[4]

Rent control impacts key questions of economic interest. This paper examines rent control's eviction restrictions, because evictions shape the demographic distribution of people and workers (Desmond [39]) and housing availability for workers migrating to productive regions (Bunten [28]; Hsieh and Morietti [62]). However, analyses finding that controls decrease mobility or cause misallocation of units to non-needy tenants assume that evictions due to landlord profit-seeking, hereafter called "economic evictions", cannot or do not occur. Systematic economic evictions would mean that rent control imposes many costs on landlords while delivering less than the promised benefit to tenants. Thus, if this widespread assumption is not the case, rent control's integrity as a tenant protection housing policy is called into question.

In this paper, two research questions regarding rent control and economic evictions are addressed. First, controlled landlords engage in profit-seeking behavior by selectively evicting tenants despite barriers to prevent this: strong tenant protections, greater legal scrutiny on eviction proceedings, limitations on grounds for evictions, and legal buyouts of the tenant's lease. Second, despite policy incentives for landlords to remain in the controlled market, some economic evictions occur when landlords respond to price signals by using evictions to switch to the uncontrolled sector. A finding that controlled landlords are willing to switch markets when rents rise is important for understanding rent control's long-run viability, because almost all cities have a ban on applying controls to new buildings.

Empirically testing for economic evictions means demonstrating that evictions rose in re-

[4] Section 2.1 has an overview of San Francisco's Rent Ordinance and has a more detailed policy discussion of rent control and eviction laws in other jurisdictions. See Table 2.5 for the sources for the coverage figures.

sponse to a known price shock. The price pressure must be well-identified and exogenous to avoid biased estimation from tenant self-selection into rent control on unobservable characteristics (Glaeser [49]; Early [42]; Ault, Jackson, and Saba [8]). This paper uses San Francisco data because it has a unique, spatially-varying, well-identified price shock: the network of commuter corporate shuttles stops operated by Google, Apple, Facebook, and Electronic Arts (EA), which transport employees from various sites around San Francisco to Silicon Valley and is a highly valued employee benefit (Dai and Weinzimmer [37]). The value of the shuttle's transit amenity is shown to be capitalized into housing prices, and substantial enough to plausibly incentivize the marginal landlord to evict. The paper then uses shuttle stop placements as a proxy for time- and area-varying free-market rent increases, and investigates whether evictions increased in buildings near a shuttle stop.

I find evidence that landlords engage in economic evictions. First, I find that the commuter shuttles raised the price per owner-occupied housing unit by $51,356, yielding a price increase of 10.5%. The first result is significant because it suggests that landlords stand to recoup at least their fixed eviction costs in the long-run after they fully switch into the uncontrolled market. The second result indicates that the price percent change in owner-occupied housing exceeded the controlled units' maximum allowable annual rent increase. Generalizing the percent increase across housing markets, the shuttles plausibly create a true economic incentive for landlords to pursue an eviction to capitalize in higher rents they cannot realize through other means. The shuttle system by its full extent in December 2013 was responsible for an additional 218 controlled buildings with at least one at-fault eviction per year. These evictions create greater tenant turnover in the controlled market than would otherwise exist, but leave the stock of controlled housing unchanged. However, that same shuttle-induced rent increase is estimated to cause 51 additional market withdrawals from the controlled market per year, likely permanently. Extending findings from other papers, the pattern of evictions indicates that the current policy structure is inimical to using rent control to create a more income-equitable allocation of housing in San Francisco, while also exacerbating housing market distortions.

Section 2.1 gives an overview of San Francisco's rent control and eviction laws and gives evidence for the study's external validity. Section 2.2 is a literature review that motivates the research questions by highlighting extant work on rent control and its relevance to economic evictions, while generating predictions of how landlords will use evictions in response to a price shock. Section 2.3 details the data used to test the hypothesis. Section 2.4 explains the empirical strategy. Section 2.5 gives results from the empirical strategy, and Section 2.6 is a discussion of the policy and economic considerations of the results. Section 2.7 concludes.

2.1 San Francisco's Rent Ordinance and Its External Validity

This section recaps San Francisco's, and other cities', rent control and evictions policies. Regulatory details are referenced throughout the rest of the paper to help tie the results back to the policy. The external validity of the study is also briefly discussed.

2.1.1 The Ordinance

There are four key provisions of San Francisco's Rent Ordinance:

1. **Rent increases are capped at 60% the rate of inflation.** Many maintenance costs can be passed through to the tenant. Limited hardship provisions ensure that landlords can earn a profit.

2. **Security of Tenancy.** Landlords cannot refuse to renew the lease of a tenant in

compliance. After the original lease expires, tenancy becomes month-to-month. Lastly, eviction notices must be for "just-cause" and approved by the San Francisco Rent Board.

3. **Vacancy decontrol-recontrol (or vacancy decontrol).** Tenants and landlords may negotiate the base rent, and only subsequent increases are controlled (capped). After a tenant vacates, the landlord can negotiate a new base rent with the next tenant and controls only apply to subsequent rent increases.

4. **No new controlled buildings.** Only buildings built before June 13, 1979 and with 2 or more units are subject to rent control.

As stated in Provision (2), rent controlled landlords must have a "just cause" for an eviction. The 15 grounds for a "just-cause" eviction are given in Table 2.1. Seven for an at-fault tenant, who is in some way in breach of the lease,[5] and eight where the tenant is not at fault, or "no-fault". Six of the no fault evictions are not part of this study because they are either temporary, very rarely permitted by the city, or are suspended from use.[6] Landlords most commonly use two no-fault eviction types: owner move-in (OMI) evictions, where the landlord wants to reclaim the unit for themselves or for a close relative, and Ellis Act evictions, where the landlord retires the units from the rental market.

OMI and Ellis Act evictions are advantageous to landlords because tenants can be evicted even when fully lease compliant and, unlike demolitions or rehabilitations, there are minimal bureaucratic delays on eviction notice issuance. Critically, they also create opportunities

[5]More information is in Table 2.1, but the seven grounds are 1.) on-payment or late payment on rent, 2.) general breach of the lease, 3.) nuisance, 4.) illegal usage of unit, 5.) refusing to quit after previously agreeing to end tenancy, 6.) refusing to grant landlord lawful access, and 7.) sole remaining tenant is an unapproved subtenant.

[6]These include temporary eviction for lead abatement or capital improvements; revoking of "Good Samaritan" status for tenants who are fleeing natural disasters; and converting rental units to condominiums, which was suspended in 2013 after being allotted very rarely by lottery in the previous decade. Demolitions are more commonly granted, but only after an extensive permitting process and are thus strongly endogenous to local conditions.

to exit the controlled market. Once the tenants are evicted, it is easier to gain a permit to convert the units to condominiums or demolish the building entirely which has the additional advantage of the new building exemption (Provision (4)) meaning that new rental units will be uncontrolled regardless of the building's future uses.

The city compensates for the lack of red tape for OMI and Ellis Act evictions by making them very costly. Table 2.3 shows how San Francisco passed various policy changes between 2002 and 2013 to regulate controlled evictions. Since 2006, all tenants are entitled to relocation payments, with landlord surcharges for protected tenant classes, such as elderly, disabled, or legal minor tenants. Table 2.4 shows how these relocation payments grew with time so that by December 2013, a landlord had to pay roughly \$5,200 for each evicted tenant, capped at about \$15,620, with a protected surcharge of about \$3,470. Other rules include suspending vacancy decontrol on withdrawn units for up to 3 years after an OMI eviction and 10 years after an Ellis Act eviction if the landlords rerent the units. Landlords can only do one OMI eviction per building and the set-aside unit is marked on the deed. A post-Ellis Act vacant building faces additional restrictions. A ten year period is marked on the deed where the new building exemption is suspended for the property. If the landlord demolishes the old units and build new ones during this time, rent control will apply until the waiting period expires.

Non-controlled landlords can evict tenants without cause, pursuant only to the lease and relevant state and city statutes. However, San Francisco is clear that controlled unit evictions should only happen "in good faith" (San Francisco Administrative Code §37.9(8)),[7] but fluctuations in eviction counts follow changing economic conditions. Table 2.2 shows the yearly number of controlled unit evictions for all at-fault evictions, OMI's, Ellis Act evictions, demolitions, and the annual change in the Bureau of Labor Statistics Consumer Price Index (CPI) for the rent of the primary residence for the San Francisco Bay Area. This table also

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

shows why a study that just used a city-level time series would not properly identify a causal effect on evictions from rent increases. Changes in at-fault evictions track changes in the rents, but it is impossible to know at this level of variation whether tenants, landlords, or both are changing their behavior as rental growth rates change.

Evictions can be avoided altogether through a buyout agreement. Court cases for at-fault evictions are costly, time-consuming, and uncertain, so that landlords prefer to buy out a tenant first (Downs [41]). Legal buyouts do not mean that the economic eviction rate is zero. Evictions are a landlord's credible threat to achieve a higher rate of successful buyouts, but then can be used as a final resort on recalcitrant tenants. Legal buyouts' significance for this study is only to artificially lower the observed economic eviction rate.

Buyout agreements could more seriously bias the results if there are reasons to believe that tenants are more (or less) likely to sign them when rents rise. Unfortunately, San Francisco only started regulating and publishing detailed information on buyouts in 2015.[8] Further, the literature is silent on how to sign this bias. From first principles, rising rents could make tenants more fearful of eviction because finding comparable units is now more expensive. They would then be more inclined to sign a buyout, biasing downwards the sign on any exogenous local shock to rents. Thus, the results are likely to understate both the impact of rent control and the impact of a rent shock on controlled landlords.

In summary, rent control works by locking landlords and tenants into a base rent that erodes annually in real terms. Tenants can only be evicted for violating the lease or under limited, expensive circumstances when lease-compliant. The only path to full decontrol is to demolish the current building and rebuild it, pending approval for a city-granted demolition permit.

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

2.1.2 External Validity

San Francisco's rent control provisions are not unique. Los Angeles, Oakland, San Jose, New York, and Washington D.C. have very similar rules. Table 2.5 lists the rent control and eviction rules in these cities analogous to the one in San Francisco listed in 2.1.1. All five have at least limited vacancy decontrol coupled with yearly rent increase restrictions, new building exemptions, and allow tenants to indefinitely renew their leases. Only San Jose does not require a just cause for an eviction, but even there, landlords have to go through a city-mandated arbitration process to evict.

Thus, if economic evictions are found in San Francisco, they are found elsewhere. The most significant difference is that the economic eviction rate is probably elevated in San Francisco. The city is the only major locale where inflation outpaces controlled rent increases by design.[9] Economic evictions are therefore more strongly incentivized in San Francisco because landlords face steeper losses with the passage of time.

2.2 Literature and Theoretical Review

Three questions motivate this paper: why are controlled landlords incentivized to do an economic eviction? Do landlords perform more evictions when prices rise? Do rising prices incentivize no-fault evictions (leaving the controlled market) or at-fault evictions (staying in the controlled market)? The last question in particular is critical for understanding the long-term viability of rent control. Per Section 2.1.1, no-fault evictions allow landlords to switch their properties to the uncontrolled sector. If no-fault evictions increase with rents, rising prices will accelerate the demise of the controlled housing stock. This literature

[9] This claim does include those remaining units in New York City under first-generation rent control.

review explores answers to these questions from other papers and highlights this paper's contributions.

2.2.1 Existing Literature on Rent Control and Evictions

The only extant analysis of rent control and evictions is by Miron [84]. His model assumes that rents are capped and economic evictions are impossible, so that landlords' constrained eviction ability incurs higher costs due to adversely selected tenants.[10] Miron considers an adversely selected tenant to be a legally-protected tenant who abuses the tenancy security laws by being a nuisance or frequently delinquent on rent. However, under rent control, adverse selection is not just about tenant quality; it is also about tenants' expected duration.

Consider San Francisco's annual rent increase cap, which is equal to 60% the rate of inflation (Provision (1)). Landlords' losses increase every period market rents rise faster than the allowed increase.[11] If a tenant can be induced to leave, the landlord can lock in a higher base rent with a new tenant thanks to vacancy decontrol (Provision (3)). Controlled landlords in a rising market thus want "short-stayers" so they can frequently reset base rents, and avoid "long-stayers" who compound losses with each subsequent rental period. Coupled with the other kinds of adverse selection mentioned above, rent control landlords run a substantial risk of acquiring an expensive, undesirable tenant they cannot be rid of.

Rent control distorts the housing market because controlled landlords prefer short-stayers but cannot directly observe the tenant's duration preference. Long-staying tenants, knowing this, are incentivized to present themselves as short-staying tenants. Landlords must then maximize their profits in the face of uncertainty engendered by imperfect information on the

[10] Miron [84] and Hubert [63] use "economic eviction" to mean pricing a troublesome tenant out of their unit. This is not possible under tenancy rent control by design, so I believe it is more useful to discuss it in this context as an eviction driven exclusively by profit-seeking behavior on the part of the landlord.

[11] Even in a static market, landlords will suffer losses due to erosion in real rents.

tenant's true type (Basu and Emerson [16], [17]); Arnott and Igarashi [6]; Hubert [63]; Iwata [66]), raising costs for everyone.

Miron concludes that landlords will turn to one of two recourses when they cannot evict: charging significantly higher first-year rents to compensate for adverse selection's expected costs or withdrawing from the rental market altogether. However, Miron's paper, like all others in the literature, considers only the extreme case that evictions are entirely prohibited. Section 2.1.1 and Table 2.1 are clear that evictions are not forbidden, only limited to just causes in most rent control cities. Nonetheless, knowing the outcome at the "corner solution" of $P(Eviction) = 0$ is useful for considering cases where $P(Eviction) > 0$, which are now considered below.

2.2.2 Predicting How Evictions Respond to Price Shocks

The above section establishes why prohibited evictions are costly for landlords, but does not directly answer how rent control incentivizes evictions. The two sections below draw on the literature to hypothesize how landlords will utilize evictions when market rents change.

2.2.2.1 "3 Days to Make Whole Or Quit": At-Fault Evictions and Price Shocks

A reasonable hypothesis from Section 2.2.1 is that controlled units have a higher baseline at-fault eviction rate than uncontrolled units. "Bad" tenants are disproportionately drawn to controlled units because of tenant security laws (Provisions (1) and (2)), and rent control puts financial pressure on landlords to turn over tenants (Provisions (1) and (3)). Any study on evictions needs to account for differential eviction rates caused by selection into rent control on unobservable tenant characteristics.

There are some financial and bureaucratic barriers to tactical at-fault evictions in the controlled sector,[12] but these regulations likely exacerbate the moral hazard and adverse selection problems. Self-aware "bad" tenants can outbid "good" tenants for these protections, and those in controlled apartments have fewer incentives to remain good-quality tenants. *Ex ante*, these barriers are probably not prohibitive enough to completely deter economic evictions and likely exacerbate the underlying adverse selection problem. Thus, a reasonable prediction is that price shocks raise an already elevated eviction rate in controlled units.

2.2.2.2 "Everybody Out!": No-Fault Evictions and Price Shocks

Drawing reasonable hypotheses on no-fault evictions is harder, but even more important because they impact the long-run housing supply. High upfront costs mean that the baseline no-fault eviction rate is almost surely lower in controlled units. A price shock to uncontrolled units likely lowers their no-fault eviction rate, because simple supply and demand suggests that landlords are unlikely to withdraw or repurpose units when rents rise.

On the other hand, no-fault evictions for controlled units might rise with rents because these evictions clearly have a "pull" component. If landlords believe that rents will persistently increase or a price shock is permanent, they'll anticipate that greater long-term profits in the uncontrolled sector will cancel out the short-to-medium term losses from the eviction. Further, the literature suggests that controlled landlords have a "push" component to use no-fault evictions, *particularly* when prices are rising.

That last prediction is drawn from two papers on tenancy rent control by Basu and Emerson [16], [17]. They point out that landlords' adverse selection problem worsens when rents

[12] Landlords must file in writing with the Rent Board the notice to quit, justify the grounds given, and give an opportunity to make whole the breach in the lease. An exception exists if the unit is being used for illegal purposes. The notice to quit in that case is unconditional. They must also inform the tenant of their rights under the Rent Ordinance in writing. If the eviction is ruled to lack a just cause, the landlord has to pay at least three times the actual damages, including mental and emotional distress (San Francisco Rent Ordinance Administrative Code § 37.9(f))

increase because only tenants preferring long durations are willing to pay the controlled unit premium. Worse, rising rents also endogenously increase chosen long durations (Basu and Emerson [16]). Several papers have empirically and theoretically demonstrated that tenants do pay a premium for rent control and security of tenancy, including Raess and von Ungern-Sternberg [99], Skelley [108], Arnott ([5]), Turner and Malpezzi [114], and Arnott and Igarashi [6].[13] Nagy [91] found using New York City data that landlords collected the same average rent per tenancy across controlled and uncontrolled units, which could only be true if tenants were willing to pay a higher base rent at the outset.

How serious of a problem are long-stayers for controlled landlords? If landlords conclude they are very likely to be saddled with a costly long-stayer, Miron [84] and Basu and Emerson [17] explain they will choose market exit over finding a new tenant. The crux of the problem is that unlike rent delinquents and nuisances, controlled landlords lack any direct remedy for long-stayers. Tenants have indeed been repeatedly found to disproportionately have long tenures in all forms of rent control (*e.g.* Linneman [81]; Gyourko and Linneman [56]; Munch and Svarer [87]; Nagy [90], [91]; Basu and Emerson [16]). Long durations are incentivized not just by increasingly steep rent discounts when the market as a whole is rising, but also by security of tenure itself (Basu [15], Iwata [66]). Unfortunately for landlords, not only are long-stayers disproportionately attracted to rent control, but Ault, Jackson, and Saba [8] found that up to 80% of the extra tenancy duration is due to tenants endogenously choosing to extend their tenancies. Tenants have been documented going to some lengths to retain their controlled units, such as by increasing their propensity to accept a local job (Svarer, Rosholm, and Munch [111]) or accepting longer commutes (Krol and Svorny [77]) when switching jobs. All these problems will be exacerbated as rents rise, and in fact, might also spur further declines in the controlled housing stock.

Thus, the weight of the evidence in the literature is that some controlled landlords will choose

[13] The findings in these models rest on the assumption of absolute security of tenancy, but there is no reason to believe that even partial security of tenancy would not be valued by tenants.

to decrease their medium-run housing supply when rents rise. The *ex ante* hypothesis is that both types of evictions will thus rise with rents.

2.2.2.3 Eviction Costs

The final consideration is how much will rents need to rise to plausibly incentivize the marginal landlord to evict. There is very little evidence on average eviction costs, but costs benchmarks can be established from the laws.

Two benchmarks for whether prices rise enough to incentivize evictions are as follows:

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

The first ensures that controlled landlords cannot realize any gains in market rents and are incentivized to push out tenants. Between 2003 and 2013, the highest allowable annual increase was 2.1%.[14] The second condition ensures that landlords could plausibly profit from a no-fault eviction. Exceeding at least the fixed cost of relocation payments is the most plausible minimum condition to push a landlord into a no-fault eviction. The empirical exercise shows that at a minimum, the hedonic value of a shuttle stop exceeds the most expensive relocation payment per unit of $26,055,[15] and increases the value of a unit by more than 2.1%.

[14] "Allowable Annual Rent Increases", The Residential Rent Stabilization and Arbitration Board of the City and County of San Francisco, 2016.

[15] From the 2013 amounts given in Table 2.4. Assumes three people in the unit to hit the maximum cap, and all three of whom are in a protected class.

2.3 Data

The outcome of interest in this paper are evictions, and San Francisco has several publicly-available sources for issued eviction notices. Data for controlled, "just-cause" evictions come courtesy of the Rent Board of the City of San Francisco. The Rent Board publishes information on all just-cause evictions since January 1st, 1997 and are available at public request. The data set includes the address of the eviction, the date the eviction was served, and the reason for the eviction. The provided Rent Board dataset totaled 30,992 evictions. Errors in the recorded addresses were corrected by consulting Google Maps[16] and a dataset of addresses from the City of San Francisco's Planning Department. All merges involving addresses therefore occur at the building or parcel level and not for the specific housing unit. For less than 0.5% of the evictions, the address could not be determined from the record, and these observations are excluded. Furthermore, an additional 9% of addresses did not match with any known address in the city's database. After consulting with the San Francisco Rent Board, they believe these are mostly illegal units, and therefore unmatchable to the official address list.

Data on issued eviction notices in uncontrolled units is also required. Plaintiffs filing an unlawful detainer action will include in their court documents scanned copies of their eviction notices and other relevant information. These scanned documents were manually examined and appended with the controlled evictions dataset to create a master dataset of all sector eviction notices.

Evictions then need to be matched to characteristics of the building the tenant was issued a notice in. Data for building characteristics and housing values come from the City and County of San Francisco's Office of the Assessor-Recorder. Supplemental information on addresses, latitude and longitude, and other property parcel characteristics was obtained

[16] `maps.google.com`

from the website of the City and County of San Francisco.[17] The dataset covers assessments made yearly from June 2003 to June 2014, for 2,373,721 observations, ranging from 188,333 parcels in 2003 to 205,229 parcels in 2014. Some records were self-evidently corrupted, and were cleaned through manual inspection and by comparing other records for the same parcel. All identified public housing complexes are dropped.

The property survey is only conducted annually, and does not report specific demolition, construction, or building opening dates. Property change dates were thus obtained from the City of San Francisco's Department of Building Inspections.[18] Data on the certificates of final occupancy from January 2001-December 2013 were purchased from the same department.

Gaps and inconsistencies in the property survey were remedied by drawing on building characteristics attached to the permits and building certificates. Discrepancies in reported sales prices were compared with publicly-reported information from the real estate website Zillow.[19]

Table 2.6 presents information about the housing market in San Francisco. San Francisco's housing stock is disproportionately older, and apartment buildings comprise just over half of the stock. The majority of housing units are also rent-controlled, and the decade between 2003 and 2013 saw a decrease in the number of rent-controlled housing units, apartment units, and residences. As the housing stock modernized, rental units were replaced with owner-occupied single-family housing. The other notable feature is that on average, the nearest shuttle stop is less than a half-mile away, meaning that condos did in fact experience significant "treatment" from the placement of shuttle stops.

Any parcel or eviction that is in Golden Gate Park, the Presidio, or Yerba Buena and

[17] `data.sfgov.org`

[18] Demolition permits can be found at `http://sfdbi.org/demolition-permits-filed-and-issued`. Building permits can be found at `http://sfdbi.org/building-permits-filed-and-issued`. For 2001-2003, these permits are freely available upon request.

[19] `www.zillow.com`

Treasure Islands is dropped. The former two are dropped because they are park properties and the latter because they are small, non-contiguous parts of the city. This leaves 85 neighborhoods in the evictions analysis covering 126 months for 10,710 observations. To test for economic evictions at a more disaggregated level, another panel was constructed from the same source of rental buildings with two units or more comprising 37,000 buildings across 126 months.

2.4 Empirical Approach

This paper asks whether rent-controlled landlords engage in economic evictions, either to replace the existing tenant with a new tenant who pays a higher base rent or as a means of exiting the controlled market entirely. The ideal empirical strategy would be to estimate an equation of the form:

$$\begin{aligned} Eviction_{it} = &\zeta_0 + \zeta_1 Rent_{it}^{RC} + \zeta_2 RentControl_{it} + \zeta_3 \left(Rent_{it}^{RC} * RentControl_{it} \right) \\ &+ \zeta_4 Time_t + \zeta_5 F_i + \zeta_6 OtherControls_{it} + \epsilon_{it}, \end{aligned} \tag{2.1}$$

where $Rent_{it}^{RC}$ is the prevailing rent for a vacant controlled unit in building or area i during time period t; $Time_t$ is a fixed effect for time period t; F_i is a fixed effect for the building or area; and $OtherControls_{it}$ are other characteristics relating to rent control and evictions.

However, including $Rent_{it}^{RC}$ creates prohibitive challenges. Even if an exhaustive dataset of rents existed, rents are endogenous to evictions by this paper's premise. This makes consistent estimation of Equation (2.1) impossible. Since rents are both systematically unobservable and endogenous, this paper instead proxies for changes in the market rent level using changes in local transit amenities created by the local roll-out of Google, Apple, Facebook, and Electronic Arts' commuter shuttle stop program.

Figure 2.1 shows how the shuttle stops spread throughout San Francisco, starting from September 2004, when Google placed the first two stops, to the study period's end in December 2013. The shuttles are an important employee benefit, because the distance from San Francisco to worksites in Silicon Valley can be as long as 50 miles. The shuttles have wireless internet, and many employees can now use their practically door-to-door, free commutes to get work done (Dai and Weinzimmer [37]). Proximity to a shuttle stop is thus a transit amenity a technology company employee might value highly. Although the shuttles are a private good, demand for access should be strong enough to impact neighborhood housing markets.

A more extensive description of the history and data collection behind the shuttles can be found in Appendix A.1, but a brief overview illustrates why they are a useful strategy to proxy identifiable local rent increases. The shuttles were first initiated by Google in 2004 and went to just two locations: Glen Park BART Station and a park and ride stop near Candlestick Park. Apple started its own system in 2007, Facebook in 2009, and Electronic Arts in 2012.[20] By 2009, the shuttles had come to cover many city neighborhoods either completely or in part, particularly in the eastern half of the city. The shuttle stops clearly favored the city's east and northeast, and so like $Rent_{it}^{RC}$, the shuttle stops are likely endogenous to local conditions. Therefore, the paper instruments for shuttle stop locations by exploiting exogenous constraints on their placements. The commuter shuttle stops can only be located at large public stops, called "bus zones" hereafter, because the shuttles are otherwise too large for San Francisco's streets. Patterns for shuttle stop placement are detailed below in Section 2.4.1.2.

The empirical investigation is thus conducted in two stages. The first stage establishes that the transit amenity from the commuter shuttles is capitalized into local housing prices

[20]Yahoo initiated its own service in 2005, followed by Genentech in 2006 and others after the recession (Dai and Weinzimmer [37]). Genentech and Yahoo! shuttle stops are not included in this study, but existing information indicates that they overlapped with Google, Facebook, and Apple stops almost completely.

using condominium sales data rather than rental units. Following ordinary least squares estimation, an instrumental variables (IV) approach instruments for shuttle stop placement. The value of the amenity is found to be plausibly large enough to incentivize a landlord to evict because it exceeds the values given in Conditions 1 and 2 in Section 2.2.2.3 as being necessary to plausibly incentivize landlords to evict. The second stage of the investigation tests for economic evictions in controlled units by exploring how eviction rates, counts, and probabilities change in response to greater commuter shuttle coverage and thus higher prices. Estimates on changes in evictions after exposure both to the observed and fitted shuttle placements are then presented.

2.4.1 Hedonic Price Effect of the Commuter Shuttles

In the absence of rent data, condominium sales in San Francisco from July 2003 to December 2013 are used to estimate the effect of the shuttle stops on housing prices. A series of hedonic price regressions show that the transit amenity from the privately-provided commuter shuttles is capitalized into local housing values and that this transit amenity is substantial. The key identification assumption in the hedonic price regressions is that there are no other unobserved shocks to the outcomes coincident with shuttle placement that affected pricing outcomes. This assumption will later be dropped when shuttle stops are instrumented.

This exercise establishes that the new transit amenity very likely put upward pressure on prices exogenous to other changes in the condominium market. The inference about concomitant rent pressure is made even though changes in condominium prices cannot be directly used to determine the equivalent rent hike. Condominium and apartment unit supply are endogenous (Sinai [107]),[21] even when there is rent control (Häckner and Nyberg [58]). Condominium sales prices also provide a useful way of appraising the net present value of the

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

transit amenity. They are germane to the paper's hypothesis because landlords will consider this information when deciding whether to use a no-fault eviction to exit the rent control market. Thus, while this analysis cannot determine exactly how the shuttles change the controlled rent premium, quantifying it from condominium sales is sufficient to establish that the premium exists in apartment rents.

2.4.1.1 Estimating the Commuter Shuttles' Transit Amenity Value from Condominium Sale Prices

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

$$\begin{aligned} Price_{it} =& \alpha + \beta_1 \mathbb{1}\{Shuttle_{it}\} + \beta_2 New_{it} + \beta_3 Baths_{it} + \beta_4 SqFt_{it} + \beta_5 YearBuilt_{it} \\ &+ \beta_6 Beds_{it} + \beta_7 OtherTransit_{it} + \beta_8 Nbrhood_i + \beta_9 Y_t + \beta_{10} M_t + \epsilon_{it}, \end{aligned} \tag{2.2}$$

where $\mathbb{1}\{Shuttle_{it}\}$ is an indicator for whether condominium i is a half-mile from the nearest Google, Apple, Facebook, or Electronic Arts shuttle stop. This is the treatment variable of interest. Control variables include $OtherTransit_i$, a collection of measures of how far the condo is from other forms of transit;[22] $Baths_{it}$, the number of baths in the condo; $SqFt_{it}$, the total area of the condo in square feet; $Beds_{it}$, the number of bedrooms in the condo; $YearBuilt_{it}$, the year the condominium was built; and New_{it}, whether the sale occurred

[22] This includes information on how far the average condominium is from alternative forms of transit, including distances to BART and Caltrain, to the light rail transit stops, and to major north-south thoroughfares. Specifically, these are the distances in miles to the condominium's nearest BART station, Caltrain station, MUNI Metro station (but not cable cars), and nearest major thoroughfare segment. The major north-south thoroughfares are 19th Ave/State Highway 1 and extensions, US Highway 101 and extensions, I-80, and I-280. I also control for the distance to the central business district, under the assumption from the monocentric city model that prices are highest closest to it, and distance to the geographical center of the city, allowing for the possibility that prices might rise more away from the center (and closer to the coast). Lastly, another variable includes the count of the number of public transit stops within a half-mile of the condominium.

within the first two years after the unit was built. The regressions also control for fixed effects: $Nbrhood_i$, a vector of neighborhood dummies and Y_t, a vector of year dummies, and M_t a vector of sales month dummies. The sample is restricted to single-family condos ($> 99\%$) built between 1849 and 2013 that have at least 291 sq. ft.[23]

In addition to the linear models described above, a second pair of regressions is run with the log of the price as the outcome variable. The linear-linear model is informative about whether the shuttles raised per-unit prices by enough to exceed relocation costs (Condition 2) and the log-linear model tells whether the shuttles yield a percent increase in price above the 2.1% allowable rent increase benchmark (Condition 1).

One issue with the city's transactions data is that officials made transcription mistakes.[24] These mistakes occasionally introduced outliers into the sample. Several strategies were employed to remedy these mistakes, notably cross-checking property information against the online real estate company Zillow's database. Remaining outliers in the lower tail were winsorized at the 1st percentile (following the suggestion of Bollinger and Chandra [25]). For outliers in the upper tail, they appeared to be exclusively cases where the city misclassified whole buildings as single unit, single-family condominiums. These observations were trimmed from the sample by establishing from the Zillow database that no condominium has ever been sold for more than $12,000,000 and had a square footage exceeding 6,000 square feet. 171 properties were excluded on the basis of excessive square footage and 305 transactions were excluded on the basis of a sales price of greater than $12,000,000. This left 31,150 in-sample transactions.

Table 2.7 shows selected descriptive statistics of the condominiums. Coefficients and robust standard errors from an estimation of Equation (2.2) are reported in Table 2.8. Figure

[23] 291 sq feet is the smallest condo sold in San Francisco, source: `http://sanfrancisco.cbslocal.com/2015/04/15/san-franciscos-smallest-condo-just-sold-for-415k-soma-south-of-market/`

[24] These include listing the entire building sales price as the individual condominium's sales price, missing condominium characteristics, bunching of build dates around certain years, such as 1900, and other issues.

2.2 shows the distribution of condominiums by neighborhood. Concerns about possible endogeneity between condominium supply and shuttle placement are addressed in the next section.

2.4.1.2 Instrumenting for Shuttle Placement

Consistent estimation of Equation (2.2) requires that the errors are uncorrelated with the outcome variables. Figure 2.1 strongly suggests that shuttles were not randomly placed and it cannot be ruled out *ex ante* that placement likelihood changed with local housing market conditions. If so, this condition would be violated and coefficient estimates would be biased. For example, areas with more at-fault evictions could be correlated with greater delinquency and ambient nuisance, and are avoided by upper middle-income technology workers.

The identification strategy in light of this potential endogeneity is based on exploiting an exogenous constraint on shuttle stop placements. The constraint is that the shuttles can only utilize public bus stops long enough on each side of the street to accommodate the 50-foot plus motorcoaches, called "bus zones".[25] Figure 2.3 shows the location of the 870 bus zones that match this description, which are about 25% of the SFMTA's total public bus stops. The greatest concentration of eligible bus zones is in the far northeast part of the city, near the central business district, and extending directly west to Golden Gate Park and directly southwest into the Inner Mission. Outside of these areas, eligible bus zones are much sparser.

Satisfying the exclusion criterion for using the eligible bus zones rests on the assumption that while local residents' location decisions may be influenced by being near a public bus stop,

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

residents are almost surely indifferent to the closest bus stop's *length*. The 50 foot constraint is relevant for shuttle placement but not for local prices or anything else correlated with evictions.

The bus zones' distribution alone is not enough to predict shuttle placement in the city for two reasons. The first reason is that even with the 50 foot restriction, there are many more eligible bus zones than shuttle stops. The second reason is that the bus zone distribution changes very little over time. Exploiting the static distribution for the instrument thus means predicting three aspects of the shuttle system from bus zones: which eligible bus zones were selected, when would an eligible bus zone have been selected, and rules for how coverage would have changed time period to time period. The first is done by identifying patterns for selecting among eligible bus zones exogenous to local housing markets. The second is done by introducing time variation into the patterns through creating interactions between bus zone characteristics and time variables. The last is done by introducing a gradient for how the shuttle system grew between time periods and within the established patterns.

2.4.1.2.1 Pattern for Selecting Among Eligible Bus Zones

Bus zone locations are very static throughout this time period, but selection on bus zone characteristics risks violating the exclusion criterion due to unobserved changes in public bus service impacting local housing conditions coincident with shuttle stop placement. This concern will be addressed for each bus zone characteristic selection.

To create an instrument for shuttle stop placements, it is necessary to know why certain bus zones were favored above others. Since the shuttles need to be able to reach Silicon Valley, proximity to the one of the north-south thoroughfares was prioritized. Bus zones accessible to the greatest number of people were also highly preferred.[26] Shuttles thus tended to stop in

[26] See Appendix A.2

bus zones close to centralized commuter transit points not offering transit service to Silicon Valley: MUNI Metro light rail stops.

Figures 2.5 and 2.6 overlay the commuter shuttle stop maps (streamlined from Figure 2.1), with the underlying transit options in San Francisco at each point in time. The first shuttle stops were along the BART system at Glen Park and a park and ride lot near Candlestick Park (Figure 2.1a), but the orientation of the system changed quickly once it became more popular. Figure 2.5 unmistakably shows that the shuttles clearly prioritized access to the thoroughfares in every phase of their growth. Figure 2.6 shows that after 2004, shuttle stops prioritizing the MUNI Metro lines originated around Civic Plaza,[27] where all the light rail lines converge and spread south and west along the "spokes" of the system. Stops oriented toward just the thoroughfares clustered first in the far north where there are few other transit options, before spreading south and west.

In Figure 2.5, an eligible bus zone is defined as being throughfare-adjacent if it is within a half-mile of a north-south thoroughfare. The exclusion criterion is satisfied after selecting on this eligible bus zone characteristic because the MUNI buses only serve San Francisco proper. This means that location and service changes in thoroughfare-adjacent eligible bus zones are not going to be made in response to increased commuting preferences to Silicon Valley along these thoroughfares.

In Figure 2.6, an eligible bus zone is defined as being metro-adjacent if it is both within a a half-mile of a MUNI Metro stop and within a mile of a north-south thoroughfare. The radius for eligible bus zones is expanded specifically for this selection because the companies are willing to extend their network a little further into the city if it meant conveniently accessing a commuter gathering point. The exclusion criterion is satisfied because both the MUNI buses and MUNI metro only serve San Francisco proper, so that any location changes or

[27] This is not directly shown, but the Civic Center stop started by Google in 2005. Other MUNI-Metro adjacent bus zones were added in 2006, as seen in Figure 2.6b

unobserved changes to service for either MUNI buses or MUNI metro, separate or jointly, are independent of increased demand for commuting to Silicon Valley. Housing demand might increase along the thoroughfare-adjacent MUNI lines, but not in a pattern coincident with shuttle stop placement, because this bus stop sub-distribution is poorly correlated with the housing stock's proximity to either the MUNI Metro (unit-weighted pairwise correlation=-0.061) or the thoroughfares (unit-weighted pairwise correlation=-0.048).

2.4.1.2.2 Introducing Time Variation

The absence of time variation is the second obstacle. Two methods are used to account for time variation in network growth. The first interacts each spatial location instrument with a monthly linear time trend to reflect consistent network expansion over time. Shuttle coverage growth is uneven over time, increasing greatly between 2006 and 2009, with more modest changes thereafter. The second method thus interacts each spatial location instrument with year dummies to capture yearly shocks to network placement, similar to the strategy used in Neumark, Zhang, and Ciccarella [93]. Together, these help predict how "favored" stops fitting the pattern described above were over time.

2.4.1.2.3 Direction of Shuttle Stop Growth

Needed last is a way to predict how placement likelihood changed over time within the fixed transit corridors. Ignoring in both cases the initial shuttle stop placements in September 2004 in Glen Park and Candlestick Park, there are clear directions for shuttle stop growth. The first shuttle stop coverage pattern shows a north to south-by-southwest gradient across thoroughfare-adjacent eligible bus zones (Figure 2.5). The second shuttle coverage pattern shows a northeast to southwest gradient across eligible bus zones near both a thoroughfare and near a MUNI metro station (Figure 2.6). Thus indicators for how far east and how far north each bus zone is are used as instruments for the gradient.[28] The east measure used rises monotonically with how far east the bus zone is (in the longitudinal sense), and likewise the north measure used rises monotonically with how far north the bus zone is (in the latitudinal sense).

Thus, for each building, the above exogenous bus zone characteristics are exploited to calculate whether the building would have shuttle coverage in a given time period. Whether nearby bus zones are thoroughfare or MUNI-adjacent is interacted with time trends and fixed effects. This reflects how the likelihood for placement increased within the aforementioned transit corridors over time. Each term is then interacted with how far east and how far north each bus zone is, reflecting that the coverage probability increased over time for southern and western bus zones. This approach also permits a straightforward interpretation For example, the expectation is that thoroughfare-adjacent bus zones in the east (having a high "east" value) are going to be assigned the highest probability of getting early shuttle coverage. As more western stops get selected for a shuttle, the average east value of a shuttle stop will decline over time. Thus, the expected coefficient on the time-trended thoroughfare-adjacent flag interacted with the east value will be negative.

[28] The simplest way to simulate movement in a cardinal direction is to use the Universal Transverse Mercator (UTM)'s easting and northing coordinates. They are used in place of latitude and longitude in UTM projected coordinate system, which overlays the earth with unique ellipsoid zones so that distances on the globe can be measured directly in meters. Easting in this case is how many kilometers east the point is from the zone's point of origin, which is the intersection between the zone's central meridian and the Equator. Likewise, northing is how many meters north the point is from the zone's point of origin. UTM is a remarkably accurate system for calculating distances in this way, as one study found that it is accurate to within 9 mm over the entire ellipsoid (Karney [70]). San Francisco lies in UTM Zone 10 far north of the Equator, so an arbitrary point in the ocean just southwest of the peninsula was chosen as a new point of origin (specifically Easting=543500, Northing=4173500 in UTM Zone 10) and the easting and northing coordinates were rescaled and converted to kilometers so that interaction terms would not become arbitrarily large.

2.4.1.2.4 Regressing for Shuttle Coverage

The first stage regression takes the form of:

$$\begin{aligned} Shuttle_{it} &= \pi_0 + \pi_1 BZChars_{it} + \pi_2 Dir_i \\ &+ \pi_3 BZChars_{it} * Dir_i + \pi_4 BZChars_{it} * Dir_i * t + u_{it} \end{aligned} \tag{2.3}$$

where each variable prefixed with "BZ" was calculated as a spatially-weighted average of the characteristics all the eligible bus zone within a half-mile of each condominium.

The $BZChars_{it}$ are:

1. $\sum_{k=1}^{B_t} w_{ikt} \mathbb{1}\{\leq 1/2 \text{ Mile to Thrufare}_{kt}\}$ is the spatially-weighted sum of nearby bus zones within 1/2 mile of a thoroughfare.
2. $\sum_{k=1}^{B_t} w_{ikt} \mathbb{1}\{\leq 1 \text{ Mile to Thrufare}_{kt} \ \& \ \leq 1/2 \text{ Mile to Metro}_{kt}\}$ is the spatially-weighted sum of nearby bus zones within a 1/2 mile of a light rail stop and are thoroughfare-adjacent.
3. $\sum_{k=1}^{B_t} w_{ikt} Length_{kt}$ is the spatially-weighted sum of the lengths of nearby bus zones.
4. $\sum_{k=1}^{B_t} w_{ikt} Shelter_{kt}$ is the spatially-weighted sum of nearby bus zones with a shelter.

where B_t is the total number of eligible bus zones in time t. w_{ikt} is the double-power distance weight assigned to each bus zone k based on its proximity to condominium i. This is calculated as

$$w_{ikt} = \begin{cases} \left[1 - \left(\frac{d_{ik}}{d}\right)^2\right]^2 & 0 \leq d_{ik} \leq d \\ 0 & d_{ik} > d \end{cases}$$

where d here is the 1/2 mile cut-off for being in the vicinity of condominium i and d_{ik} is the distance between condominium i and bus zone k.

Dir refers to either $East$ and $North$ for either the nearby bus zones' average east or north position. These averages were then aggregated to the neighborhood level when necessary. Each term is thus the interaction between the average bus zone's easting or northing coordinate, the spatially weighted fraction of nearby bus zones' characteristics, and time trends to reflect the directional gradient. X_{it} contains all the exogenous variables from Equation (2.2).

2.4.1.3 Hedonic Results

In the linear specification shown in Table 2.8, shuttles lower condominium prices by about \$11,643, but the log price specification indicates that shuttles raise prices by about 2.4%. Price changes from the shuttles fail both eviction cost conditions (see Section 2.2.2.3), but as stated in Section 2.4.1.1, these regressions could suffer from bias from outliers. I thus also include in Table 2.8 two quantile regressions at the median as a check against outliers. The quantile regressions show that when outliers are factored out, the shuttle raised prices by about \$27,100 and yielded a price increase of about 4.1%, satisfying both eviction cost conditions.

Table 2.9 reports the first-stage estimates for the shuttle stop instrument. As predicted in Section 2.4.1.2, MUNI Metro-adjacent bus zones in southwestern San Francisco became more likely to host a shuttle stop as time increased. In Table 2.10, the coefficient on linear time-trended MUNI Metro-adjacent eligible bus zones interacted with the Northing and Easting coordinate is negative, which means that the distance between the fixed point and the bus zones hosting a shuttle stop steadily shrank (equivalent to moving southwest-ward). The same coefficient on thoroughfare-adjacent bus zones interacted with a time trend and the easting coordinate is positive, as the shuttle stops steadily drifted northward along the thoroughfares. The coefficient on time-trended bus zone lengths is positive, as successively larger bus zones were selected over time. These results confirm the instrument's intuition.

The second stage results estimation results indicate three things: the instrument is not weak; the original hedonic price regressions understate the shuttle's hedonic impact; and that the shuttles almost certainly changed market prices by enough to incentivize evictions. The Anderson-Rubin Wald F statistic on the instrument is 49.05, permitting rejection of the weak instrument null hypothesis. The hedonic price underestimation in ordinary least squares is likely due to the shuttles congregating in gentrifying areas, which initially tend to have lower-than-median housing prices. Lastly, the 2SLS estimates on the shuttle stops show that prices changed by enough to exceed the annual allowable rent increase (Condition 1 10.5%>2.1%) and the maximum relocation payment per unit (Condition 2 - $51,356.84>$26,055). The median regressions move in the opposite direction, suggesting that the median condo was less affected than the average condo. While the log-linear estimate on shuttle exposure is just shy of the maximal allowable rent increase, 2.0% is still a greater price increase than has been permitted by the city in all but 1 year since 2003.[29] Having established that the shuttles changed prices (and presumably rents) on average by enough to incentivize economic evictions, the paper now proceeds to estimating how evictions changed in response to the shuttles' price shock.

[29]"Allowable Annual Rent Increases", The Residential Rent Stabilization and Arbitration Board of the City and County of San Francisco, 2016.

2.4.2 Testing for Economic Evictions

The empirical strategy is pursued on multiple outcomes for economic evictions. The two eviction types are investigated separately: all at-fault evictions and no-fault evictions (Ellis+OMI's). At-fault evictions are used to test whether landlords use economic evictions to exploit vacancy decontrol. Owner move-in and Ellis Act evictions summed test whether landlords respond to a price shock by withdrawing their units from the market. At-faults and no-faults are not combined because they have very different cost structures and incentives and may in fact be substitutes. Any evidence for economic evictions of either type will be interpreted as evidence in favor of the hypothesis.

The regressions control for changes to rent control and evictions policies over the study's observational period.[30] Year/month effects control for general policy exposure, but most policies were only applicable to controlled buildings. Each regression thus controls for the policies interacted with rent control status, under the assumption that eviction policy changes are endogenous with respect to the city's overall eviction rate, but exogenous to changes in any individual building or neighborhood. The key identification assumption in these regressions is that there are no other unobserved shocks to the outcomes coincident with shuttle placement and shuttle placement interacted with rent control status that affected eviction outcomes.

A first series of regressions is performed on an unbalanced monthly panel of rental buildings, where the outcome is whether the building had an eviction in a given month ($\mathbb{1}\{Eviction_{it}\}$) and rent control and shuttle coverage are indicators. This panel contains 36,592 buildings that are commonly used as rentals: apartments, flats, and multi-family dwellings observed for at least a year and up to 126 months. A second series of regressions is performed using an unbalanced yearly panel of all 128,990 residential buildings, including single-family detached

[30] These changes can be found in Tables 2.3 and 2.4.

homes, because owners are just as free to offer these for rent as apartment landlords.

Panel fixed effects and instrumental variable estimates are reported. The instrument is the same as reported in Section 2.4.1.2, except that the interaction between rent control and the shuttles could also be endogenous. Shuttles may have favored controlled apartments because technology employees are willing to pay the premium, anticipating that high job security necessitates prolonged tenures. Figure 2.4 shows the fraction rent controlled per neighborhood, and a comparison with Figures 2.1 and 2.3 shows that areas with a high rent control fraction are both well-covered by the shuttles and have many bus zones. Thus, Equation (2.3) is augmented by interactions with the rent control measure. While in the long-run, rent control status is endogenous to evictions, the lags are so considerable that at the time scale observed in this study, rent control status can be safely treated as exogenous.

2.4.2.1 Estimating Building-Level Eviction Occurrence

A panel fixed effects model is employed to estimate the impact of price shocks on evictions at the individual building level. This estimation sample is comprised of 36,592 rental buildings that observed for at least a year that could have at least one rental unit.[31] Evictions are relatively rare on a monthly basis: a total of 4,987 at-fault and 2,181 no-fault notices over the course of the study period were issued in 3,590 of the in-sample buildings.[32] This equation

[31] This was defined as all residential buildings that are not government housing or single-family dwellings, condominiums, co-ops, or residential hotels.

[32] There are many more eviction notices given out in total, as can be seen in Table 2.2. Many of these are in fact for illegal units in the San Francisco's single-family housing stock, and so are outside the scope of this study.

takes the form of

$$
\begin{aligned}
&Eviction_{it} = \delta_0 + \delta_1 Shuttle_{it} + \delta_2 \left(Shuttle_{it} * RentControl_i\right) + \delta_3 Transit_{it} \\
&+ \delta_4 (Policies_t * RentControl_i) + \delta_5 YYMM_t + \delta_6 Building_i \\
&+ \delta_7 Nbrhd_i * Year_t + \delta_8 Nbrhd_i * t + \epsilon_{it},
\end{aligned} \tag{2.4}
$$

where $Eviction_{it}$ is an indicator for whether an eviction occurred in building i in month t. $Shuttle_{it}$ is an indicator for whether the building is within a half-mile of a shuttle stop. $Policies_t$ refers to the changes in eviction policies detailed in Table 2.3 and are interacted with rent control status, $RentControl_{it}$, which captures how these policies might have had the strongest impact on dampening controlled evictions.[33] $OtherTransit_{it}$ differs from 2.4.1.1 in that it only has measures for distances to the MUNI Metro and interactions with rent control status - all other distance measures are dropped in this specification because they are not time-varying. $TimePeriod_t$ is either a vector of year/month fixed effects in the monthly panel, controlling for month-specific shocks to eviction rates, or a vector of year fixed effects in the yearly panel. $Building_i$ is a vector of building-specific fixed effects to control for all time-invariant building characteristics. Lastly, the disaggregated nature of Equation (2.4) allows the identification assumption to strengthened by adding $Nbrhd_i * Year_t$ to control for neighborhood-specific yearly shocks that might be coincident with shuttle stop placement.

The key variables of interest are $RentControl_{it}$, $Shuttle_{it}$ and the interaction term between the shuttle coverage measure and the rent control measure because it captures how rent-controlled landlords react to increases in market rents. Since the transit amenity's value should be highest within the immediate vicinity of a stop, economic evictions occur if the interaction is positive. Standard errors are clustered at the neighborhood level,[34] and each

[33]For the imposition of relocation payments on owner move-in evictions in 2006 and the extension of relocation payments to all Ellis'd tenants in 2005 described in Table 2.4, I create dummy variables also interacted with rent control status for both policy changes.

[34]While this is something of a strong assumption, it is not unreasonable to assume that buildings are much

observation is weighted by its initial number of housing units. Results are reported in Table 2.11.

2.5 Results

Key results are found in Table 2.11. Table 2.11 shows the building-level ordinary least squares and IV results after applying the instrument for shuttle coverage. The top panel (Columns 1-6) shows no evidence for economic evictions for at-fault evictions and some evidence for economic no-fault evictions at the 10% level in the monthly panel (Column 2). This is unsurprising given that the endogeneity bias in shuttle locations *understates* the strength of the shuttles' price shock, which can be seen when comparing the hedonic results from Table 2.8 to the IV results in Table 2.10. While not significant, the monthly panel (Columns 1-2) shows the predicted signs for at-fault evictions (Column 1): rent control positively predicts an eviction and exposure to the rent increase proxy increases the probability of an at-fault eviction in controlled buildings. For no-fault evictions, being controlled statistically significantly increases the probability of an eviction by about 0.04% and this increases further when exposed to a shuttle stop by 0.0007%.

One interesting result from Column (1) is that shuttle coverage significantly decreases at-fault evictions across uncontrolled buildings by 0.8%. The interpretation here is that price increases make uncontrolled landlords less likely to evict. Landlords may be anticipating that they can successfully exploit the next lease renewal to price out an undesired tenant rather than having to evict them.

As a check that shuttle exposure and evictions are not being mistimed in the yearly panel, the shuttle variables (and instruments) are lagged back a period. Results for a one-year lag on the

more independent across neighborhoods than not, while being possibly highly correlated within the same neighborhood.

shuttle and its interaction term are reported in Columns 5-6. Neither the contemporaneous regression nor the regression with lagged coverage have consistent results.

The instrumental variables estimates in Columns 7-12 in Table 2.11 show more robust evidence for economic evictions. Assuming that the percent change in rents is the same as the percent change in condominium prices after exposure to a shuttle, Table 2.11 shows that the monthly probability a controlled building will undergo a no-fault eviction increases by 0.019% when rents increase by the implied hedonic value of the shuttles. Similarly, the results for at-fault evictions in Column 7 show that landlords become almost 1% more likely to have at least one eviction per month if they are within a half-mile of a shuttle stop. The yearly probability for no-faults rises to 0.22%, which is close to what the estimated monthly probability times 12 equals (0.00019*12=0.00228). This shows that the finding in favor of economic no-fault evictions is at least somewhat robust to sample composition considerations. The estimates in the lagged panel (Columns 11-12) are not statistically significant, but the signs remain the same across the two specifications.

To estimate the city-wide eviction impact using the results from the all-rental monthly panel, start by taking the 22,251 controlled buildings within a half-miled of a shuttle as the base. From Table 2.11, the yearly number of "extra" controlled buildings with no-fault evictions has increased by 22,251*0.0022=50.73. The equivalent extra number among at-fault evictions is 218 buildings. This represents a substantial shrinkage in a fixed stock of controlled housing year on year.

2.6 Discussion

This paper shows that economic evictions occur in at least certain circumstances, in spite of being assumed away in all previous rent control studies. The question now is: what does this mean for the rent control literature and rent control policy?

The first and most obvious conclusion is that landlords can exploit rent control's loopholes to deny tenants their *de facto* duration subsidy. Tenant protections are a cornerstone of the rent control policy package, and the results of this paper shows that they are weaker than others have believed. In practice, rent control proponent's argue that its costs are justified by the security it offers to economically vulnerable populations. The existence of systematic economic evictions shows that fewer tenants actually reap the benefits than are advertised, and as the number of controlled units shrink, fewer will do so in the future.

Incorporating economic evictions into rent control models raises many questions, but there are two I will outline briefly as particularly relevant areas for future research. The first is how evictions change rent control's distributional impact. Rent control's beneficiaries do not seem to be disproportionately low-income and minority (Navarro [92]; Sims [106]; Turner and Malpezzi [114]). Glaeser [49] and Gyourko and Linneman [56] point out that controlled housing is likely non-price rationed because its supply is highly inelastic and price signals cannot be used to otherwise clear the market. Landlords are thus empowered to lean on their personal biases, so that on net, poor families do not seem more likely to have controlled units.

However, landlord bias may not be the only cause. Both no-fault and at-fault evictions could work together to skew the income distribution of controlled tenants upwards. No-fault evictions steadily shrink the controlled housing supply, causing the controlled unit premium to rise. Systematically, the long-run pressure on base rents might price out poorer tenants, even poor long-stayers. At-fault evictions may then whittle down the number of existing poor tenants, when landlords seize on the greater propensity for poor tenants to experience failures to pay rent. The two eviction types would then be working in tandem to weight the controlled tenant pool towards the wealthy.

Counterbalancing these pressures is that rent control is confined to older, likelier shabbier buildings, which are more likely to command lower prices in a heterogeneous housing market.[35] Unfortunately, the literature is not extensive enough to rule out the possibility that in a city experiencing strong rent pressures, controlled unit allocation to poor residents has a negative long-term outlook. The estimated 218 extra controlled buildings with at least one at-fault eviction due to the shuttle system means many residents are also displaced without relocation payments. Given the tightness of the housing market in San Francisco, many would be unable to find new controlled housing. Desmond [39] shows in a demographic study of evictions in an uncontrolled market that at-fault evictions disproportionately happen to poor, female, and minority tenants. Rent control thus does not seem to retain many of the people it seeks to help, and circumstantial evidence shows that one of its perverse incentives targets the very people it seeks to aid.

On the whole, economic evictions are a rent control feature policymakers would be hard pressed to end. Obvious legislative solutions like legislating new controlled housing risk dampening the long-run housing supply (McFarlane [83]). Banning all evictions outright would almost surely be met with intense landlord resistance, and large-scale housing unit withdrawal (Miron [84]). Increasing penalties on no-fault evictions will likely increase the number of units that are off the market in limbo. It is left to future research to establish how economic evictions change the social welfare benefits of rent control.

2.7 Conclusion

This paper tested the widespread assumption in the rent control literature that economic evictions do or cannot occur under these policy regimes. Exploiting unique properties of San

[35] The literature on rent control and landlord and tenant maintenance offers mixed evidence for the quality of controlled buildings (Moon and Stotsky [85]; Gyourko and Linneman [57]; Arnott and Shevyakhova [7]; McFarlane [83]; Kutty [78]), so it is unclear as of this writing what the distributional equilibrium would be when evictions are factored in.

Francisco, namely an identifiable, locally differentiated change in prices, eviction responses to free market rent changes were tested. Landlords were found to engage in short-term unit-clearing evictions that kept their units in the controlled market and also evictions that allowed them to begin the process of switching to the uncontrolled market, even at the expense of decreasing their medium-term housing supply. The shuttle-induced rent increase for vacant units of about 10.5% yields an additional 218 controlled buildings with at least one at-fault eviction and 51 unit-withdrawing evictions in the controlled market per year. In light of these findings, future research could revisit the distributional and housing market aspects of rent control policy. Policymakers looking to protect the controlled housing supply would be advised to make the provisions for market exit ("no-fault" evictions) more expensive.

The 15 Grounds for “Just Cause" Eviction in San Francisco

TABLE 2.1

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

Table 2.1 enumerates the reasons a landlord may reclaim a rent-controlled unit. The “At-Fault" evictions refer to the 7 ways a tenant may be evicted for breaching the rental contract in some fashion, and “No-Fault" refers to the 8 ways a tenant may be evicted even if not in breach of the lease.

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

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

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

[d] “Good Samaritan" status is temporary housing for tenants fleeing a natural disaster.

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

Yearly Counts of “Just-Cause" Eviction Notices by Eviction Type: 2003-2013

TABLE 2.2

	Owner Move-In	**Demolitions[a]**	**Ellis Act[a]**	**At-Fault**	**TOTAL**	**Rental CPI[b]**
2003	275	79	95	725	1174	2.9%
2004	263	49	238	969	1519	2.6%
2005	200	37	246	1270	1753	2.9%
2006	180	33	197	1306	1716	3.4%
2007	156	31	199	1499	1885	4.1%
2008	132	28	170	1422	1752	3.5%
2009	103	27	36	1321	1487	2.2%
2010	116	25	58	1444	1643	0.2%
2011	107	39	45	1564	1745	1.7%
2012	154	32	77	1611	1874	2.6%
2013	230	86	180	1903	2399	2.7%
TOTAL	1641	377	1446	14309	18957	

Table 2.2 presents counts of eviction notices for all at-fault evictions and the three largest categories of no-fault evictions in San Francisco from 2004-2013. For owner move-in and at-fault evictions, these are generally applied to only one unit. For demolitions and Ellis Act evictions, these almost exclusively occur for an entire building, so potentially many units (and tenants) are impacted. There are three other categories of no-fault evictions, but none of these have more than 30 evictions over the course of the study period.

[a] These are counts of buildings where an Ellis Act or demolition notice to quit was given. The total number of people and units affected is a few multiples of these counts.

[b] “Rental CPI" is the annual percent change in the Consumer Price Index for the Rent of Primary Residence in the San Francisco-Oakland-San Jose, California Combined Metropolitan Statistical Area reported by the Bureau of Labor Statistics.

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

Policy Changes Regulating Evictions in San Francisco: January 2002-December 2013

TABLE 2.3

Description	Start Date	End Date
General Eviction Rules		
Require all eviction notices except those for non-payment of rent to be in writing and filed with the Rent Board. The grounds cited in an eviction notice must be adhered to regardless of any. separate agreement between tenant and landlord.	2002	
Landlords who wish to terminate that tenancy are no longer required to give 60 days notice, only 30-days notice, for tenants who have resided in the premises for one year or more.	1/1/2006	
Owners of properties with two or more residential units must disclose to any prospective purchaser the legal grounds for terminating the tenancy of each unit vacant at the close of escrow and whether the unit was occupied by an elderly or disabled tenant at the time the tenancy was terminated.	6/6/2006	
Reinstated the prior requirement of a 60 day notice to terminate a tenancy without a tenant fault good cause for any tenant or resident residing in the unit for a year or more.	1/1/2007	12/31/2009
A tenant who has resided in the unit for at least one year, and has a child under the age of 18 who also resides in the unit, may not be evicted during the school year for an OMI eviction.	3/14/2010	
Tenant may not be evicted for violation of a unilaterally imposed change in the terms of a tenancy unless the tenant previously accepted it in or the newly imposed term is authorized by the Rent Ordinance.	12/14/2011	2/1/2012
Allows a landlord to evict a tenant for violation of a unilaterally imposed change in terms where the change is required by law	2/1/2012	
Condo conversion evictions are suspended	8/1/2012	
Ellis Act		
Landlords must state in Ellis Act eviction notices that tenants have the right to relocation payments and the amount which the landlord believes to be due.	7/25/2005	1/30/2006
Landlords are no longer required to state the amount of relocation payment the landlord believes to be due to the tenant	1/31/2006	
Owner Move-In		
Landlords seeking to challenge a tenants' protected status for an OMI eviction have to file a petition rather than seeking a court order.	2006	

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

Relocation Payments for No-Fault Evictions: February 2000-February 2014

TABLE 2.4[a]

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

Table 2.4 shows the mandated relocation payments given to tenants for Ellis Act evictions and all other no-fault evictions. “Low Income Tenants" are the payments originally only given to poor tenants before August 2004 for Ellis Act evictions before being extended to all tenants. “General Tenants" are the relocation payments that were given to any controlled tenant. All amounts are in nominal US dollars..

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

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

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

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

Major City Rent Control and Evictions Policies, October 2016

TABLE 2.5

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

[a] Sources: Los Angeles Municipal Code Chapter 151. Coverage figure comes from the Office of the Mayor of Los Angeles, January 26, 2016. "Mayor Garcetti Announces New Access to Information on L.A.'s Rent-Stabilized Buildings.", last accessed October 14, 2016.
[b] Sources: Oakland Municipal Code §8.22 et seq. Coverage figure comes Sam Levin, July, 1, 2015. "When Landlords Target Tenants in Rent-Controlled Buildings." East Bay Express.
[c] Sources: The New York State Rent Regulation Reform Act of 1997, 1997 New York Laws 116; The Rent Act of 2015, 2015 New York Laws 20; The New York State Emergency Tenant Protection Act of 1974, 1974 New York Laws 576 §5-a. Coverage figure comes from Sieg, Holger and Chamna Yoon, 2016. "Waiting for Affordable Housing in New York City." Working Paper.
[d] In both cities, landlords can appeal for a rent increase on new base rents of up to 30% if rents in comparable units are shown to be higher.
[f] Sources: San José Municipal Code, Apartment Ordinance, Chapter 17.23. Coverage figure is from San José Municipal Ordinance No. 29730, p. 1.
[g] Sources: Code of the District of Columbia, Chapter 42. Coverage figure comes from Tatian, Peter A. and Ashley Williams, 2011. "A Rent Control Report for the District of Columbia." The Urban Institute, Washington, D.C.
Note: All cities exempt new buildings, but developers in NYC can get a tax rebate on new buildings if they agree to controls for 10 years. After 10 years, all new vacancies are fully decontrolled. Source: NYC Rent Stabilization Code, Sections 2520.11(o), (p), (r), and (s).

Key Characteristics of San Francisco's Housing Stock

TABLE 2.6

	2003	2008	2013
Total Housing Units	377,182	389,787	405,021
Average Year Built	1935	1937	1938
As Percent of Total Housing Stock			
Rent-Controlled	55.5%	53.5%	52.0%
Within 1/2 Mile of Shuttle Stop	0.0%	56.5%	59.7%
Condominiums	9.5%	12.1%	14.2%
Apartment Buildings	57.5%	56.2%	55.3%
Residences	29.7%	28.7%	27.7%

Table 2.6 presents averages of key characteristics of the housing stock in San Francisco for three selected years. The sample comprises all buildings in San Francisco in the years specified with a housing property code that indicates it is a condominium, house, apartment building, flat, or townhome, summed by year in Row 1. Condominiums are denoted as buildings in their own right, because they are recorded individually by the city even if they are in a multi-unit building. Average Year Built (Row 2) is the mean year built by property. Rows 4-9 give the fraction of housing units that are rent-controlled; the fraction that are within a half-mile of a shuttle stop as of June of that year; the fraction that are condominiums; apartment buildings; residences; and the fraction that have no-fault evictions, respectively.

In-Sample Sold Condominium Characteristics, July 2003-December 2013

TABLE 2.7

	Mean	Median
Sales Price	$ 1,750,049	$ 797,900
# of Baths	1.63	2
# of Beds	1.87	2
Sq. Ft.	1170.32	1096
Year Built	1975.9	1996
Distance to the Central Business District (mi.)	2.12	1.61
Distance to BART (mi.)	1.03	0.88
Distance to CalTrain (mi.)	1.71	1.60
Distance to MUNI (mi.)	0.59	0.40
Distance to Thoroughfare (mi.)	0.45	0.33
Distance to Any Company Shuttle (mi.)	1.64	0.52
Within a Half-Mile of Shuttle	0.47	0

Table 2.7 presents summary statistics on housing characteristics of condominiums sold in San Francisco between July 2003 and December 2013. The sample is defined as single-use, single-family condominiums with no more than 3.5 bathrooms and 4 bedrooms and a minimum of 291 square feet. “Within a Half-Mile of Shuttle" is coded 0/1 in the underlying data, so that the mean of “Within a Half-Mile of Shuttle" represents the fraction across the entire time period of shuttles that were within a half mile of a shuttle stop.

Hedonic Price Regression on Condominium Sales: July 2003-December 2013

TABLE 2.8

	OLS		Quantile Reg	
Sales Price in $2013	**Linear**	**Log**	**Linear**	**Log**
$\mathbb{1}$\{**Shuttle**\}	-11643.59	0.024	27066.31***	0.041***
	(16944.43)	(0.017)	(9758.35)	(0.013)
Sq. Ft	690.66***	0.0006***	533.35**	0.0005**
	(95.71)	(0.00004)	(71.15)	(0.00001)
Adj R^2	0.49	0.68	0.44	0.66
N	30,150	30,150	32,453	32,453
Distance Controls	Y	Y	Y	Y
Other Property Controls	Y	Y	Y	Y
Month Dummies	Y	Y	Y	Y
Year Dummies	Y	Y	Y	Y
Neighborhood Dummies	Y	Y	Y	Y
Zone Dummies	Y	Y	Y	Y

The dependent variable in Table 2.8 is the reported sales price for a single-family, single-use condominium in 2013 dollars. Point estimates are obtained by regressing on an indicator variable for whether the condominium is within a half-mile of a shuttle stop ("$\mathbb{1}$\{Shuttle\}"), as well as other transit/location controls, property characteristics, neighborhood effects, indicators for the condominium's zoning area, and sales year and month effects. Distance controls include the distance to the nearest north/south thoroughfare, BART station, Caltrain station, MUNI Metro light rail station, city geographical center, and the central business district, as well as indicators for being a half mile within each point/transit node. Other property controls include the number of beds, the number of baths (including half-baths), the number of public bus stops within a half mile, and an indicator for whether the condominium is brand-new. Zone controls refer to various historic and restricted development districts throughout the city as defined by the City of San Francisco Planning Department.

Cluster-robust standard errors are reported in parentheses.

† p<0.10, * p<0.05, ** p<0.01.

Instrument Regression Results, July 2003-December 2013

TABLE 2.9

	First-Stage
$\mathbb{1}\{\leq 1/2 \text{ Mile to Thrufare}_{it}\}$	-0.22201
	(1.21371)
$\mathbb{1}\{\leq 1 \text{ Mile to Thrufare}_{it} \ \& \ \leq 1/2 \text{ Mile to Metro}_{it}\}$	0.94903
	(0.94232)
$Length_{it}$	0.00006**
	(0.00003)
$\mathbb{1}\{\leq 1/2 \text{ Mile to Thrufare}_{it}\} * North_i * t$	0.00005**
	(0.00002)
$\mathbb{1}\{\leq 1/2 \text{ Mile to Thrufare}_{it}\} * East_i * t$	0.00002
	(0.00003)
$\mathbb{1}\{\leq 1 \text{ Mile to Thrufare}_{it} \ \& \ \leq 1/2 \text{ Mile to Metro}_{it}\} * North_i * t$	-0.00017***
	(0.00005)
$\mathbb{1}\{\leq 1 \text{ Mile to Thrufare}_{it} \ \& \ \leq 1/2 \text{ Mile to Metro}_{it}\} * East_i * t$	-0.00001
	(0.00001)
Adj R^2	0.65
Other Property Controls	Y
Year/Month Dummies	Y
Neighborhood & Zone Dummies	Y

The dependent variable in Table 2.9 is the reported sales price for a single-family, single-use condominium in 2013 dollars. Point estimates are obtained by regressing on an indicator variable for whether the condominium is within a half-mile of a shuttle stop "Other Propety Controls" include the distance to the nearest north/south thoroughfare, BART station, Caltrain station, MUNI Metro light rail station, city geographical center, and the central business district, as well as indicators for being a half mile within each point/transit node. It also includes the number of beds, the number of baths (including half-baths), the number of public bus stops within a half mile, and an indicator for whether the condominium is brand-new. Zone controls refer to various historic and restricted development districts throughout the city as defined by the City of San Francisco Planning Department.

Cluster-robust standard errors are reported in parentheses, * p<0.10, ** p<0.05, *** p<0.01.

2nd Stage IV Results for Hedonic Price Regressions

TABLE 2.10

	OLS		**Quantile Reg**	
Sales Price in $2013	**Linear**	**Log**	**Linear**	**Log**
$\mathbb{1}\{$**Shuttle**$\}$	51356.84	0.105*	12748.03†	0.020***
	(36927.72)	(0.055)	(6891.88)	(0.008)
Sq. Ft	688.13***	0.0006***	534.29	0.0005***
	(95.54)	(0.00004)	(71.10)	(0.0001)
1st-Stage Wald F-Stat	49.05	49.05		
A-R p-value	0.00	0.00		
Hansen's J-stat	36.05	38.99		
Hansen p-value	0.11	0.06		
Distance Controls	Y	Y	Y	Y
Other Property Controls	Y	Y	Y	Y
Month Dummies	Y	Y	Y	Y
Year Dummies	Y	Y	Y	Y
Neighborhood Dummies	Y	Y	Y	Y
Zone Dummies	Y	Y	Y	Y

Table 2.10 reports regression results for the second stage of the instrumental variable regression described in Section 2.4.1.2. The second stage regresses condominium sales prices (expressed in 2013 dollars) on fitted values for $\mathbb{1}\{\text{Shuttle}\}$ and also includes all other controls described in Equation (2.2) and described in Table 2.8. Quantile regression results are reported at the median.

Cluster-robust standard errors are reported in parentheses, * $p<0.10$, ** $p<0.05$, *** $p<0.01$.

Eviction Probability Estimates, Panel Fixed Effects and IV Estimates, Jul 2003-Dec 2013

TABLE 2.11

Panel Fixed Effects Estimates

	Monthly Panel		**Yearly Panel**			
			(t)		(t-1)	
	At-Fault	**No-Fault**	**At-Fault**	**No-Fault**	**At-Fault**	**No-Fault**
	(1)	(2)	(3)	(4)	(5)	(6)
Rent Control	0.0010	0.0004***	0.0286**	-0.0064	0.0109	-0.0085*
	(0.0031)	(0.0001)	(0.0116)	(0.0041)	(0.0122)	(0.0048)
Shuttle	-0.0081*	-0.00003	-0.0087	0.0001	0.0010	0.0000
	(0.0048)	(0.00004)	(0.0084)	(0.0002)	(0.0125)	(0.0002)
Interaction	0.0052	0.00007*	0.0054	0.0010	-0.0020	0.0004
	(0.0048)	(0.00004)	(0.0089)	(0.0006)	(0.0115)	(0.0005)
Adj R^2	0.24	0.10	0.32	0.05	0.33	0.05

IV Estimates

	Monthly Panel		**Yearly Panel**			
			(t)		(t-1)	
	At-Fault	**No-Fault**	**At-Fault**	**No-Fault**	**At-Fault**	**No-Fault**
	(7)	(8)	(9)	(10)	(11)	(12)
Rent Control	0.0028	0.00034***	0.0255**	-0.0060	0.0128	-0.0075
	(0.0033)	(0.00008)	(0.0121)	(0.0041)	(0.0142)	(0.0050)
Shuttle	-0.0225**	-0.00003	0.0000	-0.0002	0.0133	-0.0007
	(0.0110)	(0.00010)	(0.0216)	(0.0007)	(0.0212)	(0.0006)
Interaction	0.0098**	0.00019**	-0.0133	0.0022*	-0.0153	0.0007
	(0.0046)	(0.00008)	(0.0167)	(0.0011)	(0.0165)	(0.011)
First-Stage F Stat	29.52	29.52	79.83	79.83	37.51	37.51
Hansen's J	42.20	37.54	36.67	29.10	34.98	28.44
p-value	0.29	0.49	0.53	0.85	0.61	0.87
YYMM FEs	Y	Y	N	N	N	N
Year FEs	N	N	Y	Y	Y	Y

Table 2.11 displays the panel fixed effects and instrumental variable estimates for building-level eviction probabilities of an at-fault eviction and no-fault eviction for the stated time period. Point estimates are obtained by estimating coefficients for the model given in Equation (2.4). In the monthly panel, all three variables of interest are indicators for the presence of rent control, being within a half-mile of a shuttle, and the interaction between the two. In the yearly panel, $Shuttle_{it}$ is the fraction of year t that the building was adjacent to a shuttle stop. Columns (5), (6), (11), and (12) lag the shuttle and its interaction back one period as a precaution against mis-timing between the observed eviction and shuttle exposure.
* p<0.10, ** p<0.05, *** p<0.01.

Figure 2.1: Evolution of Apple, Electronic Arts, Facebook, and Google Shuttle Stops

September 2004 **April 2006** **May 2008**

September 2009 **December 2013**

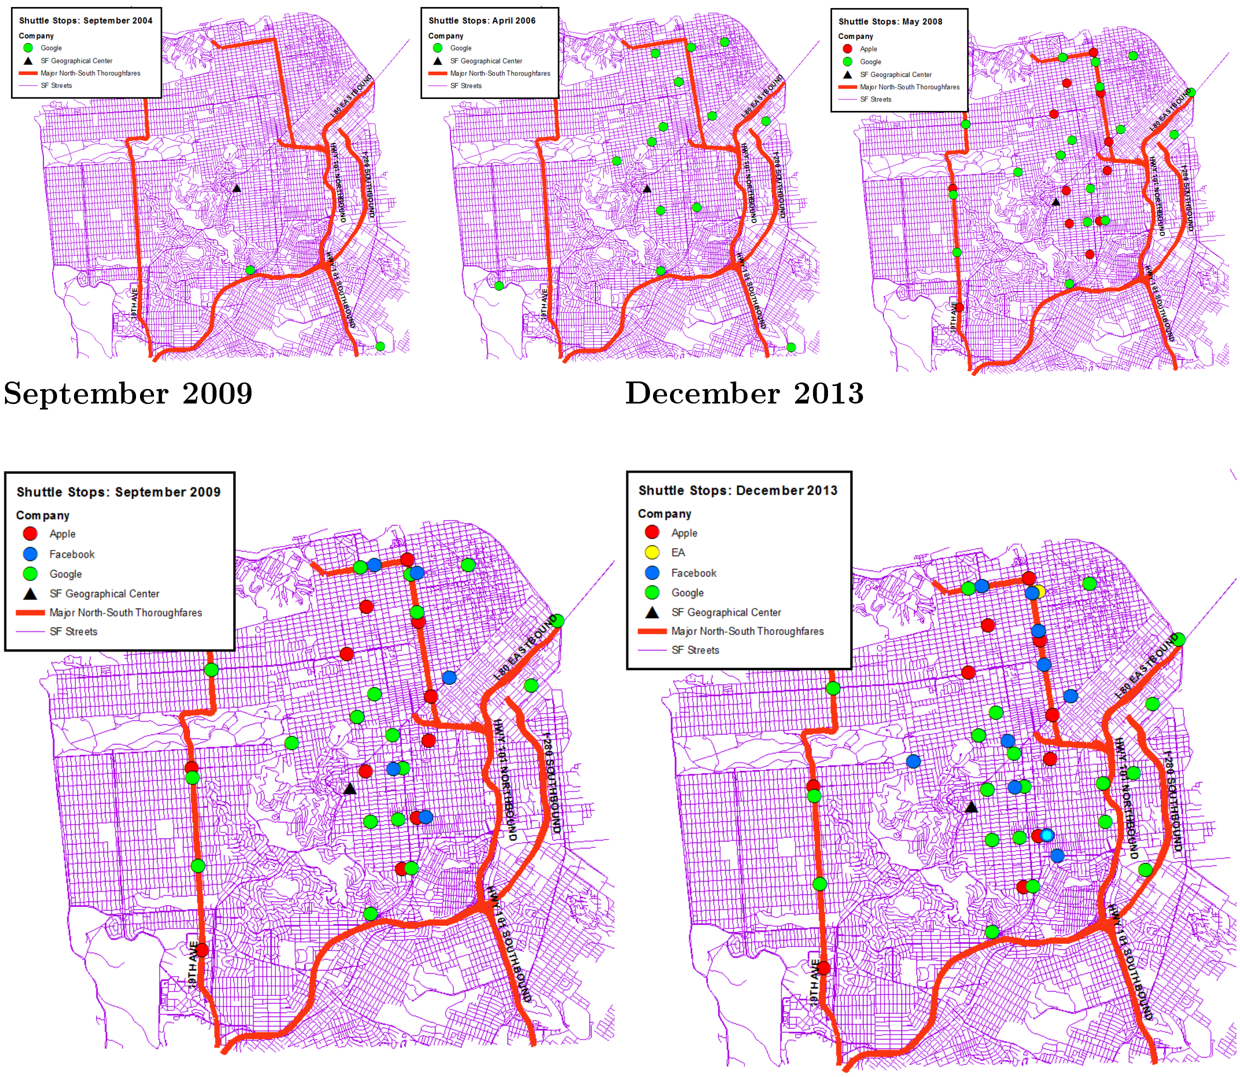

Figure 2.2: Overall Spatial Distribution of Condominiums

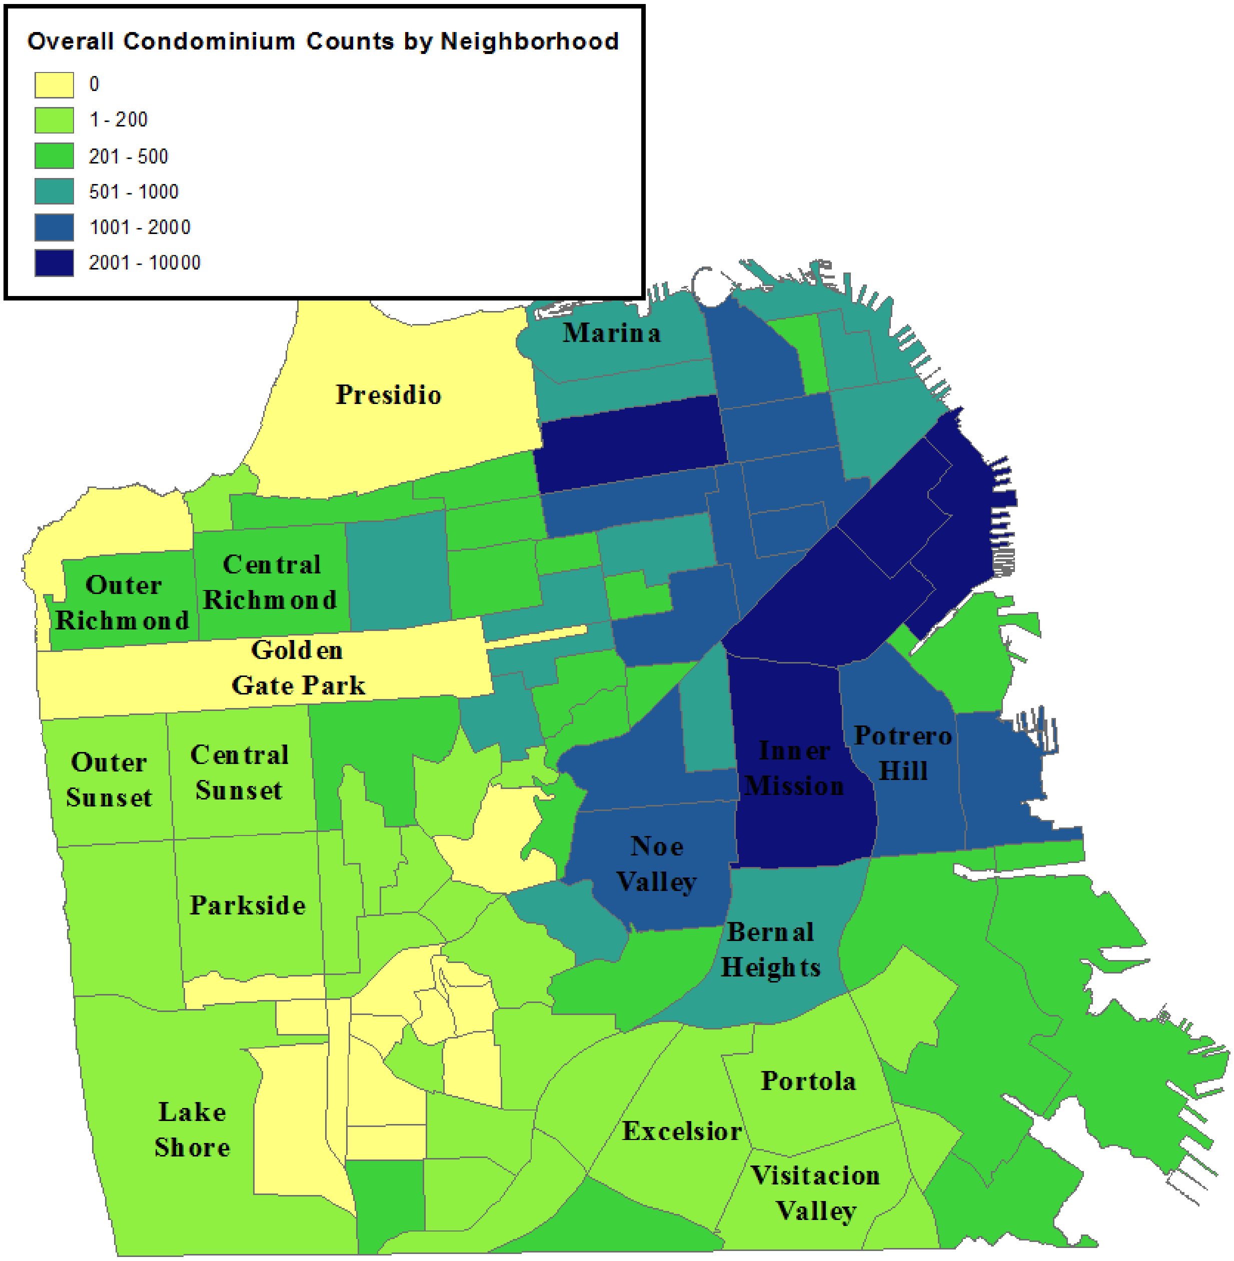

Figure 2.3: San Francisco Transit Networks as of September 2004

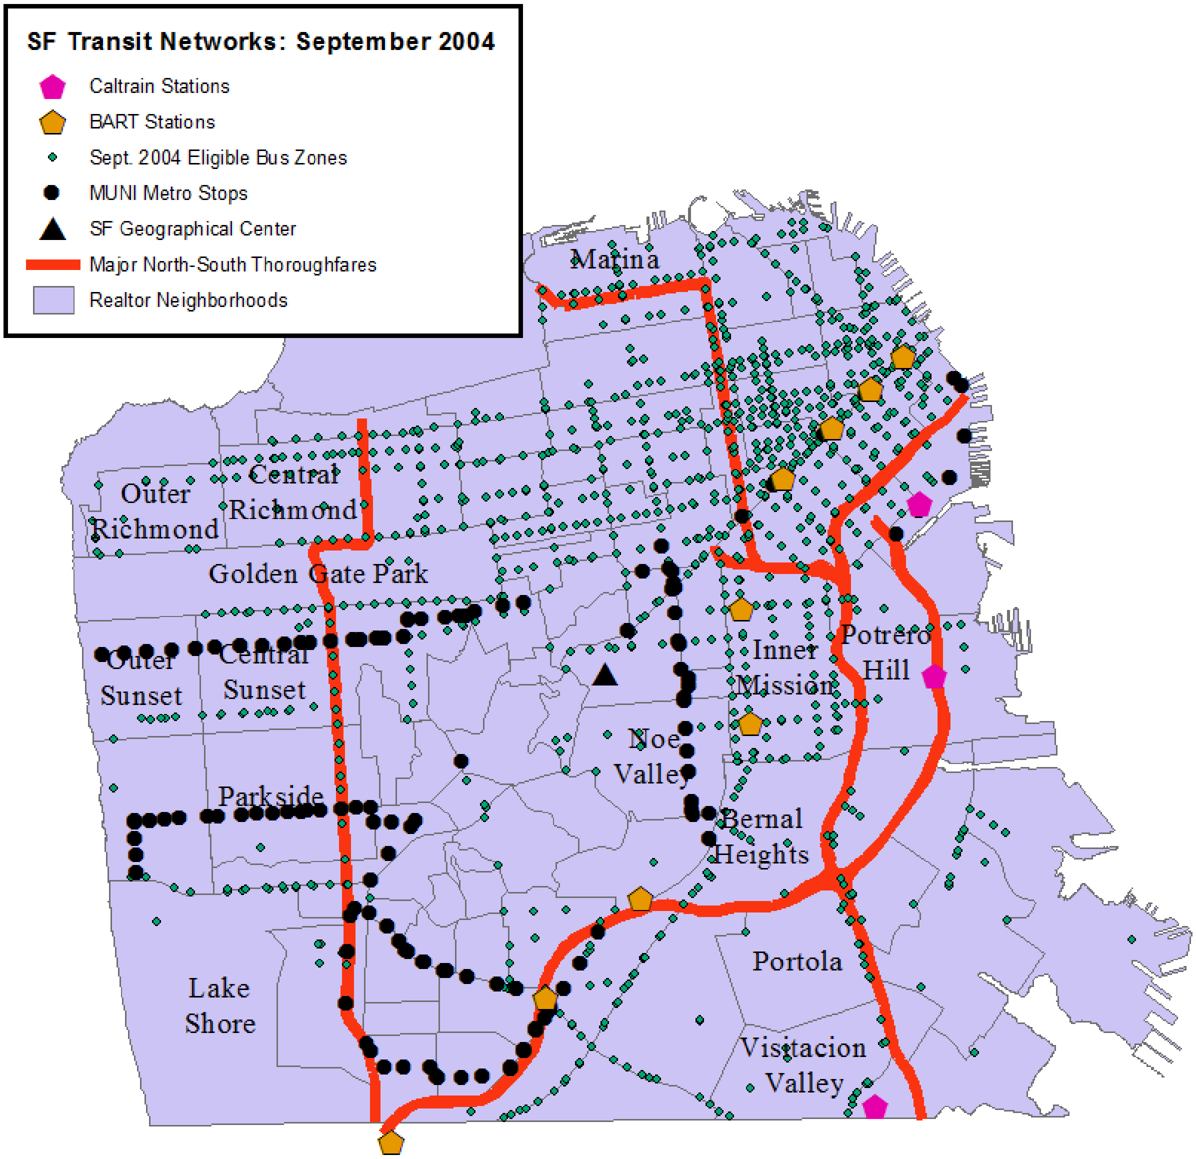

Figure 2.4: Non-Condominium Residential Buildings Rent Controlled By Neighborhood

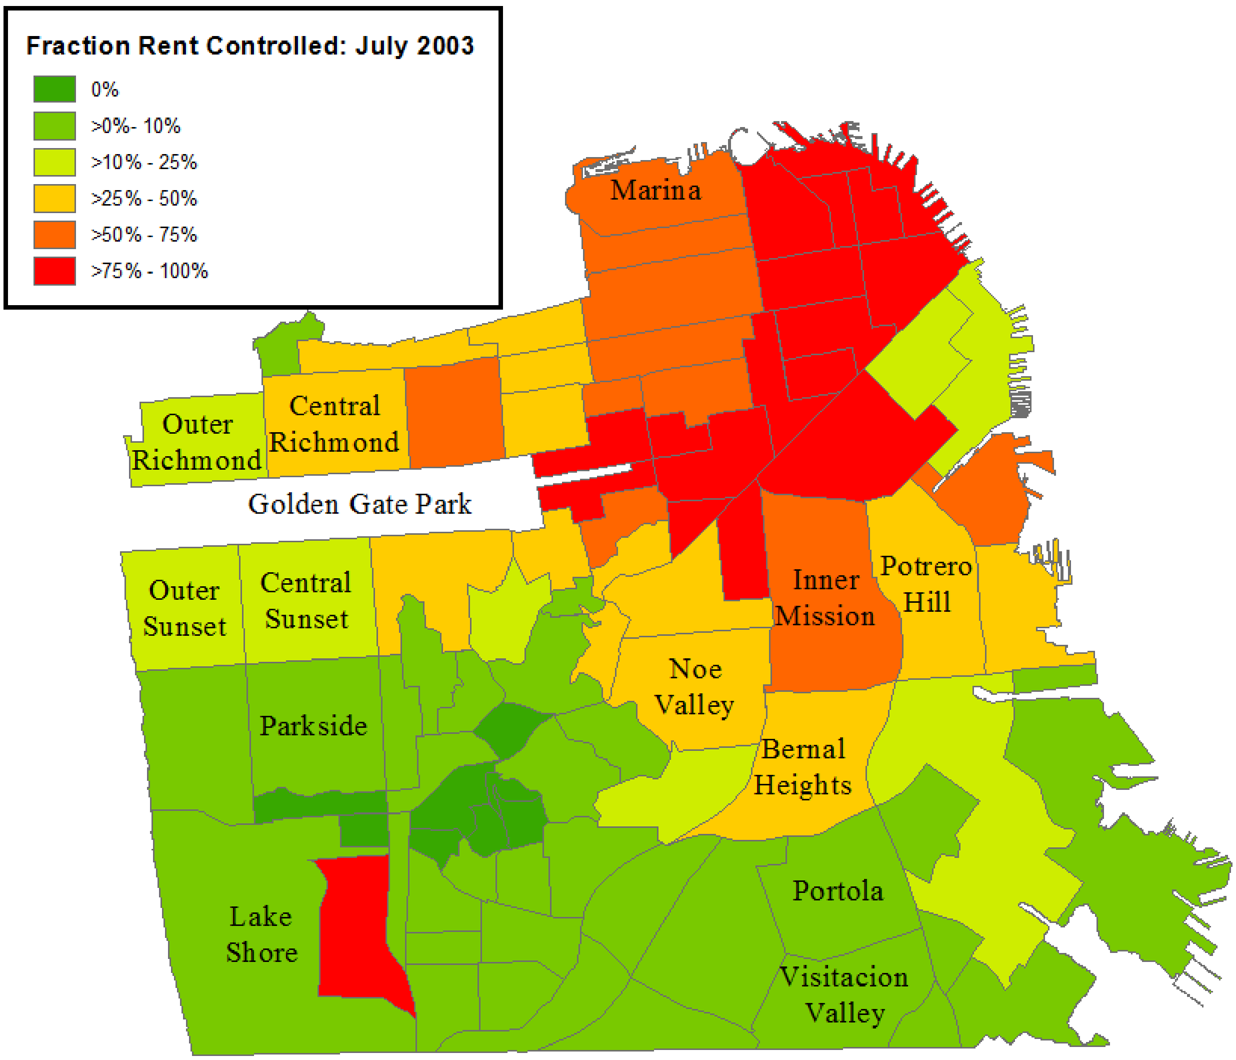

Figure 2.5: Growth of Shuttle Stops Along Thoroughfare-Adjacent Bus Zones, December 2004-December 2013.

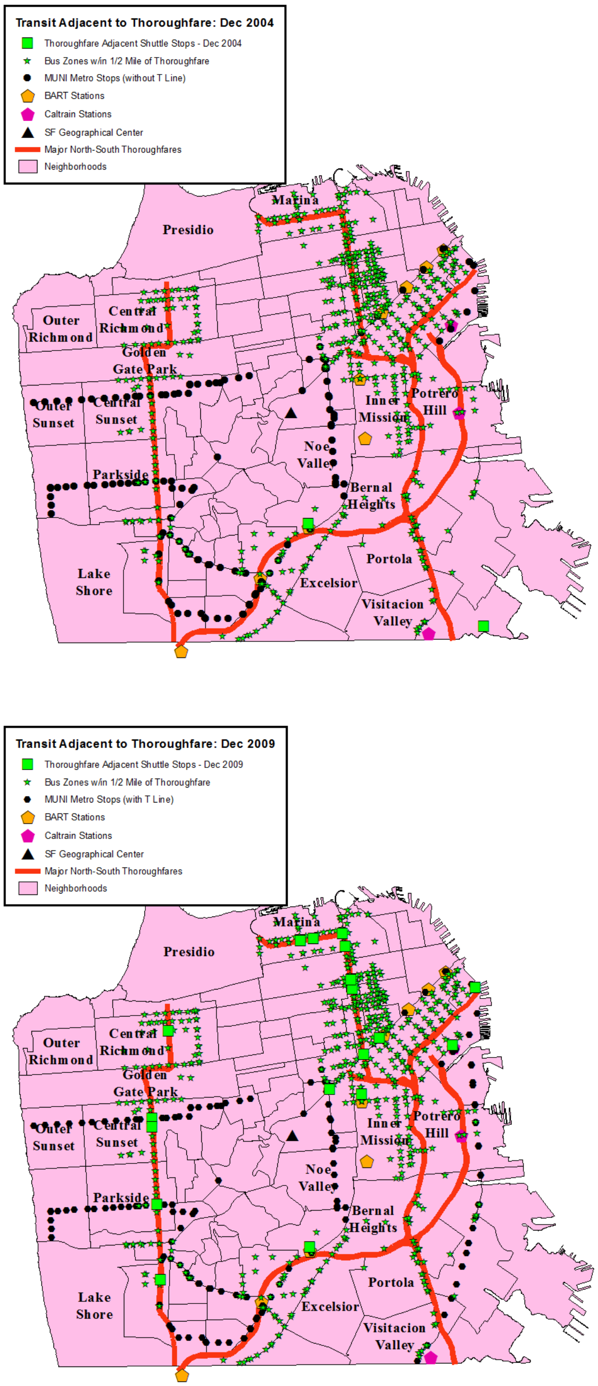

Figure 2.6: Growth of Shuttle Stops Along MUNI Metro-Adjacent Bus Zones, December 2004-December 2013.

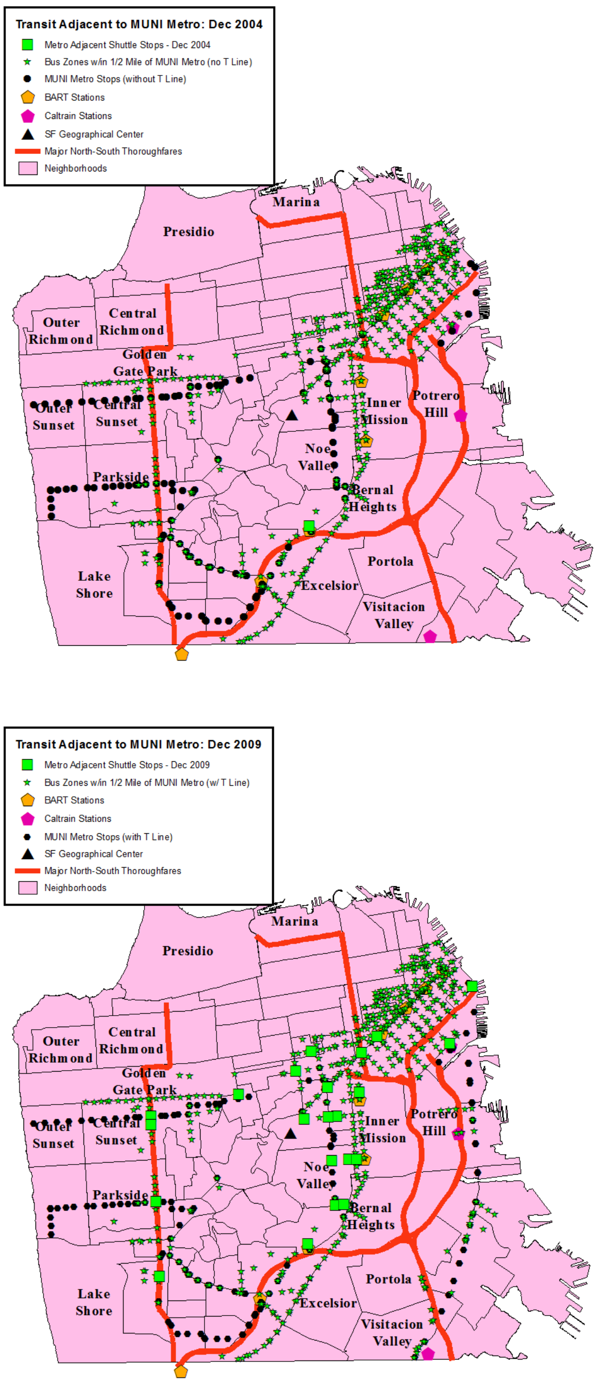

Chapter 3

Grandchildren and Grandparents' Labor Force Attachment

Introduction

Fears of a looming workforce population contraction have been partially allayed by rising labor force participation among older workers. The rise in labor force attachment in this group since the early 1990's (from 31% to 40%) has been ascribed to changes in educational attainment, retirement incentives, and improvements in life expectancy, among other things (Maestas and Zissimopoulos [82]). However, labor force attachment gains seen between the early 1990's have stalled out since the Great Recession: the labor force participation rate for workers 55 and older has remained essentially unchanged at 40% since 2009.[1] While there are many causes, researchers have largely overlooked the role of the changes in family size and composition that have occurred across the industrialized world since the Sexual Revolution.

[1] U.S. Bureau of Labor Statistics, Civilian Labor Force Participation Rate: 55 years and over [LNS11324230], retrieved from FRED, Federal Reserve Bank of St. Louis; https://fred.stlouisfed.org/series/LNS11324230, April 24, 2017.

This curious omission is in spite of the broad reach of grandparenthood: according to a Pew survey, 51% of people aged 50-64 have grandchildren (Taylor et al. [113]). Grandchildren could be playing an influential, though heretofore understudied, role in older workers' labor supply and retirement-timing decisions.

In this study, I address three questions about grandchildren's impact on grandparents. The first is whether grandchildren's presence changes grandparents' labor force attachment. The second is which labor market activities change and which grandparent characteristics are most sensitive to grandchildren's presence. If grandparents change their behavior immediately, this could be a labor substitution effect with the new parents. If instead, the impact is greatest at age 62 or later, then the grandchildren effect is more likely about consumption/leisure tradeoffs around retirement timing. To answer these first two questions, I use the Panel Study of Income Dynamics (PSID), an intergenerational extended survey of US families, to empirically test how individuals respond to grandchildren. Previous studies have shown that grandmothers in particular give their daughters time transfers when a newborn arrives (Johnsen [68]; Compton and Pollak [36]; Posadas and Vidal Fernandes [98], among others), but unanswered is whether these time transfers impact the grandmothers' (or grandfathers') labor supply itself. The outcomes of interest are the propensity to be retired, the propensity to be in the labor force, and the annual number of hours worked were chosen as the results most relevant to policymakers. In the PSID data, women are not asked about their labor force status until several years after the survey started, so for consistency, I use instead an indicator for whether they reported non-zero working hours in the previous year.

The third question is to what extent changes in older worker's labor force participation (LFP) rates are a product of trends in grandparenthood. This third question is tested using Current Population Survey (CPS) data, which is used to generate estimates of national- and state-level labor force participation rates by age. I use Centers for Disease Control and Prevention (CDC), Vital Statistics, Health and Retirement Study, and Retirement History

Longitudinal Survey data to create grandparenthood measures that are merged on to the CPS data to test how LFP rates have responded to changes in extended families.

A panel fixed effects model with the survey data might yield biased results from endogeneity between grandparents' labor force characteristics and fertility timing. For example, some potential new parents may wait to have children when they judge their own parents are most able to help, biasing upward estimates of how much help grandparents give. Endogeneity concerns are addressed in this paper via an instrumental variables approach grounded in a long literature: the repeal of barriers to women's access to reproductive technologies in the 1960's and 1970's. For the CPS labor force participation trends regressions, I add variation in statewide Prohibition laws to motivate the instrument for the oldest in-sample grandparents (born 1892-1895). Laws or court rulings enabling women to freely access either oral contraceptives or abortion were enacted in waves from between roughly 1960 to 1976, with some additional legislation occurring thereafter adjudicating access for minors. Exploiting the variation in the timing and age of medical consent and the ability to access alcohol across states, I instrument for the grandchildren numbers and timing to arrive at unbiased, consistent estimates for how grandparents labor supply changes in response to grandchildren. Key strengths of this empirical strategy are that the drivers of the instrument are well-documented; state-level policy changes are clearly exogenous to micro-level labor supply decisions; and it also controls for both total fertility and fertility timing so that changes through both channels are identified.

I find that both grandparents and grandmothers decrease their labor force attachment in response to grandchildren. Grandfathers are 19.4% more likely to be retired and work 363 fewer hours a year, while grandmothers are 8.5% more likely to be retired and 13% less likely to work any hours in the previous year. When factoring in retirement eligibility, grandfathers eligible for early retirement work 553 fewer hours a year and are 21.5% more likely to be out of the labor force than the grandchildless. Grandmothers are 41.7% more likely to be retired between age 62 and their normal retirement age than those without grandchildren. The number of grandchildren also matters, so that the marginal grandchild causes grandfathers to increase their retirement probability by 17.8%, work 219 fewer hours a year, and be 9% less likely to be in the labor force. For grandmothers, these numbers are 21.7% more likely to be retired, work 432 fewer hours a year, and be 19.1% less likely to report non-zero working hours.

There is also evidence of a diminishing marginal grandchild response for grandfathers. When factoring in retirement eligibility, the marginal grandchild makes grandfathers 27.0% more likely to be retired, but this shrinks to 15.9% more likely when early retirement eligible, and further still to 5.4% when full retirement eligible. Similarly, the marginal grandchild also has a diminishing effect on grandfathers' annual hours worked and propensity to be in the labor fore as they advance to early retirement and then to full retirement eligibility. Grandmothers, on the other hand, have a growing marginal grandchildren effect, where a marginal grandchild makes them 24.4% more likely to be retired before retirement eligible, and which rises to 33.5% when the marginal grandchild arrives during early retirement and 32.2% during full retirement, with a similar effect seen for annual hours worked and propensity to report non-zero hours worked.

The individual level results point to a differential grandparent response over non-grandparents, but the deduction that grandchildren might have played a roll in the fall and rise in labor force participation of older men between the 1960's to the Great Recession is not supported by the model. I find that whether grandchildren change older men's LFP in the aggregate is sensitive to the specification of birth cohort controls, but that there is robust evidence for economically meaningful interaction effects between grandparenthood, retirement eligibility, and Social Security benefit levels. Each additional average grandchild when grandfathers are early retirement eligible lowers their national labor force participation rate by 0.51 points (assuming average values of other controls). If the penalty for taking early retirement were

$100 smaller, then an additional average grandchild would lower the LFP rate by 0.56 points, so that for each $100 reduction in the penalty, the growth in the extended family by another grandchild lowers the LFP rate by 0.05 points. In general, my results indicate that grandfathers do not differentially retire at higher rates when they are early retirement eligible over the grandchildless, but that they are more likely to respond to reductions in the early retirement penalty if they have grandchildren.

Grandfathers, however, are very responsive in the aggregate to taking full retirement over their grandchildless peers. All specifications of both grandfather measures and controls agree that grandfathers exit the labor force at full retirement at higher rates: each 1 point increase in the fraction grandparent at the full retirement age yields a 0.32 point decline in the LFP rate and a 0.29 drop at the early retirement age (assuming average values of the other controls). With a one child increase in the average number of grandchildren, this becomes a 2.94 point drop. Interactions with the primary insurance amount (PIA) received upon reaching full retirement age and the delayed retirement credit (DRC) are also statistically significant, but the coefficients are so small that this variation only accounts for a fractional part of the trends in LFP rates observed since 1962. Simulations of ultra-high fertility (permanent Baby Boom), medium high fertility (no *Roe v Wade*), medium low-fertility (no Baby Boom), and ultra-low fertility (permanent post-Great Recession baby bust), show that while changes in grandparenthood do not explain the 1970-2009 fall and rise in the LFP rate, the fall would have been even steeper because the Baby Boom partially masked that the latent labor force participation of older male workers during this time would have been 12 points or about 25% higher in the early-to-mid 1960's than what was actually observed.

The paper proceeds as follows: Section I reviews the existing literature on grandparents, their labor supply, and trends in both phenomena and how they may be interrelated. Section II describes the PSID and its grandparent samples, the CPS, and other data sources used to estimate national trends in LFP and grandparenthood. The research design and empirical approach for the individual-level PSID estimations and the results of those estimations are discussed in Section III. Section IV discusses the empirical approach and results for the national-level trends chiefly from the CPS data. Section V has results of counterfactual simulations on various alternative grandparenting trend scenarios to illustrate the overall significance of grandchildren. Section VI presents various robustness checks on the results from Sections III and IV and Section VII concludes.

3.1 The Case for Grandparenthood's Effect on Labor Force Participation

3.1.1 Trends in Labor Force Participation Among Older Workers

The key research question is whether being a grandparent changes older worker's labor force attachment and what the economy-wide implications are. Policymakers throughout the developed world are interested in what drives older worker's participation because population aging has put pressure on the solvency of social insurance systems (Organisation for Economic Development and Cooperation [97]). In the United States as of 2015, 24% ($888 billion) of the federal budget goes to Social Secuirty alone, with another 17% ($546 billion) going to Medicare, the old age health insurance program for people 65 and older (Center for Budget and Policy Priorities [32]). Thus, collectively, spending on retirees is now over 40% of the federal budget, so that the future of federal expenditures is sensitive to trends in labor force participation among current and future retirees.

Recent trends offer some room for optimism. Figure 3.1 shows the national trends in labor force participation among workers 55 and older. Since World War II, LFP among workers 55 and older steadily declined from 1948 to 1970 from 43.3% to 39.0% (-0.2% per year), before more steeply dropping off between 1970 to 1987 from 39% to 30% (-0.5% per year).[2] However, since the late 1980's and early 1990's, LFP for older workers reversed and rose until the Great Recession (from 30.1% in 1994 to 40.0% in 2009), and essentially leveled off at about 40% until the present.

Nonetheless, while participation has recovered from its early 1990's lows, LFP in this age group was higher even as late as 1960. This is true in spite of the fact that life expectancy at age 65 was 5 years higher in 2014 (79.3 versus 84.3 years, National Center for Health Statistics [48]). Further, existing evidence indicates that older workers are as healthy or healthier now than they were 50 years ago. The fraction of adults aged 55-64 and 65 and over who smoke is one-third of what it was in 1965 (National Center for Health Statistics [48]). The fraction of adults ages 40-59 reporting a work-limiting health condition or a disability has been roughly stable since 1988 (Autor [9]), and the rate of adults claiming disability insurance for heart disease and cancer declined between 1983 and 2003 (Autor and Duggan [11]).

Several papers have tried to explain these shifts in LFP, but the paper closest to this one is Blau and Goodstein [22], who explore a variety of factors to explain the post-war fall and rise in labor force participation among older workers. While they ascribe the post-1990 rise largely to greater educational attainment and reduced Social Security-generosity, reasons why labor force participation fell remained unaccounted for. Ultimately, the problem, as stated by Blau and Goodstein [22] (p. 356), remains:

> "Two key points remain unresolved by the findings reported here: what caused the long decline in LFP among older men and why is Social Security more important

[2]Policymakers responded in 1977 and 1983 by decreasing the generosity of Social Security with seemingly little impact on labor force participation according to Kreuger and Pischke [76].

> in accounting for recent LFP increases than in explaining the previous decline? The first question has been studied for many years without much success, and unfortunately, our results do not suggest any new avenues of research."

While Blau and Goodstein [22] gave a comprehensive overview of factors that drive labor force attachment, several other papers have also studied what drives the changes in older worker's participation over the past 70 years that either affirmed their findings or explored motivators that were beyond the scope of their study. The rise since the 1980's has been ascribed partly to a society-wide shift from defined benefit to defined contribution retirement plans (Hurd and Rohwedder [65]; Heiland and Li [59]), changes in Social Security rules (Behaghel and Blau [18]; Blau and Goodstein [22]; Gustman and Steinmeier [55]; Hurd and Rohwedder [65]), trends in technical skill accumulation among older workers (Burlon and Vilalta-Bufi [24]), gains in educational attainment in successive birth cohorts (Burtless [29]; Maestas and Zissimopoulos [82]), and rising female labor force participation causing men coordinate retirement timing with their wives (Schirle [105]; Gustman and Steinmeier [54]).[3]

While all of these factors have some explanatory power, there are still unexplained elements of the participation time series. The nearly 40 year fall in labor force participation among older workers seems to have progressed independently of both increases and decreases in Social Security generosity (Kreuger and Pischke [76]; Blau and Goodstein [22]). Nor can rising educational attainment explain the four decade fall, because attainment increased steadily during this time period. Mean years of schooling for native-born workers was at about 9.5 for those turning 62 in 1970 (born in 1908), and for those turning 62 in 1994 (born in 1932), this had increased to nearly 11.5, at an implied rate of 0.08 years of schooling per birth cohort (Goldin and Katz [52]). One major trend that has not been explored to answer this question is what role changes in grandparenthood have played in the fall and rise of older worker's LFP.

[3] Wives are often younger than husbands, so a preference for a joint retirement would prompt men to delay retirement until their wives were also eligible.

The impact of grandchildren is potentially very broad: according to a Pew survey, 51% of people aged 50-64 have grandchildren (Taylor et al. [113]). The arrival of grandchildren could play an influential, though heretofore understudied, role in older workers' labor supply and retirement-timing decisions. These dynamics are particularly relevant for entitlements, where a new grandchild represents a future contributor recipient while simultaneously changing the labor supply of a current taxpayer and imminent beneficiary. Thus, for policymakers wanting to motivate increased labor force participation by older workers, a clearer understanding of how new generations change the behavior of existing ones can help model better public policies. The pull of grandchildren is very real, but the demographic transition in total fertility and fertility timing is simultaneously changing the nature of grandparenthood. Older worker's labor force attachment started rising 20 years after the fertility transition of the late 1960's and early 1970's and then plateaued at a higher level about 18 years later.

3.1.2 Post-War Grandparenthood

Over the same post-WWII period, grandparenthood has risen and fallen with the national birthrate. Figure 3.2 shows that births fell through 1920's and into the Great Depression, before rising after World War II and spiking in the late 1950's as part of the Baby Boom. Thereafter, the birthrate flattened in the 1970's and except for minor fluctuations over the intervening years, has largely hovered around 65-75 births per 1,000 women aged 15-44.

Not only are the implied number of grandchildren changing over time, but when people become grandparents is, too. Couples throughout the world are choosing to have fewer children and to have them later (Morgan [86]; Bloom et al. [23]; Caldwell [31]). For context, Taylor et al. [112] found that in the United States in 1990, a greater share of births were to teenagers than to women 35 and older, but by 2008, the reverse was true. The Centers for Disease Control and Prevention and other agencies do not track grandparenthood, but published birth and marriage data give strong clues as to how older workers' families are evolving.

First, the age of first marriage declined for men and women between 1890 (26.1 and 22) and 1949 (22.7 and 20.5), flattened out around 23 for men and 20.5-21 for women, before starting to rise steadily from about 1975. Figure 3.3 shows these trends, and that median age of first marriage recently overtook 29 for men and 27 for women. These shifts are significant because the married fertility rate has always been higher than the unmarried fertility rate, and that until about 1970, over 90% of all births were to married women (Kendall and Tamura [72]). Marital shifts in turn impacted both the age of first birth and the fraction of women remaining childless. The median women born in 1910 was first married at 22 (in 1932) and had her first baby at 23 (in 1933), and about 20% reached 45 (in 1955) childless. In contrast, the next generation of women born in 1935 was first married at 21 (in 1956), had her first baby at 22 (in 1957), and only 11.4% reached 45 (in 1980) childless (Kirmeyer and Hamilton [73]). Since the average man married a woman roughly three years younger than himself between 1920 and 1940, these figures can be extrapolated to imply that the average man born in 1907 had *at least* a 20% childless rate versus a man born in 1932 who had a roughly 12% childless rate.

The next generation of men and women show a different pattern: the median woman born in 1960 was first married at 22-23 (in 1982-1983), and had her first baby at 25 (in 1985), and about 15.6% reached 45 (in 2005) childless (Kirmeyer and Hamilton [73]). Comparing these statistics to the LFP rates in Figure 3.1, a man born in 1907 would turn 55 in 1962, when the LFP rate for those 55 and older is 40%, but a man born in 1932 turned 55 in 1987, when the corresponding LFP rate was 30%. When the likely partners of the 1960 cohort (born themselves in 1958) reached 55 (in 2013), the LFP rate was back up to 40%. However, these rough statistics cannot accurately convey either what fraction of older workers were actually grandchildless and what the joint distribution of grandchildlessness and labor force participation was, but these numbers help motivate the potential empirical connection between grandchildren's presence and labor force attachment.

Figures 3.4a and 3.4b plot the LFP rates by selected age groups against the fraction that are grandparents in each age group by year and the average number of grandchildren in each age group by year, respectively.[4] The figures here do not show a particularly tight link between grandparenthood trends and the labor force participation for those 50-61, partly because it is not possible to look at LFP trends in these age brackets before and during the Baby Boom. Neither, however, do they contradict the idea that there might be some causal connection. However, for those 65-69, grandchildren peak just before LFP rates in those groups reaches its nadir. While the alignment in trends is not exact, the graphs strongly suggest that at least for workers 65 and older, and possibly for those 62-64, some relationship might exist between grandchildren trends and their labor force attachment.

3.1.3 The Literature on Grandparents and Grandchildren

While the past 70 years has witnessed large shifts in grandparenthood and older workers' labor force activity, the labor economics literature on grandparenthood is thin. Existing work offers mixed indications on what kind of labor market response by grandparents is most likely. Ho [60] found using Health and Retirement Survey (HRS) data that grandparent responses seem to vary according to their marital status and financial resources. Grandparents are most likely to help with newborns, and grandparents living in close proximity provide larger time transfers. Married grandparents are both more likely to be employed and to give financial help, although to what extent that is due to married couples having more resources or being able to provide both time and financial assistance to the new parents is unclear. In

[4] The grandparent statistics reported here were generated from the methods described in Appendix B.2.

comparison, single grandparents made no time or financial adjustments in response to new grandchildren. Because the study did not attempt to instrument for the adult children's fertility, it is hard to know which of these results might be significant when endogeneity bias is removed from the estimates.

Most other studies have focused just on questions of time transfers. In part, this reflects what grandparents desire themselves. A Pew Survey (Taylor et al [113]) reported that spending time with grandchildren is what the elderly most value about getting older. Hochman and Lewin-Epstein [61] found from survey data of elderly Europeans that grandparents are more likely to report a desire to retire early. This result was higher in countries that have less generous public childcare policies, suggesting that grandparents do respond to the childcare needs of their children by decreasing labor force attachment. The intuition behind the second finding is confirmed in Compton and Pollak [36], Posadas and Vidal-Fernandez [98], Aparicio-Fenoli and Vidal-Fernandez [4] who find that grandmothers providing childcare for new parents increases the mother's labor supply. This help may not make much difference in the grandmother's own labor supply, as Whelan [116] found that as long as the grandmother's help was for less than 12 hours a week, labor supply was not affected.

However, these studies do not account for the possibility that fertility timing and grandparent labor force characteristics are jointly determined. Namely, adult children's fertility decisions may be based on the likelihood they will receive grandparent assistance. If adult children believe they will need assistance, they could time their childbearing to correspond to the grandparent's ability to help. This possibility could bias estimates of the grandparents' labor response, because it is then unclear if the arrival of grandchildren causes a change in grandparent's behavior or the grandparent's willingness to provide financial or childcare help influences adult children's decision on when to have their own children.

Two previous studies attempt to address the endogeneity bias by using instrumental variables to estimate how grandchildren affect grandparent labor force attachment. First, Wang and Marcotte [115] use PSID survey data study how grandparents who are raising grandchildren change their labor force behavior when the grandchildren move in. Their interest is chiefly in comparing three-generation versus skipped-generation households, so their instrument includes the existence and number of grandchildren.[5] They find that compared to independent-living grandparents, grandparents co-residing with grandchildren are more likely to increase their labor force participation. However, the narrowness of the research question means that it is of limited use for understanding the relationship between grandchildren and grandparents labor force attachment, as only about 7% of grandchildren live in a grandparent-headed household according to the Population Reference Bureau.[6]

The second and more comparable study is a working paper by Rupert and Zanella [101], which estimates the impact of the first grandchild on grandparents also using the PSID. Their study finds that becoming a grandparent reduces the annual number of hours worked for grandmothers by at least 170 hours, but no significant effect was found for grandfathers. Rupert and Zanella instrument for arrival of the first grandchild by exploiting variation in the sex of oldest adult child of the grandparents. Their empirical strategy rests on the fact that on average women marry and bear children at younger ages than men, meaning that parents of adult daughters will exogenously be more likely to become grandparents at younger ages than parents of adult sons.

Their paper is informative but has several shortcomings that are addressed here. The first is that the authors eschew the PSID's sampling weights, arguing that conditioning on the covariates that the sampling weights account for is preferable to weighting. This approach introduces two problems: conditioning annual labor supply on annual income risks introducing endogeneity bias, and the oversample of low-income households was done on additional

[5]The rest of their excluded instruments are state-level characteristics: teenage pregnancy and incarceration rates plus the generosity of state kinship foster care arrangements.

[6]Paola Scommengna, Population Reference Bureau, "More U.S. Children Raised by Grandparents",`http://www.prb.org/Publications/Articles/2012/US-children-grandparents.aspx`, last accessed January 30, 2016. The PSID likely has a substantial subsample of these families due to the low-income OEO oversample that was originally included.

characteristics, such as race and location. Their results thus risk introducing selection bias on properties not accounted for in the covariates but are otherwise accounted for in the sampling weights. This study accordingly uses the PSID's sampling weights.

The second shortcoming is that when looking at the impact of a marginal grandchild (*i.e.*, the impact of each additional grandchild), they hold that the endogenous decision is to become a first-time parent, but subsequent children and siblings' fertility are both exogenous. However, there is no justification given for this assumption. Their paper finds that conventional estimates understate grandparent labor market changes, but their subsequent finding that additional grandchildren increase labor force participation possibly have endogeneity bias. This study explicitly accounts for endogenous fertility of all grandchildren regardless of birth order.

Lastly, their instrument's validity with respect to the exclusion criterion is undetermined. As they openly acknowledge, the literature is inconclusive on whether it can be assumed that the sex of the first-born child exerts no impacts on the parents' labor supply. They run several empirical tests to support the instrument's validity but due to the sampling issue discussed above, it is not clear that the matter is settled. Thus, it is a clear innovation to use instead state policies which more clearly satisfy the exclusion criterion.

3.2 Data Description

The sample of grandparents and their families is drawn from the PSID, a dataset that follows about 4,800 households initially sampled in 1968 and their lineal descendants. The original sample is composed of two subsamples: a nationally representative sample of 2,930 families (called the SRC Sample) and an oversample of 1,872 low-income families (the SEO Sample).[7]

[7] In 1990, the sample was updated to include 2,000 post-1968 immigrant families (exclusively of Latino origin), but they were dropped in 1995. In 1997, the sample was again refreshed by adding 500 post-1968

The PSID follows the family members of the original sample households as they move out, marry, and form families of their own, resulting in about 70,000 individuals appearing in at least one survey. This survey design makes the PSID a uniquely rich source of information on intergenerational dynamics, especially because the PSID supplements the main survey with auxiliary datasets on marriage and childbirth histories. Between 1968 and 1997, the survey was conducted annually, and from 1999 to the present has been conducted biennially.

The PSID makes available a series of files that enable identification of all surveyed descendants of a given individual through their Family Identification Mapping System (FIMS). Using the FIMS, I have identified the adult children and grandchildren of each grandparent, and then merge on the survey responses of each respondent. My panel has 2,373 grandmothers and 1,712 grandfathers across 38 survey years. Location and age information in the PSID also allows me to code with a high level of precision the likely abortion and contraception access status that the female respondents had. For a complete overview of how abortion and contraception access was encoded, see Appendix B.1.[8] In addition to observing demographic characteristics, such as marital status, age, race, and educational attainment, the dataset also measures respondent's key labor market characteristics: retirement status, annual hours worked, and labor force and employment status.

Individuals who were between the ages of 22 and 54 in 1968 and were the current or future parents of at least one child were chosen as the sample of potential grandparents. Being aged 22 in 1968 as the minimum age cutoff was chosen to minimize confounding variation between education and labor force characteristics. In 1968, the vast majority of adults had at most a college education, so almost all in-sample individuals would have completed their educations and moved into the workforce. The maximum 1968 age of 54 was chosen because

immigrant families. Because the instrument is dependent on the individual being observed between 1968-1980, these families are not included in this study.

[8]As detailed below, most states regulated access on the basis of age, but a few did so on the basis of educational attainment (minor HS graduates can buy contraceptives in Alabama and Pennsylvania) or marital status (Alabama, Florida, Maine, Missouri, New Jersey, Texas, West Virginia). Coverage can thus be ascertained with a high degree of accuracy in the PSID that other studies might overlook.

it permits me to observe most individuals' labor force participation before they retire. The only other condition put on the sample was to exclude observations from Kansas from the grandfather analyses because only 3 grandfathers were initially sampled.

Table 3.1 shows selected summary statistics on grandfathers and grandmothers. Individuals in the sample were observed to become first-time grandparents in their late 40's,[9] and then retired about 10 years thereafter. Differences in mean ages between grandfathers and grandmothers reflect that families were usually sampled as a household, so that the age gap between husbands and wives got "passed through" into the sample.

Descriptive data from the PSID on grandparent-to-adult child time transfers is presented in Table 3.2 courtesy of the PSID's 2013 Family Rosters and Transfers module. Adult children with their own children received on average about 25 more hours in time transfers a year from both sets of grandparents than childless households. The grandparent's marital status and sex matters, as does the sex of the adult child. Married grandparents are more time-generous than unmarried grandparents, and the mother's parents are more generous than the father's. In almost all cases, except for single grandfathers, potential or actual grandparents indeed give more time transfers to adult children with their own children than those without.

For national-level labor force participation trends, I use March Current Population Survey (CPS) micro-data to create a synthetic panel dataset, and supplement it with data drawn from Social Security Administration (SSA). As in Blau and Goodstein [22], I aggregate individual-level records on men aged 55-69 from the CPS into cells defined by year, birth year, and Census Division. I then supplement it with men aged 50-54 to provide more data on the impact of grandparenthood on the labor force attachment in this cohort. The resulting panel covers 74 birth cohorts (1892-1965) between 1962 to 2015.

For each birth cohort, I calculated the fraction who were grandparents and their average

[9] Adult children who were not living with the Head and Wife of household in 1968 are not consistently surveyed by the PSID, so this statistic is biased upwards somewhat.

number of grandchildren at both the birth cohort-age-education group and birth cohort-age-state level using Health and Retirement Study and Retirement History Longitudinal Survey data. I was not able to use PSID data to estimate grandparenthood measures for the national-level dataset, primarily because it is not a large enough sample of older individuals to generate credible grandparent statistics at the birth cohort level. Instead, I combined two data sources that oversample older individuals longitudinally to estimate this fraction. The first is the Health and Retirement Study (HRS) data which sampled roughly 20,000 older individuals in successive birth cohorts from 1992 to 2014. The second is the Retirement History Longitudinal Survey (RHLS), the predecessor of the HRS, which sample 11,000 plus individuals chiefly born between 1906 and 1911 biennially from 1969 to 1979. Unfortunately, only the 1975, 1977, and 1979 questionnaires asked about the number of living grandchildren but the two datasets combined provide important evidence on the evolution of grandparenthood over time. Appendix B.2 has more detail on how this measure was constructed by using the data points to estimate the fraction grandparents and their average number of children for the two panels.

I then use Blau and Goodstein's method to create simulated work lifetime earnings histories and use these to generate expected Social Security old age and disability benefits payments for either retiring at ages 62, 65, and 70, or dropping out of the labor force and claiming disability payments from ages 50-64. More detail on how these were performed can be found in Appendix B.3.

3.3 Individual-Level Estimation with the PSID

In this section, I test whether and how grandchildren alter grandparents' behavior by means of a fixed effects panel regression. Regressions for grandmothers and grandfathers are estimated separately. The left-hand side variable is the grandparent's labor market outcome:

retirement status, annual hours worked, labor force status (grandfathers) and non-zero hours reported (grandmothers).[10]

The descriptive evidence suggests several ways changes in grandparenthood may be altering labor force participation. Fewer people become grandparents and they do so later than in the mid-20th Century, and that effect may account for greater labor force participation among older workers. Thus, I create an indicator to test for the impact of being a grandparent on the outcomes of interest. Another factor is that family sizes have shrunk since the mid-20th Century, so that there may be a total fertility effect that causes people to work longer when their extended families are smaller. I study this by estimating the grandparent response to the number of children that each adult child has. The last channel is whether a grandchild effect exists differentially by grandparent age. If the grandchild channel has the strongest impact on the retirement decision, then there should be only a small effect for workers who are not eligible yet for Social Security's early or full retirements.[11] This effect would also explain why there would be a "lag" between when people become grandparents (typically in their early 50's) and an effect a decade or more later. I thus interact grandparenthood status with early and full retirement eligibility indicators.

[10] Labor force status is not reported for wives in every year in the PSID, so an indicator for whether the grandmother reported some working hours is used as a stand-in. Compared to a measure of being in the labor force, it codes to zero grandmothers who were unemployed and looking for work (and are technically in the labor force), and it will code to 1 grandmothers who report some hours worked, but are students, retired, or homemakers. For those years where labor force status is available (1976 onwards), the correlation between a indicator for being in the labor force and an indicator for reporting non-zero annual workings hours is 0.5468.

[11] Some people do retire earlier, particularly if they have defined benefit pension plan, but most choose to do so when Social Security eligible. See Hurd and Rohwedder [65] in particular for more discussion.

3.3.1 Empirical Strategy for Individual-Level Estimates

The first grandchildren impact channel is whether becoming a grandparent influences labor force attachment. The equation to estimate this channel takes the form of

$$\begin{aligned} Outcome_{gst} = & \beta_0 + \beta_1 \mathbb{1}\{Grandparent_{gst}\} + \beta_2 GPDemVars_{gst} \\ & + \beta_3 ACDemVars_{igst} + \beta_4 Year_t \beta_5 State1968_s \\ & + + \beta_6 (State1968_s * Year_t) + \beta_7 GP_g + u_{gst}. \end{aligned} \tag{3.1}$$

The unit of observation is the grandparent and the key variable of interest is the indicator for grandparent status, $\mathbb{1}\{Grandparent_{gst}\}$, which was created by finding the birth year of the oldest grandchild. $Outcome_{igst}$ is either grandparent g's annual number of hours worked, retirement status, age of retirement, age of death, or whether the grandparent is in the labor force in year t in state s with adult child i. Each regression can thus be thought of the adult child's fertility choice's effect on the grandparent. The right hand side is populated with the demographic information of both the grandparent and the eldest adult child, plus state, year, state-by-year, and grandparent fixed effects. I use only the eldest daughter or daughter-in-law's controls, under the assumption that the eldest grandchild will be born to the eldest non-senior female in the family.

$GPDemVars_{igst}$ is a vector of demographic information about the grandparent, which includes a dummy for whether the father or mother of the adult daughter is eligible for full Social Security benefits; a dummy for whether the head of household became age-eligible for early Social Security benefits;[12] age and age-squared, reflecting that often labor force at-

[12] People can become eligible for partial benefits at 62 as long as they have worked a sufficient number of quarters, but the work requirement is difficult to accurately estimate in the PSID, so this dummy is measured only as a function of age. It's less likely that grandmothers would have been eligible to receive early retirement benefits, in particular, so in their regressions, this is changed to be the head of household's eligibility. The head of household is the husband if present, and the wife or single woman if not.

tachment first rises and then falls with age; and marital status. Time-invariant grandparent characteristics are not included, because the grandparent fixed effects would cause them to drop out. To avoid endogeneity bias between education level and labor supply, educational attainment measures are not included, but by construction, there is very little change in educational attainment in sample.

$ACDemVars_{igst}$ is a vector of the adult child's demographic information. It includes age, marital status, the adult child's sex, and the wife's age if the adult son is married. $State1968_{is}$ and $Year_t$ are vectors of state and year dummies.[13] State fixed effects control for time-invariant characteristics common to all residents who lived in state s in 1968, year fixed effects control for year-specific shocks, and state-by-year fixed effects thus control for state-specific yearly shocks . These could include state-specific employment or economic shocks common to all individuals in a given year that would influence labor force attachment coincident with fertility timing, also affected by economic conditions (Amialchuk [2]; Black et al. [21]; Schaller [104]). Each regression is run separately for grandmothers and grandfathers. Grandparent fixed effects, GP_g, are included to control for unobserved, time-invariant characteristics of grandparents and their relationships with their children.

Since the greatest impact of being a grandparent may be observed when older workers are retirement-eligible, Equation (3.1) can be augmented by interacting the grandparent

[13] As stated in the data description, the state here is the individual's 1968 state. These also fall out of the model when grandparent fixed effects are added.

indicator with indicators for being Social Security-eligible:

$$
\begin{aligned}
&Outcome_{gst} = \beta_0 + \beta_1 \mathbb{1}\{Grandparent_{gst}\} + \beta_2 ACDemVars_{igst} \\
&+ \beta_3 \mathbb{1}\{Grandparent_{gst}\}\mathbb{1}\{EarlySSEligible_{gst}\} \\
&+ \beta_4 \mathbb{1}\{Grandparent_{gst}\}\mathbb{1}\{FullSSEligible_{gst}\} \\
&+ \beta_5 GPDemVars_{gst} + \beta_6 Year_t + \beta_7 State1968_s \\
&+ \beta_8 (State1968_s * Year_t) + \beta_9 GP_g + u_{gst}.
\end{aligned}
\tag{3.2}
$$

Total fertility effects are analyzed with the panel fixed effects model below:

$$
\begin{aligned}
&Outcome_{igst} = \beta_0 + \beta_1 ChildCount_{igst} + \beta_2 GPDemVars_{igst} + \beta_3 ACDemVars_{igst} \\
&+ \beta_4 State1968_{is} + \beta_5 Year_t + \beta_6 (State1968_{is} * Year_t) + \beta_7 GP_g + u_{igst}
\end{aligned}
\tag{3.3}
$$

The key variable of interest is $ChildCount_{igst}$, the number of children individual i has in year t. The unit of observation in these regressions is at the adult child level rather than the grandparent level, to reflect the fact that the fertility decision is made by the grown children. This design also makes instrumenting for fertility more tractable. If instead I attempted to instrument for the total number of grandchildren, each adult child would require an age and birth year-dependent instrument for their fertility, so that the number of covariates would change grandparent to grandparent which is not feasible in this setting. This design allows for consistent instrumenting for total fertility and fertility timing while preserving the ability to observe the labor supply change from the marginal grandchild.[14] The other advantage is that it allows me to include adult children of any birth order, in contrast to the panel in Equation (3.1).

[14] One alternative is running the estimation strategy on just 1, 2, or 3 adult child families at a time, but the PSID is not a large enough national sample to permit cross-sections at this fine of a level without creating too many small cell sizes.

The family's PSID-provided 1968 sampling weight is adjusted to reflect the number of times a grandparent appears in this dataset, which is simply equal to the number of adult children they have. The other variables in this regression are otherwise the same as in Equation (3.1).

3.3.1.1 Endogeneity of Timing and Number of Grandchildren

If, however, adult children are basing the fertility decisions on anticipated changes in grandparent's labor supply, then Equations (3.1)-(3.3) cannot be consistently estimated. As discussed in the introduction, parents might time their fertility with anticipated changes in their parents labor force status, so that a panel fixed effects model would overestimate the impact of grandchildren. Similarly, they might wait to have children for when their parents achieve financial stability, so that the models underestimate the impact of grandchildren.

My identification strategy in light of this likely endogeneity is based on changes in legal barriers to abortion and contraception access that occurred throughout the US in the 1960's and 1970's. The identifying assumption is that there were no other state-by-year variables that also affected fertility coincident with the repeal of the access barriers. The number of children an adult woman has is modeled as being a function of access to oral contraceptives and abortion on-demand, the distance to an abortion early-legalization state, and eight lags on each policy.[15]

These policy changes are used to instrument for all three key variables discussed in the previous section. The first-stage regression for $ChildCount_{igst}$ is:

$$
\begin{aligned}
ChildCount_{igst} &= \pi_0 + \pi_1 PillAccess_{ist} + \pi_2 AbortionAccess_{ist} \\
&+ \pi_3 AbortionAccess\_LT250_{ist} + \pi_5 PillAccessLags_{ist} \\
&+ \pi_6 AbortionAccessLags_{ist} + \pi_7 AbortionAccessLags\_LT250_{ist} + \nu_{igst},
\end{aligned} \tag{3.4}
$$

[15]More information on the policy changes can be found in Sections B.1.1 and B.1.2.

where $PillAccess_{ist}$ is the fraction of year t that adult daughter i in state s could buy oral contraceptives under the age of 21; similarly, $AbortionAccess_{ist}$ codes the fraction of year t that an undesired conception could occur and then later aborted. $AbortionAccess\_LT250_{ist}$ is used to code access by grouping the distance a state is to either California, New York, or Washington D.C.,[16] because non-residents who were age-eligible could get an abortion. The three variables measure how far a pregnant woman would have to travel for an abortion, under the assumption that legalization's impact would be strongest in neighboring states. $AbortionAccessLags\_LT250_{ist}$, $PillAccessLags_{ist}$, and $AbortionAccessLag_{ist}$ are vectors of one- to eight-period lags for each policy variable.

To illustrate the intuition behind the lags, recall that *Roe* was decided in January 22, 1973. Women who conceived in all of November or December 1972 (and part of October) were eligible to end those pregnancies. Thus, for eligible women living in states whose statutes were invalidated by *Roe* are coded as having access for 71/366=19.4% of 1972. Conceptions between October 1972-January 1973 would have resulted in births in July 1973-October 1973, just after the PSID had concluded most of its 1973 interviews. Thus, had those conceptions been carried to term, the children would have first "appeared" in the 1974 survey. Especially for the young grandchild measure, whether you had access 5 years ago to abortion or contraception will partly determine whether you have a 4 year old in year t. Including the coding for the consecutive lags going back eight years accounts for all the possible timing combinations between conceptions, the ability to abort them, and when the PSID surveys were conducted, and allows for a fertility delaying effect that abortion and

[16] The categorizations were done by Levine et al. [80] and Ananat, Gruber, and Levine [3] on the basis of how the maximal distance a person would have to drive to get an abortion within half a day (<250 miles) or greater. Those papers do not code DC as a repeal state, as I do, so I made the requisite recategorizations. Joyce, Tam, and Zhang [69] offers compelling evidence that New York's lack of residency requirement, in particular, acted as an exogenous shock on birth rates in neighboring states. In my study, women are coded by age on the basis of how close they are to the closest early legalization state they are eligible to get an abortion at. For example, Washington State legalized abortion in December 1970, but minors needed parental permission. In May 1971, California legalized access for minors, so minors in Washington in 1971 are coded as having $AbortionAccess\_GT250_{ist} = 319/365 = 0.874$ while their adult counterparts are coded $AbortionAccess_{ist} = 1$. Washington State's policy had a residency requirement, so I assume that its legalization had no impact on women in neighboring states.

contraception permit.

The grandparenthood indicator, $\mathbb{1}\{Grandparent_{gst}\}$, is instrumented for after accounting for the change in the unit of observation. Grandparenthood status is now a function of the eldest daughter or daughter-in-law's exposure to changes in contraception and abortion access barriers. This takes the form of

$$\begin{aligned}\mathbb{1}\{Grandparent_{gst}\} &= \pi_0 + \pi_1 PillAccess_{gst} + \pi_2 AbortionAccess_{gst}\\ &+ \pi_3 AbortionAccessLT250_{gst} + \pi_4 PillAccessLags_{gst}\\ &+ \pi_5 AbortionAccessLags_{gst} + \pi_6 AbortionAccessLagsLT250_{gst} + \nu_{igst},\end{aligned} \tag{3.5}$$

where the policy variables described in Equation (3.4) are now the exposure for the eldest daughter or daughter-in-law for grandparent g to the changes in access.

3.3.2 Individual-Level Results

3.3.2.1 Panel Fixed Effects Estimates

Estimation results for Equations (3.1)-(3.3) are in Table 3.3, which reports the effect of being a grandparent, the marginal effect of an additional grandchild, and interactions between these measures with grandparent g's Social Security eligibility on four labor market outcomes: being retired, annual number of hours worked, and being in the labor force (grandfathers) or reporting non-zero working hours (grandmothers).

Being a grandparent does have a significant labor force detachment effect. Retirement propensity increases for both grandfathers (by 8.3%) and grandmothers (3.7%), and annual hours worked decreases for grandfathers by 142 hours and for grandmothers by 41.7

hours. When Social Security eligibility is factored in, neither grandfathers nor grandmothers seem to have a differential response to grandchildren when their ability to retire changes. The point estimates on grandparenthood status do change slightly when these interactions are included, but only grandfather's labor force participation effect goes from non-significant to significant, showing that they are 2% more likely to exit the labor force when they are grandparents.

The marginal effect of each additional grandchild for grandfathers is that they are 3.5% more likely to be retired, 1.1% less likely to be in the labor force, and to work about 41.5 hours less annually with each additional grandchild. Grandmothers have a similar labor market response to the number of grandchildren: working 41.4 fewer hours annually for each additional grandchild, becoming 1.9% more likely to retire, and 1.8% less likely to report non-zero annual hours worked. The "In Labor Force" measure and "Non-Zero Working Hours" measure are not directly comparable, but broadly speaking, both grandmothers and grandfathers decrease their labor force attachment and labor supply with each additional grandchild.

In contrast to the grandparenthood status regressions, a grandchild's marginal effect is expressed differentially on the grandparent's Social Security eligibility status. Each grandchild increases the propensity to be retired by 5.3% when the grandparent is younger than 62, but only 3.6% more likely to retire when the grandchild arrives between ages 62 to the full retirement age (FRA), and just 1% more likely to retire when the marginal grandchild arrives after the FRA.

This shrinking grandchild effect is likely driven by two forces. The first is that the cumulative differential effect of retirement might have pushed out most of the people who would respond to grandchildren before age 62. Recall that this is a person-by-year panel, so that the result holds that, on average, a grandfather is 8.3% and a grandmother 3.7% more likely to retire in a given year prior to age 62 than the grandchildless. If grandchildren start arriving around ages 40-45, the cumulative effect may leave a smaller pool of grandfathers who will then differentially uptake retirement at age 65. The second force is that so many people uptake retirement upon hitting either age 62 or their FRA that the remaining variation attributable to grandchildren may be relatively small.

3.3.2.2 First-Stage Results

The results in Table 3.3 cannot be understood as a causal labor supply effect until the endogeneity concern is addressed. Table 3.4 shows the first-stage estimates for Equations (3.4) and (3.5). For the access to contraception and abortion to be a valid instrument for the number and timing grandchildren, the results should show evidence that exposure to the policies changed fertility timing and total parity. Reported is the effective F-statistic for the result of a weak instrument test estimated with cluster-robust standard errors using the procedure described in Kleibergen [74]. Under that test, all specifications can reject the null that they are weak instruments using the Staiger and Stock [109] "rule of thumb" of rejecting an instrument as weak if $F > 10$.[17]

The coefficient estimates on the policy variables largely affirm the intuition that the pill and abortion decreased total fertility and also induced delayed childbearing. The coefficients for predicting $\mathbb{1}\{Grandparent_{gst}\}$ are negative on both contemporaneous policy variables. Having access to the pill in year t decreases the chance of being a grandfather by 5.7% and a grandmother by 4%. For abortion, these figures are 6% and 4.4%, but neither coefficient is statistically significantly different from zero. In the lags, the pill coefficients are all positive for grandparenthood status, and become larger in further lags. By the 8th period lag, both the pill and abortion indicators are positive and significant. Having access to the pill in $t-8$

[17] To date, there are not formal critical values used for the Kleibergen-Paap rk Wald F statistic. Some sources use instead the Stock and Yogo [110] critical values for the Cragg-Donald Wald F statistic, which assumes i.i.d. errors. Nonetheless, the F statistics reported in the table exceed the threshold for 5% maximal IV relative bias.

increases the chance a person is a grandfather by 12.7% or a grandmother by 12.8%, and for abortion, these numbers are 13.7% and 2.4%, respectively. Thus, abortion and contraception both lower the person's probability of being a grandparent, but these odds are somewhat mitigated if the policy has existed for several years.

The coefficients' interpretation in the $ChildCount_{igst}$ regressions shows a stronger dampening effect for the policies on total fertility. Access to the pill and abortion negatively predict the number of children the adult child has, with access to abortion predicting lower total fertility out to at least the third lag and access to the pill out to the second lag. Access to the pill in year t predicts 0.158 fewer grandchildren per adult child for grandfathers and -0.148 fewer grandchildren per adult child for grandmothers. Access to abortion in year t predicts 0.104 fewer grandchildren per adult child for grandfathers and -0.106 fewer grandchildren for grandmothers. If the adult child was living in an adjacent state in the immediate pre-*Roe* years, then they would have 0.166 fewer grandchildren per adult child for grandfathers and 0.138 fewer grandchildren per adult child for grandmothers. Like the estimate for Equation (3.5), in the 8th lag, the coefficients become positive, relatively large, and statistically significant, indicating that abortion and the pill exerted a delayed childbearing effect.

3.3.2.3 Second Stage Results

The results of the instrumental variables regressions reported here offer several insights into how the evolution of grandparenthood in the past seven decades may have influenced the fall and rise in older worker's LFP. Results with the instrumented values of $\mathbb{1}\{Grandparent_{gst}\}$ and $ChildCount_{igst}$ are reported in Table 3.5. It's clear from a comparison with the results in Table 3.3, that a panel fixed effects model underestimates the impact of grandchildren. Being a grandparent makes grandfathers and grandmothers substantially more likely to be retired, at 19.4% and 8.5%, respectively, and this effect is reinforced as retirement becomes more feasible. Grandfathers are 23.3% more likely to retire, and this is increased to 41.7% when they become eligible for early retirement at age 62, and are 29.8% more likely to be retired when they reach their FRA.

The evidence from annual hours worked and labor force status also shows that grandfathers supply less labor than the grandchildless. Grandfathers work 362.9 fewer hours a year compared to the grandchildless. Before early retirement eligibility kicks in, this working gap is 409.2 fewer hours a year, and then it rises to working almost 962.6 fewer hours a year when the older workers are eligible for early retirement through Social Security. Interestingly, the gap then shrinks to only 182.5 when these older workers arrive at their FRA, but the interaction between grandfatherhood and full Social Security retirement eligibility is not statistically significant. Unlike retirement status or annual hours worked, being a grandfather does not make you statistically significantly less likely to be in the labor force. The exception is that when you are eligible for early retirement, the interaction between grandfatherhood and elibility is 21.5% at the 10% level, with a total working gap of 28%.

Grandmothers exhibit somewhat more mixed behavior patterns to grandchildren than grandfathers. They are 13.4% more likely to be retired than their grandchildless peers, and this increases to being 55.1% more likely to be retired when they become eligible for early retirement. Yet, at their FRA, the gap between grandmothers and older women without grandchildren reverses: women without grandchildren become 26.4% more likely to be retired than those with grandchildren, but this is not statistically significant. Grandchildren do not appear to have a statistically significant effect on grandmother's annual hours worked, although their effect is surprisingly positive, and is even more positive when grandmothers arrive at their FRA. The annual hours worked results are inconclusive, especially since the likelihood that grandmothers report non-zero working hours decreases by 16.1% when they have grandchildren.

The estimates of total fertility effects on grandparents' labor force attachment offer somewhat different evidence for how extended families influence older workers. In general, these results also indicate that grandchildren cause grandparents to lower their labor supply. With or without Social Security eligibility interactions, grandfathers and grandmothers lower their labor supply with each additional grandchild across all outcomes. Further, the coefficients on the marginal grandchild are all significant at the 1% level.

However, what is different is that the interactions between the marginal grandchild and Social Security eligibility run counter to what was observed for grandparent status. For example, the marginal grandchild increases retirment propensity by 27%, but if this marginal grandchild arrives when the grandfather is eligible for early retirement, the combined marginal grandchild effect makes retirement only 15.9% more likely. When the grandfather reaches his FRA, the combined marginal grandchild effect drops even further to 5.4%. This pattern is repeated for the annual hours worked and for labor force status. The marginal grandchild decreases annual hours worked by 401.5 hours, but this becomes only 73 fewer hours if this grandchild arrives when the grandfather is at his FRA. For labor force status, these numbers are 17.1% less likely to be in the labor force with each additional grandchild, but only 8.9% less likely when the grandfather is at his FRA.

Grandmothers' response to additional grandchildren tells a much more straightforward story than their grandmother status results. Each grandchild makes them 21.7% more likely to be retired, work 431.6 fewer hours a year, and be 19.1% less likely to report non-zero working hours. When broken out by retirement eligibility, the impact of grandchildren is compounded at each stage and causes strunger labor force detachment. Now, the marginal grandchild causes the grandmother to be 24.4% more likely to be retired, work 499 fewer hours a year, and be 22% less likely to report non-zero working hours. When they reach early retirement status, the combined effect of the marginal grandchild is now 33.5% more likely to be retired, 607.82 fewer annual hours, and 26.2% less likely to report non-zero hours worked; at full retirement status, these numbers are 32.2%, 683.63, and 30.6%.

A cautious but plausible interpretation comparing the grandchild count results with the grandfather status results suggests that grandfathers are differentially less likely to be in the labor force at all stages, and that the n^{th} grandchild (*i.e.* more likely to arrive when the grandfather is 62 or older) has a less pronounced effect than the first or second grandchild. By contrast, perhaps, grandmothers become differentially more engaged with their families the bigger they become, perhaps because childcare needs grow non-linearly with the number of grandchildren.

From these results, it's clear that grandparents are less attached to the labor force than non-grandparents, and those with bigger extended families are even less so. It is also evident that the most economically significant responses are in the retirement and labor force regressions. Pre-retirement, I estimate that grandfathers work between 362.9-409.2 fewer hours than non-grandfathers, but this comes out to about 9-10.25 fewer working weeks a year. Not a trivial amount, but still consistent with working about 32 hours a week, which would still qualify many people for benefits associated with working full time. However, I also estimate that grandfathers are between 19.4%-23.3% more likely to be retired. This means that if the fraction of men between 55-61 who are grandfathers rises 10% (like what was seen in the Baby Boom), then approximately an additional 2% will be retired.

3.4 National Labor Force Participation Trends Estimation

The results of Section 3.3 suggest that grandchildren alter grandparent's labor supply at different rates depending on their retirement eligibility. This results informs the empirical strategy for national-level trends because it gives a starting place for the expected lag between the adult children's fertility decision and the grandparent's response.

I now build off of Blau and Goodstein [22] to estimate how changes in the supply of grandchildren change older workers' labor supply. They use data from the CPS, Social Security Administration (SSA), and other sources to model the employment decision rule for older workers. Their model accounts for how variation in Social Security benefits, disability insurance, educational attainment, and the labor force participation of spouses impact the fall and rise in labor force participation among men aged 55-69. Identification in their model occurs from variation at the year-by-birth year-by-education group level.

Their paper also focuses exclusively on older men, and for this analysis, I too will only analyze the labor force participation of men. Given the sea change in labor force attachment shown by women between 1962 and the present, credibly estimating a model for women is an exercise that will be left for future research. Further, I will also use the CPS instead of the PSID, in part because the PSID sample was unrepresentative of the nation at various points in its cycle, and the CPS is designed specifically to permit credible estimates of national-level descriptive statistics from micro data.

I augment their model by first extending the time series out to 2015, second by extending the panel to include 50-54 year old men, and third by adding one of two grandparent measures: the fraction grandparents or the average number of grandchildren. I run all specifications under the assumption that agents have perfect foresight, because it is a more "conservative" assumption from an identification standpoint. Further, papers that have assumed myopic expectations have generated results that are counterintuitive and specification-sensitive (Kreuger and Pischke [76]; Blau and Goodstein [22]). If, however, all agents have perfect foresight, then births are done anticipating grandparents' behavior so that even in the aggregate, birth rates and labor force participation rates are co-determined. I will thus instrument for changes in birth rates with the variation in reproductive technology access laws. i will also add to my instrument variation in statewide Prohibition laws, which others have found to affect births and will provide some important exogenous variation for the oldest members of my CPS sample, who were typically born between 1892-1895 and would be expecting to become grandparents mostly from 1927 onwards.

Addressing endogeneity by means of an instrumental variables approach like the one used in Section 3.3 is not straightforward in the CPS. Expanding the panel to the year-by-birth year-by-age-by-education level-by-state would allow me to identify changes in grandparenthood characteristics using state-level birth rate variation. However, this would come at the cost of generating many small cell counts in the CPS. Building a panel that properly identifies variation in all causal channels is impossible in a dataset like the CPS, which can only survey so many households at a time.

Thus, I strike a balance by conducting the analysis first purely as an extension of the Blau and Goodstein approach, adding controls for grandparent status. This model implicitly assumes that grandparents have perfect foresight about Social Security, earnings, and inflation expectations, but take changes to their grandparenthood status as randomly and exogenously determined. This analysis allows me to directly compare my approach with previous analyses. I then drop the grandparenthood exogeneity and randomness assumption and instrument for grandparent measures by compressing the education level cells used by Blau and Goodstein and reaggregating the synthetic panel to the year by birth year by state level. This sacrifices some identification in the educational attainment channel, but permits full identification of fertility effects and birth cohort effects on labor force attachment.

One tradeoff I am forced to make with this research design is in my cutoff for cell size. In the Blau and Goodstein extension, I use the same cutoff of 30 respondents minimum per cell. For my state panel, this would leave me with only 10% of my initial sample. Therefore, I instead use a cutoff of 10 respondents, which leaves me with 52% of my original sample from the CPS.

3.4.1 Labor Force Participation Model: Blau and Goodstein Extension

I begin by modifying the model created by Blau and Goodstein [22], which looks at labor force participation among older men by creating a simulated panel of older men by year by birth year by education grouping and takes the form:

$$\begin{aligned}
&LFP_{eabt} = \delta_0 + \delta_1 GP_Measure_{eabt} + \delta_2 SSB65_{eb} \\
&+ \delta_3 (SSB62_{eb} - SSB65_{eb}) + \delta_4 (SSB62_{eb} - SSB65_{eb}) \\
&+ \delta_5 AME_{eb} + \delta_6 DisabilityBenefit_{eabt} + \delta_7 Demographics_{eabt} \\
&+ \delta_8 EducationGroup_e + \delta_9 Year_t + \delta_{10} BirthYear_b + \delta_{11} Age_a + u_{eabt},
\end{aligned} \tag{3.6}$$

where $GP\_Measure_{eat}$, the key variable of interest, is either the fraction who are grandparents in each age cohort a and birth year b in year t at education attainment level e or the number of grandchildren; $Demographics_{eat}$ controls for the fractions married, previously married, white, black, U.S. Armed Services veteran, or reported being in bad health; $EducationGroup_e$ is a vector of indicators for either having less than high school education, a high school education, some college, or college-plus; $Year_t$ is a vector for year dummies; Age_a is a vector of age dummies; and $BirthYear_b$ is likewise a vector for birth year dummies.

The Blau and Goodstein empirical model approximates the decision rule for labor force participation at older ages under a life cycle model of employment and retirement where men seek to maximize the expected present discounted value of remaining lifetime utility, subject to various constraints.[18] The decision rule for Social Security participation is estimated by means of the retirement benefits a worker could receive at ages 62 ($SSB6_{eb}2$, early retirement), 65 ($SSB65_{eb}$, or full retirement), and 70 ($SSB70_{eb}$, or delayed retirement).

[18] More information on their model can be found on Blau and Goodstein [22], p. 332.

Differencing between $SSB6_{eb}2$ and $SSB65_{eb}$ models the tradeoff between early and full retirmeent, and likewise, the difference between $SSB70_{eb}$ and $SSB65_{eb}$ the tradeoff between earning the Delayed Retirement Credit (DRC) and accepting full retirement, or primary insurance amount (PIA). To separate the Social Security wealth effect from changes in lifetime earnings, the average monthly lifetime earnings, AME_{eb}, from ages 27 to 65 for the average worker in birth cohort b at education level e is included. Higher values of $SSB70_{eb} - SSB65_{eb}$ imply a stronger incentive to delay retirement, and likewise, lower (more negative) values of $SSB62_{eb} - SSB65_{eb}$ also imply a stronger incentive to delay retirement.

The model includes the average monthly Social Security Disability Insurance amount received by a worker in birth cohort b at education level e if they were to work until year $t-2$, receive no earnings in year $t-1$ and then be on SSDI from year t until age 65. The lack of earnings in year $t-1$ mimics the 5 month waiting period a worker must observe before receiving SSDI.

The model also includes the fraction of married men whose spouse's are in the labor force, *pace* Schirle [105] and Gustman and Steinmeier [54] that men may prolong their labor force attachment out of a desire to jointly retire with their wives. However, since this rate is co-determined with the men's own labor force participation rate, I instrument for the wives' labor force participation rate using each cell's average spousal birth year and spousal birth year squared. I chose these instruments to reflect the findings from Schirle [105] and Goldin and Katz [10] that successive birth cohorts of women entered and remained in the labor force, although this phenomenon has leveled off in the past decade. Since I am not attempting to establish a causal relationship in this paper between female labor force participation rate, their fertility, and their husband's behavior, taking a constant coefficients approach to instrumenting for wives' labor force participation rate is sensible and parsimonious. The instrument's F-statistic is reported in Table 3.6.[19]

[19] Blau and Goodstein acknowledge that this variable may be endogenous, but do not address this question directly. This aspect of the model can thus also be considered as an extension of their work.

The birth year dummies are particularly important with respect to identification of effects other than grandparenthood. Changes in Social Security can typically be identified either by exogenous changes in eligibility rules, non-linearities in benefit rules, or variation in lifetime earnings growth across birth cohorts. The first is perfectly collinear with birth year fixed effects, and in Blau and Goodstein [22], they report that relying on variation other than these exogenous rule changes yields counterintuitive and problematic results that do not seem to capture the variation in labor force participation. Thus, I anticipate that the results will be sensitive to what level of birth year fixed effects I include, so I run several specifications to control for birth year effects: the birth year squared; single year birth cohort fixed effects; the birth year squared plus 2 year birth cohort fixed effects; and the birth year squared plus 4 year birth cohort fixed effects.[20] In the tables, I refer to the specifications without birth cohort fixed effects as those with “Time Trends".

3.4.2 Labor Force Participation Model: Instrumenting for Grandfatherhood

However, if agents do have foresight, then the results of Section 3.3 suggest that current labor force status may be codetermined with grandparenthood status, even in the aggregate. Instrumenting for grandparenthood status via the state changes in abortion and contraception laws is impractical with the above model. Disaggregating the data by state means that 90% of the observations now have fewer than 10 respondents. Therefore, I present a second model that is disaggregated at the state level, but aggregates respondents by their education

[20]The first order birth year is perfectly collinear with year and age fixed effects, so I omit it.

level:

$$
\begin{aligned}
LFP_{sabt} &= \delta_0 + \delta_1 GP\_Measure_{sabt} + \delta_2 SSB65_{sb} + \delta_3 (SSB62_{sb} - SSB65_{sb}) \\
&+ \delta_4 (SSB62_{sb} - SSB65_{sb}) + \delta_5 AME_{sb} + \delta_6 DisabilityBenefit_{sabt} \qquad (3.7) \\
&+ \delta_7 Demographics_{sabt} + \delta_8 State_s + \delta_9 Year_t + \delta_{10} BirthYear_b + \delta_{11} Age_a + u_{sabt},
\end{aligned}
$$

where s is the index for the state of residence. $Demographics_{sabt}$ now includes four new variables measuring educational attainment: fraction with less than a high school education, the fraction with a high school degree, the fraction with some college, and the fraction with a college degree or more. Otherwise, all other variables are substantively the same.

From this model, I then instrument for $GP\_Measure_{sabt}$ using:

$$
\begin{aligned}
GP\_Measure_{sabt} &= \gamma_0 + \gamma_1 Abortion21Plus_{sabt} + \gamma_2 Abortion18to20_{sabt} \\
&+ \gamma_5 Pill21Plus_{sabt} + \gamma_6 Pill18to20_{sabt} + \gamma_7 Prohibition_{sabt} + \eta_{sabt}, \qquad (3.8)
\end{aligned}
$$

where $Abortion21Plus_{sabt} = t_a - t_{33}$ is the number of years from the time a person in birth cohort turned 33 ($t_{33} = b + 33$) until their current age (t_a) that state s had unhindered legalized abortion for women 21 and over. Likewise, $Abortion18to20_{sabt}$, $Pill21Plus_{sabt}$, $Pill18to20_{sabt}$ are the number of years from t_{33} to t_a a woman could get unhindered access in state s and year t to the reproductive technology in the variable's name for the ages indicated.

This specification recognizes that policy lags in this setting are inaccurate and imprecise: a 10 year lag on a 50 year old grandfather is going to be a weaker predictor for his grandchildren measures than it would be for a 60 year old grandfather. By the same token, a one year policy lag will only weakly predict variation in grandchildren measures for 69 year old grandfathers, but will do so much more strongly for 50 year olds. Thus, instrumenting for the grandchildren measures by means of total policy exposure by age, year, state, and birth cohort permits a more robust identification of variation in grandchildren measures.

The variable $Prohibition_{sabt}$ measures the same policy exposure as the other variables, but for whether the state had a ban on alcohol sales. I included this variable both because some studies have shown that state and federal Prohibition had significant impacts on infant health (Jacks et al. [67]), and more generally, acknowledging the body of literature that alcohol regulation can impact fertility outcomes (Nilsson [94]; Cintina [34]). More information on Prohibition, including a brief review of the literature on Prohibition and its effects on births, can be found in Appendix B.1. The appendix makes clear that like abortion and contraception access, Prohibition and its repeal has substantial state-by-year variation that can be exploited in a model like Equation (3.8).

The other reason to include a measure for Prohibition is that to calibrate my grandparent measures, I used the full sample of U.S. birth rates which extends from 1915 to the present. National Prohibition went into effect on January 16, 1920, and was repealed on December 5, 1933. Both before the ratification of the 18th Amendment to the Constitution and its repeal by the 21st Amendment, many states had statutory Prohibition. In the 10 years leading up to the 18th Amendment, the temperance movement gained traction and several states enacted state-wide Prohibition, so that by 1920, 32 states and the District of Columbia had banned the sale of alcohol (Wickersham [117]). After Prohibition ended, the states with statutory or constitutional Prohibition continued to enforce bans on the sale of alcohol, with Mississippi being the last state to repeal statewide Prohibition in 1966.[21] Thus, its inclusion should give a wider spectrum of policy variation to draw on when instrumenting for fertility.

[21] See Appendix B.1.4 for more information on state and federal Prohibition laws.

3.4.3 National-Level Estimation Results

3.4.3.1 Results from Blau and Goodstein Extension

Table 3.6 presents the results of estimating (3.6) with and without interactions with retirement eligibility, both for the fraction that are grandfathers and with the average number of grandchildren as the key variables of interest. I present four different specifications for controlling for the impact of birth cohort: birth year as a second order polynomial, birth cohort fixed effects, 2 year birth cohort fixed effects and the birth year squared, and 4 year birth cohort fixed effects and the birth year squared. Like Blau and Goodstein, my results are sensitive to how the birth cohort effect is controlled for.

All fractions are multiplied by 100 before the regressions are run, so coefficients for the remaining regressions (unless otherwise noted) are interpreted as the amount the LFP rate changes on the scale of (% in LF) $*$ 100. The coefficients in the first four rows represent the change in older workers' LFP rate in response to a 1 point increase in the fraction of older men who are grandfathers. Columns 1, 3, and 4 show that 1 point increase in the fraction of older men who are grandparents lowers the labor force participation rate by between 0.688-0.914 points. Adding interactions for retirement eligibility changes this range to a drop of 0.402 to 0.656. The interactions indicate that the labor force participation rate drops another 0.124 to 0.128 at early retirement age for each 1 point rise in the fraction grandparent, and a further 0.449 to 0.510 points for a 1 point increase in the fraction grandparent among men of full retirement age.

The marginal grandchild likewise decrease older male workers' LFP, either ranging from a 8.15 rate point drop (Column (1)) to a 12.68 point drop (Column(4)). Adding interactions with retirement eligibility changes this range from the marginal grandchild causing a 5.73 point drop in LFP (Column (1)) to a 10.31 point drop (Column (4)). Being eligible for early retirement and having an additional grandchild modestly decreases LFP further by between 0.76 to 0.85 point drop. As with the PSID individual-level estimates, being eligible for full retirement yields a more dramatic change: LFP declines by between an additional 7 to 7.4 point drop with each additional grandchild. This result is noteworthy in part because it is the most robust to birth cohort specification.

The clear outlier here is Column 2, which finds labor supply drops 34.69 points with no interactions and rises 6.12 rate points with interactions included for every 1 point increase in the fraction grandparent. Both the coefficients' magnitudes and their sign change across specifications are implausible. The results for grandchild count are likewise implausible: each additional grandchild is found to decrease labor force participation by 429 points, but with retirement eligibility interactions include, this changes to each grandchild decreasing LFP by 36 points, with only a very modest decrease of 1.6 points if the grandfather is early retirement eligible and 7.13 points if he is at or past his FRA. The problem is that the remaining identification after birth year fixed effects are included comes through year-by-birth cohort, education group-by-birth cohort, or education group-by-birth year-by-year variation, such that unobserved shocks only impacting certain segments of a birth cohort or a birth cohort only in certain years are neither well-motivated nor well-understood in this empirical framework.[22] 2 year birth cohort effects also leave little between-cohort variation, so I will only present results that either have the birth year squared or 4 year birth cohort fixed effects for the remainder of the paper.

In Table 3.7, I present the remaining coefficient estimates for the quadratic birth cohort and 4 year birth cohort effects models from Equation (3.6). Notably, there is practically no difference in coefficients whether I use $\% \, Grandfather$ or $Grandchild \ Count$ as my grandparent control. This suggests that grandparenthood is both an important factor in predicting LFP

[22]Blau and Goodstein also find that using birth cohort fixed effects yields counterintuitive and implausible results, because the remaining policy variation in Social Security benefits comes through non-exogenous rule changes in benefit calculations.

but that to a striking degree its effects are uncorrelated with most other structural and demographic covariates. Also notable is that with a few exceptions, the coefficients here share the same signs as those reported in Blau and Goodstein. The social security and monthly disability benefit amounts are scaled down by 100, so that the coefficients represent the point change in the LFP rate in response to a $100 increase in these benefits. This interpretation is also true for the lifetime average monthly earnings coefficients. The wealth effect from Social Security benefits causes the sign on $SSB65$ to be negative, and a smaller gap between the PIA and benefit levels available at 62 causes greater labor force detachment, although this is not statistically significant in my model. Likewise, a greater credit for remaining in the labor force past the FRA ($SSB70 - SSB65$) prompts a higher LFP rate.

I also find that a 1 point increase in being in bad health lowers the LFP rate by between 0.609 to 0.720 points, and that a 1 point rise in the fraction married or fraction previously married raises LFP by between 0.103 to 0.l23 points and 0.101 to 0.126 points, respectively. It is worth noting that in almost every specification, the coefficient on previously married is larger than the coefficient on married, perhaps because previously married men lose some wealth after a divorce, separation, or widowhood that they then want to regain before retirement. A rise in the fraction who are veterans also decreases labor force participation, but this effect is small and not always statistically significant, with the significant values ranging between -0.019 to -0.047 points.

The remaining controls for the log predicted wage, the monthly disability benefit, and the lifetime average monthly earnings are where my results differ most with Blau and Goodstein. In contrast to their approach, my model includes education group-by-year fixed effects to control for unobserved yearly shocks to groups with different skills. This change prompts the coefficients on lifetime monthly average earnings to become negative, suggesting that when shocks are accounted for, the income effect dominates the substitution effect when there is more lifetime wealth.

Another contrast in my findings is that when I have no retirement eligibility interactions, a \$100 rise in the monthly disability benefit decreases the LFP rate by between 0.126 to 0.175 points. However, including the retirement eligibility and grandparenthood interactions causes the sign on disability benefits to flip to positive. This puzzling result has no easy explanation, so I note here that the positive coefficients on the monthly disability benefits are never statistically significant and are small: only ranging from 0.004 to 0.067 points per \$100 benefit increase. A similar effect appears to be at work on the log of the predicted hourly wage, which indicates for the regressions with no grandparent/retirement interactions (Columns (1)-(2) and (5)-(6)) that the labor supply elasticity of predicted wages is between 1.648-2.033, but loses significance when the eligibility interactions are included, and becomes negative in Columns (3), (4), and (8). Like the monthly disability measure, I only note that the negative elasticities are small, ranging from -0.057 to -0.295.

Lastly, I report the second-stage result for the fraction of the spouses in the labor force, which ranges from 0.168 to 0.591 points. The Kleibergen-Paap Wald F-statistics reported on the last table row are all at least 93 or more, whereby the hypothesis that the instruments are weak can safely rejected.[23] The coefficients without 4 year birth cohort effects are about double (ranging between 0.452 to 0.591 points) those with the cohort effects (0.168 to 0.227 points).

The results in Tables 3.6 and 3.7 and the results from the PSID data in Table 3.5 imply that there are economically meaningful interactions between the presence and number of grandchildren and the grandparent's retirement and consumption tradeoff. To understand these tradeoffs better, I expand the interaction terms to include the social security benefit controls, $SSB65$, $SSB62 - SSB65$, and $SSB70 - SSB62$. These results are presented in Table 3.8.

[23]First-stage and ordinary least squares estimates are available in supplemental tables. Typically, the OLS estimates are around 0.1% on the spouse in labor force measure, suggesting that the endogeneity bias is caused by married women having lower baseline LFP than unmarried women.

For the sake of comprehension, I report the coefficients for key variables of interest, although the only controls that changed from Tables 3.6 to 3.7 were the grandparent measures. Compared to the results in Table 3.6, the coefficient on the fraction grandfather/early retirement eligiblity is about 60-70% lower. However, the interaction between fraction grandfather, early retirement eligibility, and the difference between the early retirement monthly benefit and the PIA is negative and statistically significant at the 1% level. With or without the 4 year cohort effects, the coefficient is about 0.020 to 0.025 points which implies that grandparents will not differentially exit the labor force at a higher rate at age 62 compared to non-grandparents, unless the early retirement penalty is lessened.

The full retirement and grandfatherhood interaction remains statistically significant at the 1% level and negative, and this relationship appears to be invariant to interactions with PIA levels or to changes in the DRC. Also important, is that the coefficients on the fraction grandparent and full eligibility interactions are about 8-7 times higher than the coefficients on the early retirement interaction (Columns (1) and (2)). Similarly, the coefficients on the average grandchild count and full retirement eligibility are about 4.25 times higher than the coefficients on average grandchild count and early retirement eligibility (Columns (3) and (4)). The conclusion from these results is that being a grandfather decreases older men's LFP, become marginally more likely to retire upon reaching early retirement, and become substantially more likely to exit the labor force than the grandchildless upon reaching full retirement regardless of changes to benefit levels.

3.4.3.2 Results from CPS State Panel

The endogeneity problem remains, because if fertility and grandparenthood systematically adjust based on expected changes in grandparents' labor force participation opportunities, then the results reported in Section 3.4.3.1 would be biased. Therefore, I recreate the synthetic panel by state (but not by education group) and instrument for my grandparent measures using Equation (3.8). Panel fixed effects results are reported in Table 3.9, first-stage results are reported in Table 3.10, and the second-stage results are reported in Table 3.11.

The state panel's fixed effects results in Table 3.9 differ somewhat from the national-level panel results from Tables 3.6 and Tables 3.8. This is expected, given that different fixed effects are used and that no endogeneity bias has yet been removed, either from the grandparent measures or the spousal LFP measure. Importantly, coefficients that are statistically significant across all tables retain the same sign in almost every case, except for lifetime average monthly earnings. In the state panel, the coefficient on a $100 increase in average earnings increases the LFP rate by between 0.15 to 0.17 points (ignoring the coefficients that are not statistically different from zero). In the previous tables, this was an inverse relationship, but appears to be driven solely by using year-by-education group fixed effects in Table 3.7. Without these, then the coefficients on lifetime average earnings are positive as they are in the state panel.

My results for instrumenting for $GP_Measure$ are in Table 3.10. For economy, I only present the results corresponding to the $GP_Measure$ specifications without interactions, whose OLS results are in Table 3.9, Columns (1), (4), (7), and (10). It is most crucial to show that the instrument is relevant in this context for $GP_Measure$, so that the instrument's validity for the extended interactions will then follow naturally.

The results here show the strongest links between the years of exposure to abortion access for women 21 and older than to any other form of birth control. Every additional year of exposure to this policy reduces the fraction grandparent by between 0.81 to 0.93 points. Similarly, every additional year of policy exposure to abortion for women 21 and over reduces the average number of grandchildren by between 0.06 to 0.07. When 4 year birth cohort effects are included, an additional year of exposure to Prohibition increases the fraction grandparent by 0.41 and the average number of grandchildren by 0.03, both at the 10% level. The second order spousal year of birth is also statistically significant in these regressions, with the negative coefficient on the spousal year of birth squared indicating that later-born spouses predict fewer grandchildren, but only when the 4 year birth cohort fixed effects are not included. Overall, these appear to be strong predictive models for the grandparent measures, as the adjusted R-squares are 0.98 across the specifications.

In Table 3.11, the Kleibergen-Paap Wald F-statistic of the instrument's strength reports values of 100.01 corresponding to Table 3.10, Columns (1) and (3) and values of 50.74 for Columns (2) and (4). The weak instrument hypothesis can thus be safely rejected at these levels.

Table 3.11 also shows the second-stage instrumental variables regression results. Notably, just as in Section 3.3.2, conventional estimates of grandchildren trends put the effect at closer to zero than estimates stripped of endogeneity bias. This reinforces the idea that families are systematically more likely to have grandchildren when the grandparents are in the labor force, even on an age-adjusted basis.

The results in Table 3.11 point to three conclusions. The first is that grandparents do not take early retirement at higher rates differentially from the grandchildless, but can be induced to do so as the early retirement penalty shrinks. The interaction between fraction grandparent and early retirement eligibility is only statistically significant when the 4 year birth cohort fixed effects are included, but it is not very large: a 1 point increase in the fraction grandparent when they are early retirement eligible reduces the LFP rate by 0.11 points (Column (5)). For the average number of grandchildren specifications, an genadditional grandchild while early retirement eligible reduces the LFP rate by 1.02 points (Column (11)). However, when benefit level interactions are included, the grandparent measure/early retirement eligibility/early retirement penalty interaction is statistically significant across all specifications at the 1% level, but the early retirement eligibility/grandparent measure interactions all become statistically indistinguishable from zero, except in the eligibility/average grandchild count interaction in Column (12). A 1 point increase in the product of the early retirement penalty and the fraction grandparent at early retirement would decrease the LFP rate by 0.02 to 0.03 points (Columns (3) and (6)). For the average grandchild count, the LFP decrease ranges between 0.35 to 0.44 points (Columns (9) and (12)). Essentially, these effects work out so that when you use the average values, a one grandchild increase at 2015 values of the early retirement penalty would actually cause even more people to wait for full retirement, but a one grandchild increase at 1965 values of the early retirement penalty would very slightly incentivize people to exit the labor force during the early retirement phase.

The second finding is that the largest and most specification-robust grandchildren impact channel is when older workers become eligible for full retirement. Each 1 point rise in the fraction grandparent when workers are at their FRA or older decreases the LFP rate by 0.43 points (Column (2)), and when including interactions with Social Security benefit levels, this coefficient becomes -0.29 (Column (3)). Adding 4 year birth cohort fixed effects amplifies this effect, so that a 1 point increase in fraction grandparent decreases the LFP rate by 0.51 points (Column (5)), or by 0.37 points when including benefit level interactions (Column (6)). A growth in the extended family also decreases the LFP rate: each rise of an additional average grandchild decreases the LFP rate by 5.07 points (Column (8)), and by 6.09 points if 4 year birth cohort fixed effects are included (Column (11)). Adding interactions with benefit levels changes this coefficient to -4.01 with no birth cohort fixed effects (Column (9)), and -5.07 with the fixed effects included (Column (12)).

The interactions between the grandparent measures, full retirement eligibility, and the benefit levels are statistically significant, but exert only a small effect on the LFP rate. An increase by one unit of the PIA/fraction grandparent product will increase the LFP rate at the full retirement age by 0.02 points (Columns (3) and (6)), and an increase by one unit of the PIA/average number of grandchildren product will increase the LFP rate by 0.22 to 0.23 points (Columns (9) and (12)). An increase by one unit off the delayed retirement credit/fraction grandparent product will decrease the LFP rate by 0.01 points (Columns (3) and (6)). Lastly, an increase by one unit off the delayed retirement credit/average number of grandchildren product will decrease the LFP rate by between 0.09 points (Column (9)) and 0.12 points (Column (12)), but is only significant in the 4 year birth cohort fixed effects specification. These coefficients lack an obvious economic interpretation, but their overall impact is very small so whatever force is driving these coefficients is ultimately only a small part of the overall variation in the LFP rate. Altogether, these results indicate that grandchildren push grandfathers out of the labor force upon full retirement at roughly 10 times the rate they do in even the early retirement period.

The third finding is that the evidence that grandchildren impact LFP trends among those younger than 62 is mixed. Only the specifications without the birth cohort fixed effects show a statistically significant negative from either grandparent measure. When all interactions are included, each 1 point rise in the fraction grandparent decreases the LFP rate by 0.32 points (Column (3)), but when the birth cohort fixed effects are included, this changes to a 0.11 point *increase* (Column (6)) that is not statistically distinguishable from zero. Likewise, an additional grandchild would decrease the LFP rate by 4.18 points (Column (9)), but this changes to a non-significant 1.56 point increase when the birth cohort fixed effects are included (Column (12)).

To understand what these contradicting results mean across the entire panel, Table 3.12 reports the marginal effect of the two grandparent measures and the three Social Security benefit measures. The marginal effect of the fraction grandparent is -0.252 (Column (1)) to -0.298 (Column (2)) and significant at the 10% level (Column (2)), but changes to be between 0.212 (no benefit interactions, Column (3)) to 0.103 (with benefit interactions, Column (4)) and not statistically significant. A similar phenomenon occurs in the grandchildren count regressions, where the marginal effect of an additional average grandchild is -3.827 (Column (5)) to -3.987 (Column (6)) and significant at the 10% level, but changes to 2.454 and 1.392 (Columns (7) and (8)) with the cohort fixed effects.

Table 3.12 also shows the marginal effect of a reduction by $100 in the early retirement penalty (SSB62-SSB65), a $100 increase in the PIA, and a $100 increase in the credit for delaying retirement. Comparing the marginal effects in Table 3.12's Columns (2) and (6) with the benefit level coefficients in Table 3.11's Columns (2) and (5), the effect of the early retirement penalty and the PIA grow in absolute value when their interactions with the grandchildren measures are accounted for. The marginal effect of the PIA is almost 1-for-1, where a $100 increase in the PIA will decrease the LFP rate by 0.92 points. The marginal effect of reducing the early retirement penalty is -0.59, which is also larger in absolute value than the Table 3.11 Column (2) on the penalty of -0.43. However, the marginal effect of the delayed retirement benefit is actually weaker than the results in Table 3.11 would indicate, where instead of being 0.84 points, when grandchildren interactions are accounted for, a $100 increase in the credit would increase the LFP rate by 0.84 to 0.86 points.

Surprisingly, the benefit values exhibit a similar pattern to the grandparent variables, losing significance when the birth cohort fixed effects are added. The significant interactions between retirement eligibility, the benefit levels, and the grandparent measures even with 4 year birth cohort fixed effects does not appear to be driven by a strong marginal effect of the benefit levels or the grandparent measures themselves. Overall, the strongest results are on the interactions between grandparenthood, retirement eligibility, and retirement benefits. Thus, grandchildren's greatest impact does seem to be concentrated around the retirement decision, but that the marginal effects of either the Social Security benefits or the grandparent measure are specification-sensitive.

3.5 Counterfactual Simulations

Given that changes in grandparenthood do drive labor force attachment decisions around the two retirement "checkpoints", but only ambiguously prior to age 62, how did the shrinking in extended families seen since the 1960's and 1970's change overall labor force participation? To put some context on the results in Tables 3.10 and 3.12, I simulate older worker's labor force participation from 1962 onward using 4 different scenarios:

1. **No Baby Boom**: I assume that the post-WWII "boom" never happened, so that the birth rate was essentially unchanged from 1939 to 1965.
2. **No *Roe***: I assume that abortion was never nationally legalized, and extend the birth rates observed in 1970-1972 outwards to the present.
3. **Ultra Low Fertility**: I assume that the birth rate for the last 100 years has been the same as the minimum one observed, which nationally was 2015's value of 12.4.
4. **Ultra High Fertility**: I assume that the birth rate for the last 100 years has been the same as the maximum one observed, which nationally was 1957's value of 24.9.

To put the greatest possible weight on the influence of grandchildren, I use the specifications without 4 year birth cohort fixed effects, but not the interactions with Social Security benefit levels, corresponding to Columns (2) and (8) in Table 3.11. If by using the results "friendliest" to grandparenthood's LFP impact, I fail to show that trends in grandchildren made little impact in older worker's labor force attachment overall, then I can conclude that grandchildren's aggregate impact is not very material to understanding these LFP trends.

Figures 3.5a and 3.5b show the counterfactual simulation results for men 55-69. Figure 3.5a shows the fraction of men aged 55-69 who were grandfathers from 1962-2015, their observed

LFP rate, the model's predicted LFP rate, and then the predicted LFP rate for the 4 counterfactual scenarios for historical fraction grandparent values. Figure 3.5b is structured similarly. Both figures show that, absent the Baby Boom, the LFP rate would have been much 5-7 points higher through 1970, before slowly converging to the model's predicted LFP rate in 1994. Surprisingly, the *Roe* scenario shows that at least in this crude counterfactual, the fall in fertility after *Roe* matters little for the LFP rate for older workers. The greatest gap is observed only in the most recent year of data, where the observed LFP rate is 4 points higher than the simulated LFP rate without *Roe* in the average grandchild count version and 6 points higher in the fraction grandparent version.

The ultra-low and ultra-high scenarios reveal two interesting LFP trends. Sustained high fertility would not have yielded a substantially different LFP rate for older workers until after 1985 - nearly 30 years after the 1957 fertility peak. Conversely, a sustained ultra low fertility rate would have risen the LFP rate by 14 points in the mid-1960's in the fraction grandparent version or 13 points in the average grandchild count version, but this difference would have narrowed to only 3 points in either version by 2015. However, if 1950's style fertility had never ended, then the LFP rate in this group would be about 10 points lower than it is today. These contrasts suggest that the grandparenthood elasticity of labor supply was likely bigger in the 1960's than it is today, possibly because other forces not previously at play are now more strongly shaping older worker's labor force attachment decisions.

One phenomenon clearly demonstrated by the fertility scenarios is that altering the grandparenthood assumptions does not meaningfully change the fall and rise pattern in older men's LFP during this period. In keeping with the results from Table 3.11, there may be age subsets that are more strongly driven by changes in grandparenthood and are more sensitive to assumption changes. Figures 3.6-3.9 show the same trends as Figure 3.5 by age groups 50-54, 55-61, 62-64, and 65-69. In the 50-54 group (Figure 3.6) and the 55-61 group (Figure 3.7), the LFP gap between the ultra high and ultra low scenarios ranges steadily between

8-10 points throughout the entire time period and regardless of whether I use the fraction grandparent (Figures 3.6a and 3.6a) and average grandchild count versions (Figures 3.6b and 3.6b).

The 50-54 age group is characterized by gently falling LFP from the early 1960's onwards, and while different grandchildren trends may have modestly changed this group's LFP levels, they would not have reversed any trends. The 55-61 age group, on the other hand, saw a sharp fall in their LFP rate around 1970 that would have been even steeper from peak to trough in lower fertility scenarios, without being meaningfully different in the trend lines. In all cases, their LFP rate flattened out after 1985 and in each fertility scenario, continuing changes in grandparenthood do not seem to have changed their behavior. For these groups, grandparenthood changes mostly the LFP level, but in spite of continuing declines in the number of grandchildren, do not seem to be driving major LFP trends in these groups.

The 62-64 and 65-69 age groups have an altogether different relationship to grandparenthood. Firstly, the rise in LFP in older workers is mostly driven by improvement in labor force attachment by workers 62 and older. Secondly, the grandparenthood effect is most robust for the interactions with Social Security eligibility. It is therefore unsurprising that the LFP gap between the ultra high and ultra low fertility scenarios is bigger in Figures 3.8 and 3.9 than in Figures 3.6 and 3.7. More surprisingly, however, is that Figure 3.8 shows that if fertility had remained at Baby Boom levels, or even more modestly at pre-*Roe* levels, the rise in the LFP rate in the 62-64 age group from 1994 to 2009 would largely not have occurred. In the ultra-high fertlity scenario, the LFP rate in this group would have remained at about 45% (with minor fluctuations) from 1990 to the present. In the No *Roe* scenario, it would have hovered around 50%. The low-fertility regime seen since the 1970's does seem to be playing a role in keeping older workers in the workforce for this age group for at least these specifications.

All of the simulations for the 65-69 age group in Figure 3.9 have a post-1994 rise in the LFP rate, and by about the same amount. The key phenomenon here is illustrated in 3.9a, where the many extra men who became grandfathers during the Baby Boom masked to some extent a strong labor participation desire by those 65-69. In other graphs, the absence of the Baby Boom would have pushed up the LFP rate by about 7-8 points in 1962, the year closest to the Baby Boom's peak. But in the 65-69 group, the LFP rate would have been about 15 points higher when accounting for extra men who were grandfathers than otherwise would have been.[24] If births had been stuck in a permanently low-fertility regime, the LFP rate among 65-69 year old men would have been 26 points higher in 1962. Without the Baby Boom, the 65-69 year old LFP rate would have declined from 52% in the late 1960's to 26% by 1994. This is a remarkable drop, but not one that can be explained by grandparenthood trends alone, because even assuming the Baby Boom didn't occur, the fall and rise still occurred.

The 62-64 and 65-69 year old simulations may be overstating the impact of grandchildren because the generating process does not account for interactions between grandparenthood and Social Security benefit changes that also occurred during this period and were found in Table 3.11 to be statistically significant. Thus, I regenerate Figures 3.8 and 3.9 using the specifications given in Table 3.11's Columns (3) and (9), with the results shown in Figures 3.10 and 3.11b. Now with the benefit-level interactions, the ultra-high and ultra-low fertility gap is much smaller, and even in the medium-high and ultra-high fertility scenarios, the 62-64 age group would have experienced a post-1994 rise in their LFP rate. In fact, even in the ultra-high fertility scenario, the LFP rate in this group would only have been about 6 points lower than what was observed in 2015, the lowest fertility year. The implication is that the increase in the early retirement penalty effectively disincentivized labor force exit in this age group so that even if extended families were larger, grandfathers would still have been more likely to delay retirement from the 1990's onwards.

[24]The grandchild count simulations in Figure 3.9b show an effect similar to the others.

Lastly, the counterfactuals show that certain stylized facts about the LFP rate are unrelated to grandchildren trends. No matter the grandchild scenario and both for the 62-64 and 65-69 group, 1994 remains as the nadir for older men's LFP rate, although different grandchild levels would have changed what that nadir was. The greatest spread in the counterfactual scenarios occurs among the 65-69 year old group, which Tables 3.8 and 3.11 agree appears to be the group most prone to factoring in grandchildren into their labor force attachment decisions.

3.6 Robustness Checks

Because many legal changes were targeted at specific age groups, it introduces the possibility that the first stage estimates are being driven by age-by-year or state-by-age shocks to fertility coincident with the policy changes. This would violate the identification assumption and yield biased results. An augmented panel fixed effects model is now estimated by adding 10 year age group-by-year and 10 year age group-by-state interaction terms.[25] I also estimate a model using 10 year birth cohort fixed effects interactions with the state and year fixed effects. In both cases, the instrument retains reasonably large Kleibergen-Paap Wald F-statistics, but the second-stage results claim that grandfathers are disproportionately more likely to return to the labor force at full retirement than the grandchildless. This implausible result in the face of other evidence may be driven by implicit small cell sizes created by the augmented fixed effects.

I also turned age into a fourth order polynomial and reestimate (3.1)-(3.3) and present the final results in Table 3.13. Grandmother's response to the marginal grandchild becomes more modest, with or without the retirement eligibility interactions, but otherwise remains

[25] Specifically, the age groups are indicators for the daughter or daughter-in-law being 0-10, 11-20, 21-30, 31-40, 41-50, and 51 and over. 5 year age bands might be more ideal, but the PSID is not a broad enough sample to credibly estimate such a model.

the same in sign and significance. Their response to being grandmothers, however, reverses itself. Now, the model claims grandmothers are 17% *less* likely to be retired. However, the same results say that before age 62, grandmothers are 14.6% less likely to be retired, but then become 27.7% more likely to be retired, with a non-significant response to full retirement and being a grandmother. This inconsistency in the results suggests that a the fourth order polynomial in age is possibly overfitting the model.

Likewise, almost all grandfather status results become insignificant after adding the cubic and quartic in age. It is worth noting that the only statistically significant response here does also say that grandfatherhood's effect is labor force detachment, in that grandfathers work 514 fewer hours at early retirement. Their response to the marginal grandchild essentially becomes non-significant, except that the do remain more likely to retire: each grandchild now makes grandfathers 9% more likely to be retired. Including retirement eligibility interactions creates a series of inconsistent results, which both suggest that grandfathers retire more but work less in response to the marginal grandchild, and then become only 0.005% more likely to retire when the marginal grandchild arrives at full retirement, while also working slightly (+5 hours) more.

All in all, expanding the age polynomial essentially affirms the grandmother results, but paints an internally contradictory picture for grandfathers. On balance, these results do suggest that most likely, grandparenthood lowers men's labor force attachment, but I believe the results in Table 3.5 are a more accurate reflection of the labor force dynamics involved because the results are much more internally consistent.

Lastly, I interacted linear time trends with education measures used in the CPS state panel, to capture the impact of different returns to education levels over time. While the marginal effects of grandparenthood in the model without birth cohort fixed effects became somewhat larger and more significant, the rest of the model estimates are essentially no different. Since I cannot estimate a first-stage equation directly when I add these effects, because the number of instruments comes to exceed the number of clusters, I retain my original model reported in Tables 3.9 and 3.11 that does not include the time trended education shares.

3.7 Discussion and Conclusion

This paper explores grandparents' labor force attachment by testing various labor outcomes on two ways grandchildren may influence their grandparents: having the status of grandparent and the total number of grandchildren. Several questions were asked and answered to study this intergenerational dynamic. The first was whether grandparents had a labor supply response to the arrival and presence of grandchildren. The second was how to characterize this response and report it by grandparent sex. The third was to examine how grandchildren trends have altered older men's labor force attachment since the 1960's, and in particular, can the rise and fall in grandchildren partially explain the fall and rise in the labor force participation rate among men 55 and older.

The answers presented here are that there is a labor force response, and that grandparents of both sexes tend to lower their labor force attachment as their extended families grow. Compared to the grandchildless, both grandfathers and grandmothers have lower labor force attachment, with grandfathers being 19% more likely to be retired and to work 363 fewer hours a year. Grandmothers are 8.5% more likely to be retired and are 13% less likely to report having worked any hours in the previous year. Grandfathers and grandmothers do react differentially to how their families grow, however, with grandfather's being 27% more likely to report being retired with a new grandchild, but only 15.9% more likely to be retired when the new grandchild arrives during early retirement eligibility, and just 5.4% more likely to be retired when the new grandchild arrives during full retirement eligibility. A similar, but not less statistically significant, pattern exists for their annual hours worked and their propensity to be in the labor force. Grandmothers, by contrast, become more likely to lessen their labor supply when the marginal grandchild arrives at successive retirement eligibility stages.

In the limited literature on this question, grandfathers' roles have not received as much attention as grandmothers', but the findings here indicate that this is an oversight, because grandfathers are robustly shown to decrease their labor force response across multiple specifications. I therefore also examine how the post-Baby Boom increase in grandchildren affected the fall in older men's labor force participation seen between 1970 and 1994, and then it's subsequent rise from 1994 to the Great Recession's advent by using the CPS to estimate representative yearly samples of male workers, aged 50-69. I find that, as in the PSID results, there are significant interactions between grandchildren and retirement eligibility, and I also find significant interactions between grandchildren, retirement eligibility, and retirement benefits, but these interactions do not explain why the labor force participation rate fell, even though it did coincide with "peak" grandparenthood. Across all alternative historical grandparenthood scenarios, I find that the fall and rise would have occurred regardless, although the prevalence of grandchildren do substantially change the labor force participation rate levels, particularly among those 62 and older.

There are some currently unaddressed issues that future drafts will consider. These include using the PSID to study how grandparents react specifically to the youngest grandchildren, in an effort to better understood what role grandchildren play in the grandparent's lifecycle. Another extension is using the CPS data to take a pass at seeing what role the decline in grandparenthood has played in the rising labor force participation among older female workers. Lastly, given the significance of Social Security and grandparenthood interactions, exploring other simulations that make changes to both parameters. This will also account for how policymakers make decisions about Social Security, such that a permanent boom or permanent bust in fertility would almost certainly provoke a response by policymakers to adjust Social Security accordingly.

This paper uses estimates adult children's fertility impact on grandparents' labor market outcomes using exogenous variation in access to reproductive technology. Grandchildren's influence is examined through the effect of being a grandparent, the total number of grandchildren per household, and the number of young grandchildren per household. Although much of the policy variation in the fertility instrument is historical, the results complement existing findings on grandparent aid to new parents by determining that the time transfers from grandparent to adult child are likely not coming out of their labor supply. This paper is also the first in the literature to document a grandfather labor market response to grandchildren, an important contribution because policymakers wishing to model how new generations affect old ones need a clear understanding of how both older worker types may change their behavior. Lastly, it advances the literature on the labor force participation trends among older workers by explicitly accounting for the role of extended families, in issue otherwise ignored until this article.

Tables

Summary Statistics for Grandparent Sample, 1968

TABLE 3.1

	Grandfathers		Grandmothers	
	Mean	**St Dev**	**Mean**	**St Dev**
Age When First PSID Grandchild Born	50.28	6.91	46.06	7.11
Age When First Retired	59.29	7.81	57.55	8.93
Fraction Retired	10.2%	0.30	5.8%	23.5%
Annual Hours Worked	2,155.80	791.07	1,596.87	1,094.90
In Labor Force (ILS)	95.2%	0.21		
Age	38.33	8.10	35.93	7.33
Number of Children	3.19	2.02	3.36	2.16
Family Income (2016 $)	$82,023.41	78,584.04	$61,038.37	48,462.59
Fraction White	55.5%	0.50	42.0%	0.49
Fraction Black	41.3%	0.49	51.9%	0.50
Fraction Married	93.5%	0.25	72.0%	0.45
Highest Education Level Attained:				
Primary School	34.3%	0.47	27.1%	0.44
Secondary School	21.8%	0.41	30.4%	0.46
HS Grad	22.2%	0.42	29.3%	0.46
Some College	7.9%	0.27	5.4%	0.23
Bachelors or More	8.7%	0.28	3.1%	0.17
Grandparents	1,712		2,373	
Adult Children	5,465		7,970	
Observations	160,761		221,818	

Table 3.1 shows selected summary statistics for in-sample grandfather and grandmother characteristics in 1968. Sample is individuals who were aged 22-54 in 1968 and who have at least one child in the PSID.

Time Transfers (in Hours) By Number of Grandchildren

TABLE 3.2

Number of Grandchildren ⇒	**0**	**1**	**2**	**3**	**Any Child**
Mother's Parents					
Married Grandparents	21.48	58.25	52.01	51.35	55.40
Grandfather Remarried	3.51	12.88	9.63	7.38	9.23
Grandmother Remarried	40.18	59.78	19.28	42.88	43.64
Single Grandfathers	25.02	19.42	5.18	2.14	13.79
Single Grandmothers	21.24	83.17	25.21	44.13	45.31
All Mother's Parents	**19.35**	**57.31**	**27.07**	**36.31**	**37.98**
Father's Parents					
Married Grandparents	21.32	16.79	51.33	68.22	47.24
Grandfather Remarried	3.04	7.21	2.39	2.12	4.21
Grandmother Remarried	27.64	173.41	11.48	35.03	64.59
Single Grandfathers	1.92	1.49	6.06	2.31	5.76
Single Grandmothers	16.04	36.22	10.15	2.22	14.20
All Father's Parents	**15.59**	**33.42**	**18.90**	**21.41**	**22.89**
All Grandparents	**34.95**	**90.72**	**45.97**	**57.72**	**60.86**

Table 3.2 shows average time transfer in hours from parents to adult children, separated out on the basis of how many children the adult children have.
Source: 2013 PSID Family Rosters and Transfers.

TABLE 3.3

Grandchild Measure ⇓	Grandfathers: Retired	Grandfathers: Hours Worked	Grandfathers: In Labor Force	Grandmothers: Retired	Grandmothers: Hours Worked	Grandmothers: Non-Zero Working Hrs
	(b/se)	(b/se)	(b/se)	(b/se)	(b/se)	(b/se)
Grandparent Status Regressions						
Without interactions						
$\mathbb{1}\{Grandparent\}$	0.083***	-142.01***	-0.017	0.037**	-41.69*	-0.022
	(0.012)	(25.69)	(0.010)	(0.013)	(24.10)	(0.013)
Adj. R^2	0.69	0.63	0.64	0.71	0.49	0.47
With interactions						
$\mathbb{1}\{Grandparent\}$	0.83***	-150.90***	-0.020*	0.035***	-45.07*	-0.024*
	(0.013)	(26.82)	(0.011)	(0.013)	(24.86)	(0.013)
$*\mathbb{1}\{Early\ SS\ Elig\}$	0.009	-21.75	0.008	0.019	86.40	0.043
	(0.038)	(82.456)	(0.038)	(0.028)	(61.42)	(0.040)
$*\mathbb{1}\{Full\ SS\ Elig\}$	0	116.28	0.031	0.005	-6.97	-0.01
	(0.012)	(90.90)	(0.034)	(0.024)	(80.58)	(0.058)
Adj. R^2	0.69	0.63	0.64	0.69	0.49	0.47
N	44,317	44,317	44,365	62,074	62,074	62,074
Child Count Regressions						
Without interactions						
Child Count	0.035***	-41.49***	-0.011***	0.019***	-41.40***	-0.018***
	(0.005)	(7.46)	(0.004)	(0.005)	(8.11)	(0.005)
Adj. R^2	0.71	0.64	0.64	0.72	0.52	0.49
With interactions						
Child Count	0.053***	-67.22***	-0.23***	0.011*	-29.48***	-0.011**
	(0.006)	(9.38)	(0.005)	(0.005)	(8.84)	(0.005)
$*\mathbb{1}\{Early\ SS\ Elig\}$	-0.017**	6.19	0.006	0.035***	-33.38**	-0.013*
	(0.007)	(13.262)	(0.008)	(0.006)	(14.41)	(0.007)
$*\mathbb{1}\{Full\ SS\ Elig\}$	-0.043***	69.01***	0.030**	0.011	-24.257**	-0.017**
	(0.012)	(21.67)	(0.011)	(0.005)	(11.83)	(0.007)
Adj. R^2	0.71	0.64	0.64	0.72	0.52	0.49
N	130,948	130,948	129,478	180,405	180,405	180,405

Table 3.3 shows the panel fixed effects regression estimates for the $\mathbb{1}\{\text{Grandparent}\}$ measure of the effect being a grandparent on labor force attachment outcomes. Likewise, the coefficients for *Child Count* measure the marginal effect of an additional grandchild on each outcome. All regressions include individual-level and state-by-year fixed effects. Regressions are weighted with the core family sampling weights provided by the PSID, adjusted for the number of adult children each grandparent has. Robust standard errors clustered at the state level are reported.

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

First-Stage Estimates of Grandchild Measures from PSID

TABLE 3.4

	$\mathbb{1}\{$**Grandparent**$\}$		**Child Count**	
	Grandfathers	Grandmothers	Grandfathers	Grandmothers
Access Policy ⇓	(1)	(2)	(3)	(4)
Pill Access	(b/se)	(b/se)	(b/se)	(b/se)
No Lag	-0.057**	-0.040*	-0.158***	-0.148***
	(0.016)	(0.016)	(0.021)	(0.024)
Lag (t-1)	0.043***	0.043***	-0.01	-0.002
	(0.011)	(0.012)	(0.014)	(0.014)
Lag (t-2)	0.016**	0.021***	-0.02	-0.003
	(0.008)	(0.007)	(0.013)	(0.015)
Lag (t-3)	0.028***	0.029**	0.016	0.029**
	(0.009)	(0.009)	(0.012)	(0.013)
Lag (t-8)	0.127***	0.128**	0.207***	0.225***
	(0.018)	(0.009)	(0.028)	(0.025)
Abortion Access				
No Lag	-0.060	-0.044	-0.104***	-0.106***
	(0.045)	(0.050)	(0.019)	(0.023)
Lag (t-1)	0.008	0.004	-0.022**	-0.037***
	(0.013)	(0.009)	(0.012)	(0.012)
Lag (t-2)	0.002	0.005	-0.029**	-0.021
	(0.014)	(0.013)	(0.011)	(0.013)
Lag (t-3)	0.024	0.012	0.000	-0.011
	(0.017)	(0.018)	(0.017)	(0.019)
Lag (t-8)	0.137***	0.024**	0.180***	0.174***
	(0.037)	(0.009)	(0.052)	(0.039)
Early Abortion Nearby State				
No Lag	-0.056*	-0.052	-0.166***	-0.138***
	(0.029)	(0.029)	(0.038)	(0.030)
F-Statistic	36.72	163.83	114.23	69.90
N	44,246	61,960	130,948	180,405

Table 3.4 shows the first-stage regression results estimating Equation (3.4) for the grandchild measure *Child Count* and Equation (3.5) for $\mathbb{1}\{Grandparent\}$. "Pill Access" and "Abortion Access" are the treatment variables for whether the adult daughter had access to the reproductive technology in year t. "Early Abortion Nearby State" is a dummy for whether individual i's 1968 state was within 250 miles of a state with legalized abortion prior to *Roe*. "F-Statistic" is the cluster-robust Kleibergen-Paap Wald rk F-statistic. Regressions include policy lags out to the 8th lag but coefficients are not reported. All regressions are weighted with the core family sampling weights provided by the PSID, adjusted for the number of adult children each grandparent has.

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

2nd-Stage IV Results of Grandparents' Labor Response to Grandchildren

TABLE 3.5

Grandchild Measure ⇓	**Grandfathers** Retired	**Grandfathers** Hours Worked	**Grandfathers** In Labor Force	**Grandmothers** Retired	**Grandmothers** Hours Worked	**Grandmothers** Non-Zero Working Hrs
	(b/se)	(b/se)	(b/se)	(b/se)	(b/se)	(b/se)
Without interactions						
$\mathbb{1}\{Grandparent\}$	0.194***	-362.90**	-0.045	0.085*	-121.35	-0.131**
	(0.056)	(134.41)	(0.055)	(0.042)	(131.30)	(0.064)
With interactions						
$\mathbb{1}\{Grandparent\}$	0.233***	-409.179***	-0.065	0.134**	209.90	-0.161**
	(0.056)	(122.37)	(0.055)	(0.051)	(139.12)	(0.063)
$*\mathbb{1}\{Early\ SS\ Elig\}$	0.184	-553.42**	-0.215*	0.417**	-359.41	-0.135
	(0.131)	(263.90)	(0.125)	(0.169)	(245.34)	(0.122)
$*\mathbb{1}\{Full\ SS\ Elig\}$	0.065	226.71	-0.040	-0.398	444.49	0.405
	(0.384)	(369.09)	(0.185)	(0.297)	(719.23)	(0.303)
N	44,317	44,317	44,317	62,074	62,074	62,074
Without interactions						
Child Count	0.178***	-219.452***	-0.092***	0.217***	-431.57***	-0.191***
	(0.029)	(65.51)	(0.027)	(0.027)	(57.70)	(0.028)
With interactions						
Child Count	0.270***	-401.51***	-0.171***	0.244***	-499.139***	-0.220***
	(0.031)	(61.17)	(0.029)	(0.027)	(50.49)	(0.026)
$*\mathbb{1}\{Early\ SS\ Elig\}$	-0.111***	118.86	0.067	0.091***	-108.678***	-0.042***
	(0.011)	(13.262)	(0.008)	(0.011)	(25.09)	(0.012)
$*\mathbb{1}\{Full\ SS\ Elig\}$	-0.216***	328.51***	0.082**	0.078***	-184.49***	-0.086***
	(0.025)	(70.53)	(0.031)	(0.014)	(24.26)	(0.014)
N	130948	130948	129478	180405	180405	180405

Table 3.5 shows the second-stage regression estimates of grandparents' labor force characteristics. The coefficients for $\mathbb{1}\{Grandparent\}$ measure the effect being a grandparent. The coefficients for *Child Count* measure the marginal effect of an additional grandchild on each outcome. Second stage estimates are from Equation (3.4) and the grandparent flag from Equation (3.5). *Retired* is 1 if the grandmother reports being retired, 0 else. *Hours Worked* is the annual number of hours worked. *In Labor Force* is an indicator for whether the grandfather is in the labor force, and *Non-Zero Working Hrs* is an indicator for whether the grandmother reported non-zero working hours in year t. All regressions are weighted with the core family sampling weights provided by the PSID, adjusted for the number of adult children each grandparent has where necessary.

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

Panel Regression of Older Men's National Labor Force Participation Rates

TABLE 3.6

	Birth Cohort Specification			
	Time Trends	1 Year FE	2 Year FE	4 Year FE
	(1)	(2)	(3)	(4)
	(b/se)	(b/se)	(b/se)	(b/se)
Without interactions				
% Grandparent	-0.688***	-34.691***	-0.702*	-0.914***
	(0.080)	(10.902)	(0.415)	(0.206)
With interactions				
% Grandparent	-0.402***	6.123	-0.473	-0.656***
	(0.086)	(10.117)	(0.373)	(0.194)
% Grandparent*$\mathbb{1}\{\textit{Early SS Elig}\}$	-0.124***	-0.119***	-0.126***	-0.128***
	(0.033)	(0.027)	(0.027)	(0.028)
% Grandparent*$\mathbb{1}\{\textit{Full SS Elig}\}$	-0.709***	-0.702***	-0.698***	-0.702***
	(0.053)	(0.048)	(0.046)	(0.046)
Without interactions				
Grandchild Count	-8.509***	-428.985***	-8.687*	-11.302***
	(0.990)	(134.819)	(5.128)	(2.551)
With interactions				
Grandchild Count	-3.590***	-35.891	-4.541	-6.901***
	(1.127)	(123.524)	(4.656)	(2.442)
Grandchild Count*$\mathbb{1}\{\textit{Early SS Elig}\}$	-2.219***	-1.562***	-1.690***	-1.794***
	(0.389)	(0.321)	(0.319)	(0.326)
Grandchild Count*$\mathbb{1}\{\textit{Full SS Elig}\}$	-8.757***	-7.133***	-7.367***	-7.639***
	(0.612)	(0.539)	(0.540)	(0.544)
N	4121	4120	4121	4121

Table 3.6 shows the panel fixed effects regression estimates of labor force characteristics, including grandparenthood status, on national labor force participation rates. The coefficients for *% Grandparent* measure the effect of an additional 1 point in the fraction of individuals in a given age-sex-year-birth year-educational attainment group cell who are grandparents. The coefficients for *Grandchild Count* measure the marginal effect of an additional grandchild on national LFP. These are second stage estimates after instrumenting for $Spouse_InLF$ in Equation (3.6) as described in Section 3.4.1. All rate and fractional variables are multiplied by 100, so that that coefficients can be interpreted as how many points the LFP rate changes per one unit change in the variable.

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

Extended Estimates from National LFP Panel Regression

TABLE 3.7

	% Grandfathers				Grandchild Count			
	No Elig. Interactions		Elig. Interactions		No Elig. Interactions		Elig. Interactions	
	Time Trends	4 Year FE	Time Trends	4 Year FE	Time Trends	4 Year FE	Time Trends	4 Year FE
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
SSB65	-0.581***	-0.392**	-0.533***	-0.434***	-0.581***	-0.392**	-0.468***	-0.358**
	(0.147)	(0.160)	(0.143)	(0.141)	(0.147)	(0.160)	(0.145)	(0.141)
(SSB62-SSB65)	-0.297*	-0.062	-0.067	-0.171	-0.297*	-0.062	0.025	-0.189
	(0.176)	(0.223)	(0.169)	(0.196)	(0.176)	(0.223)	(0.173)	(0.198)
(SSB70-SSB65)	0.583***	0.237*	0.520***	0.318***	0.583***	0.237*	0.451***	0.333***
	(0.113)	(0.129)	(0.109)	(0.116)	(0.113)	(0.129)	(0.113)	(0.117)
Lifetime Avg. Monthly Earnings	-0.083***	-0.164***	-0.071***	-0.127***	-0.083***	-0.164***	-0.048***	-0.094***
	(0.017)	(0.021)	(0.017)	(0.019)	(0.017)	(0.021)	(0.017)	(0.019)
Monthly Disability Benefit	-0.158***	-0.205***	0.067	0.004	-0.158***	-0.205***	0.052	-0.068
	(0.045)	(0.046)	(0.075)	(0.072)	(0.045)	(0.046)	(0.073)	(0.070)
Log Predicted Wage	2.033**	1.648**	-0.057	-0.295	2.033**	1.648**	0.038	-0.073
	(0.924)	(0.824)	(0.899)	(0.746)	(0.924)	(0.824)	(0.911)	(0.759)
% Spouse in LF	0.452***	0.168**	0.556***	0.200***	0.452***	0.168**	0.591***	0.227***
	(0.060)	(0.072)	(0.059)	(0.064)	(0.060)	(0.072)	(0.062)	(0.067)
% Married	0.150***	0.125***	0.160***	0.134***	0.150***	0.125***	0.174***	0.147***
	(0.034)	(0.030)	(0.035)	(0.027)	(0.034)	(0.030)	(0.036)	(0.028)
% Previously Married	0.167***	0.119***	0.182***	0.129***	0.167***	0.119***	0.200***	0.146***
	(0.040)	(0.034)	(0.041)	(0.031)	(0.040)	(0.034)	(0.042)	(0.032)
% Veteran	-0.020*	-0.040***	-0.030***	-0.044***	-0.020*	-0.040***	-0.032***	-0.040***
	(0.011)	(0.011)	(0.011)	(0.010)	(0.011)	(0.011)	(0.011)	(0.010)
% Black	-0.054	0.058	-0.075	0.067*	-0.054	0.058	-0.049	0.091**
	(0.046)	(0.044)	(0.046)	(0.041)	(0.046)	(0.044)	(0.048)	(0.041)
% In Bad Health	-0.682***	-0.728***	-0.606***	-0.657***	-0.682***	-0.728***	-0.630***	-0.678***
	(0.027)	(0.024)	(0.029)	(0.023)	(0.027)	(0.024)	(0.029)	(0.023)
Kleibergen-Paap F-statistic	166.81	93.60	153.53	94.53	166.81	93.60	147.00	93.86

Table 3.7 shows the remaining coefficient estimates for the panel fixed effects regression estimates of labor force characteristics shown in Table 3.6.

Panel Regression of LFP With Grandchildren, Eligibility, and Benefit Levels Interactions

TABLE 3.8

	% Grandfathers		Grandchild Count	
	Time Trends	4 Year FE	Time Trends	4 Year FE
	(1)	(2)	(3)	(4)
GP_Measure	-0.542***	-0.726***	-5.246***	-8.338***
	(0.093)	(0.196)	(1.211)	(2.482)
<u>*Interacted with:*</u>				
$*\mathbb{1}\{\textit{Early SS Elig}\}$	-0.078**	-0.088***	-2.064***	-1.697***
	(0.033)	(0.028)	(0.383)	(0.325)
*(SSB62-SSB65)	-0.025***	-0.020***	-0.280***	-0.251***
	(0.003)	(0.003)	(0.052)	(0.046)
$*\mathbb{1}\{\textit{Full SS Elig}\}$	-0.641***	-0.648***	-8.522***	-7.390***
	(0.055)	(0.048)	(0.620)	(0.567)
*SSB65	0.006***	0.004**	0.084***	0.091***
	(0.002)	(0.002)	(0.028)	(0.024)
*(SSB70-SSB65)	-0.003	-0.004*	-0.067*	-0.088***
	(0.003)	(0.002)	(0.036)	(0.033)
SSB65	-0.532***	-0.385***	-0.513***	-0.391***
	(0.146)	(0.144)	(0.150)	(0.147)
(SSB62-SSB65)	0.231	0.204	0.260	0.155
	(0.176)	(0.201)	(0.180)	(0.203)
(SSB70-SSB65)	0.493***	0.334***	0.493***	0.385***
	(0.119)	(0.120)	(0.122)	(0.122)
Lifetime Avg. Monthly Earnings	-0.069***	-0.136***	-0.052***	-0.114***
	(0.017)	(0.019)	(0.017)	(0.019)
Monthly Disability Benefit	0.115	-0.023	0.107	-0.036
	(0.087)	(0.077)	(0.086)	(0.077)
Log Predicted Wage	0.177	-0.142	0.204	0.035
	(0.890)	(0.742)	(0.908)	(0.760)
% Spouse in LF	0.560***	0.230***	0.597***	0.275***
	(0.061)	(0.072)	(0.063)	(0.068)
Kleibergen-Paap F-statistic	140.97	82.19	137.88	94.17

Table 3.8 shows the results of estimating Equation (3.6) while extending the interactions between *GP_Measure* and the retirement eligibility indicators to Social Security benefit levels. See Table 3.6 for more information.

State Panel Regression of Selected Older Men's Characteristics on LFP Rates

TABLE 3.9

	% Grandparent						Grandchild Count					
	Time Trends			4 Year FE			Time Trends			4 Year FE		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
GP_Measure	-0.02	-0.04	-0.10	0.08	0.09	0.03	-0.21	-0.86	-1.64	1.12	0.90	0.22
	(0.10)	(0.08)	(0.09)	(0.12)	(0.09)	(0.09)	(1.34)	(1.06)	(1.11)	(1.54)	(1.17)	(1.20)
Interacted with:												
$\mathbb{1}\{Early\ SS\ Elig\}$		0.06	0.03		0.02	-0.02		1.18	0.45		0.60	-0.23
		(0.07)	(0.07)		(0.07)	(0.07)		(0.75)	(0.66)		(0.79)	(0.71)
(SSB62-SSB65)			-0.02***			-0.02***			-0.32***			-0.32***
			(0.00)			(0.00)			(0.05)			(0.06)
$\mathbb{1}\{Full\ SS\ Elig\}$		0.01	-0.08		-0.02	-0.13		0.63	-0.83		0.08	-1.79
		(0.12)	(0.12)		(0.12)	(0.12)		(1.41)	(1.36)		(1.44)	(1.37)
SSB65			0.01***			0.02***			0.17***			0.22***
			(0.00)			(0.00)			(0.03)			(0.03)
(SSB70-SSB65)			-0.01*			-0.01***			-0.06			-0.11***
			(0.00)			(0.00)			(0.04)			(0.03)
SSB65	-1.29***	-1.32***	-1.53***	-0.03	-0.09	-0.32	-1.29***	-1.25***	-1.49***	-0.03	-0.05	-0.30
	(0.19)	(0.16)	(0.14)	(0.20)	(0.19)	(0.19)	(0.19)	(0.16)	(0.13)	(0.20)	(0.19)	(0.19)
(SSB62-SSB65)	-1.08***	-1.07***	-0.84***	-0.16	-0.18	0.19	-1.08***	-1.06***	-0.87***	-0.16	-0.17	0.16
	(0.12)	(0.12)	(0.15)	(0.23)	(0.23)	(0.24)	(0.12)	(0.12)	(0.15)	(0.23)	(0.23)	(0.23)
(SSB70-SSB65)	1.22***	1.20***	1.19***	0.27*	0.25	0.29*	1.22***	1.18***	1.20***	0.27*	0.26*	0.30*
	(0.11)	(0.11)	(0.09)	(0.16)	(0.16)	(0.16)	(0.11)	(0.12)	(0.09)	(0.16)	(0.15)	(0.16)
Lifetime Avg. Monthly Earnings	0.16***	0.15***	0.17***	0.04	0.04	0.04	0.16***	0.15***	0.17***	0.04	0.04	0.04
	(0.02)	(0.01)	(0.01)	(0.03)	(0.03)	(0.03)	(0.02)	(0.01)	(0.01)	(0.03)	(0.03)	(0.03)
Monthly Disability Benefit	-0.56***	-0.43***	-0.18*	-0.76***	-0.57***	-0.26***	-0.56***	-0.54***	-0.24**	-0.76***	-0.70***	-0.33***
	(0.06)	(0.10)	(0.09)	(0.08)	(0.10)	(0.09)	(0.06)	(0.10)	(0.11)	(0.08)	(0.10)	(0.11)
Log Predicted Wage	0.07	0.03	-0.05	-0.07	-0.07	-0.19	0.07	0.02	-0.09	-0.07	-0.09	-0.23
	(0.51)	(0.51)	(0.52)	(0.50)	(0.50)	(0.50)	(0.51)	(0.51)	(0.52)	(0.50)	(0.50)	(0.50)
% Spouse in LF	0.07***	0.07***	0.07***	0.06***	0.06***	0.06***	0.07***	0.06***	0.06***	0.06***	0.06***	0.06***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Adj R^2	0.88	0.88	0.88	0.88	0.88	0.88	0.88	0.88	0.88	0.88	0.88	0.88
N	27,415	27,415	27,415	27,415	27,415	27,415	27,415	27,415	27,415	27,415	27,415	27,415

Table 3.9 shows the panel fixed effects regression estimates of labor force characteristics, including grandparenthood status and average grandchild count, on state-level labor force participation rates. Robust standard errors are clustered at the state level. See Table 3.6 for more information.

First-Stage Estimates of Grandchild Measures from CPS

TABLE 3.10

	% Grandparent		Grandchild Count	
	(1)	(2)	(3)	(4)
Years of Exposure				
⇓ Since Age 33 To:				
Pill 21 Plus	0.02	0.12	0.00	0.01
	(0.13)	(0.15)	(0.01)	(0.01)
Pill 18-20	-0.12	-0.10	-0.01	-0.01
	(0.10)	(0.10)	(0.01)	(0.01)
Abortion 21 Plus	-0.93***	-0.81***	-0.07***	-0.06***
	(0.24)	(0.21)	(0.02)	(0.02)
Abortion 18-20	0.12	0.04	0.01	0.00
	(0.27)	(0.22)	(0.02)	(0.02)
Prohibition	0.24	0.41*	0.02	0.03*
	(0.27)	(0.24)	(0.02)	(0.02)
Spouse Birth Year	-0.01	0.03	-0.00	0.00
	(0.02)	(0.02)	(0.00)	(0.00)
Spouse Birth Year^2	-0.01***	-0.01*	-0.00***	-0.00
	(0.00)	(0.00)	(0.00)	(0.00)
Adj. R^2	0.98	0.98	0.98	0.98
N	27,415	27,415	27,415	27,415
4-Year Birth Cohort FE's	N	Y	N	Y

Table 3.10 shows the first-stage regression results estimating Equation (3.8) for the grandchild measures *Grandchild. Count* and for % *Grandparent*. All regressions include age, state, and year fixed effects, and the birth year squared, and are estimated with heteroskedastic-robust standard errors clustered at the state level.

Second-Stage Estimates of Selected Older Men's Characteristics on LFP Rates

TABLE 3.11

	% Grandparent						Grandchild Count					
	Time Trends			4 Year FE			Time Trends			4 Year FE		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
GP_Measure	-0.31*	-0.18	-0.32**	-0.04	0.30	0.11	-4.17*	-3.06	-4.18*	-0.54	3.50	1.56
	(0.16)	(0.15)	(0.15)	(0.24)	(0.19)	(0.22)	(2.14)	(2.31)	(2.24)	(3.23)	(2.94)	(3.19)
<u>*Interacted with:*</u>												
$\mathbb{1}\{Early\ SS\ Elig\}$		-0.03	0.05		-0.11**	-0.05		-0.12	0.21		-1.02**	-0.80*
		(0.05)	(0.06)		(0.05)	(0.05)		(0.44)	(0.47)		(0.41)	(0.41)
(SSB62-SSB65)			-0.03***			-0.02***			-0.44***			-0.35***
			(0.00)			(0.00)			(0.06)			(0.06)
$\mathbb{1}\{Full\ SS\ Elig\}$		-0.43***	-0.29***		-0.51***	-0.37***		-5.07***	-4.01***		-6.09***	-5.07***
		(0.11)	(0.09)		(0.11)	(0.09)		(1.04)	(0.97)		(1.03)	(0.95)
(SSB65)			0.02***			0.02***			0.22***			0.23***
			(0.00)			(0.00)			(0.04)			(0.04)
(SSB70-SSB65)			-0.01**			-0.01***			-0.09			-0.13***
			(0.00)			(0.00)			(0.06)			(0.04)
SSB65	-0.71***	-0.68***	-1.13***	-0.02	-0.18	-0.44**	-0.71***	-0.72***	-1.12***	-0.02	-0.19	-0.42*
	(0.25)	(0.25)	(0.22)	(0.24)	(0.22)	(0.22)	(0.25)	(0.27)	(0.22)	(0.24)	(0.22)	(0.22)
(SSB62-SSB65)	-0.65***	-0.43***	-0.30*	-0.19	-0.24	0.16	-0.65***	-0.49***	-0.34**	-0.19	-0.25	0.15
	(0.13)	(0.14)	(0.16)	(0.22)	(0.23)	(0.25)	(0.13)	(0.13)	(0.16)	(0.22)	(0.23)	(0.25)
(SSB70-SSB65)	0.92***	0.92***	0.94***	0.30**	0.27*	0.35*	0.92***	0.93***	0.94***	0.30**	0.26*	0.33*
	(0.12)	(0.11)	(0.09)	(0.14)	(0.15)	(0.18)	(0.12)	(0.11)	(0.09)	(0.14)	(0.15)	(0.18)
Lifetime Avg. Monthly Earnings	0.10***	0.08***	0.11***	0.04	0.04	0.03	0.10***	0.08***	0.11***	0.04	0.04	0.04
	(0.02)	(0.03)	(0.02)	(0.03)	(0.03)	(0.03)	(0.02)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)
Monthly Disability Benefit	-0.59***	-0.28***	0.10	-0.75***	-0.36***	-0.06	-0.59***	-0.14***	0.10	-0.75***	-0.22***	-0.03
	(0.06)	(0.04)	(0.07)	(0.08)	(0.06)	(0.07)	(0.06)	(0.05)	(0.07)	(0.08)	(0.07)	(0.07)
% Spouse in LF	0.24***	0.24***	0.26***	0.10***	0.09***	0.13***	0.24***	0.25***	0.26***	0.10***	0.10***	0.13***
	(0.03)	(0.03)	(0.03)	(0.02)	(0.03)	(0.02)	(0.03)	(0.03)	(0.03)	(0.02)	(0.03)	(0.03)
Kleibergen-Paap rk LM	100.01	100.93	95.50	50.74	56.84	70.83	100.01	110.19	96.23	50.74	56.57	75.05

Table 3.11 shows the regression estimates after instrumenting for endogenous variables as described in Section 3.4.2.

Marginal Effects for Interacted Variables from Second-Stage Estimates

TABLE 3.12

	% Grandparent				Grandchild Count			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	(b/se)	(b/se)	(b/se)	(b/se)	(b/se)	(b/se)	(b/se)	(b/se)
GP_Measure	-0.252	-0.298*	0.212	0.103	-3.827	-3.987*	2.454	1.392
	(0.160)	(0.159)	(0.200)	(0.222)	(2.390)	(2.282)	(3.015)	(3.181)
SSB65		-0.922***		-0.221		-0.940***		-.228
		(0.232)		(0.222)		(0.241)		(0.224)
SSB62-SSB65		-0.587***		-0.062		-0.631***		-0.080
		(0.146)		(0.245)		(0.144)		(0.245)
SSB70-SSB65		0.843***		0.228		0.862***		0.228
		(0.108)		(0.165)		(0.111)		(0.165)
Birth Cohort Time Trends	Y	Y	Y	Y	Y	Y	Y	Y
4-Year Birth Cohort FE's	N	Y	N	Y	N	Y	N	Y

Table 3.12 shows the net effect of increasing either $GP_{M}easure$ by one unit, or increasing each of the three Social Security benefit measures by $100, derived from the results in Table 3.11.

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

PSID 2nd-Stage IV Results with Age as a 4th Order Polynomial

TABLE 3.13

Grandchild Measure ⇓	**Grandfathers** Retired	Hours Worked	In Labor Force	**Grandmothers** Retired	Hours Worked	Non-Zero Working Hrs
	(b/se)	(b/se)	(b/se)	(b/se)	(b/se)	(b/se)
<u>Without interactions</u>						
$\mathbb{1}\{Grandparent\}$	-0.082	85.565	0.073	-0.170***	185.421	-0.038
	(0.066)	(134.41)	(0.071)	(0.057)	(157.633)	(0.080)
<u>With interactions</u>						
$\mathbb{1}\{Grandparent\}$	-0.060	51.085	0.067	-0.146**	149.079	-0.037
	(0.071)	(160.206)	(0.070)	(0.062)	(158.911)	(0.077)
$*\mathbb{1}\{Early\ SS\ Elig\}$	0.152	-513.742**	-0.164	0.423**	-274.900	-0.047
	(0.128)	(252.756)	(0.119)	(0.157)	(248.000)	(0.116)
$*\mathbb{1}\{Full\ SS\ Elig\}$	0.173	60.119	-0.032	-0.048	44.84	0.241
	(0.238)	(328.165)	(0.180)	(0.203)	(569.093)	(0.220)
N	44,317	44,317	44,317	62,074	62,074	62,074
<u>Without interactions</u>						
Child Count	0.090***	-70.208	-0.043	0.145***	-331.378***	-0.138***
	(0.030)	(69.384)	(0.028)	(0.075)	(56.138)	(0.030)
<u>With interactions</u>						
Child Count	0.117***	-125.896*	-0.076**	0.145***	-326.736***	-0.132***
	(0.029)	(67.295)	(0.029)	(0.025)	(50.39)	(0.027)
$*\mathbb{1}\{Early\ SS\ Elig\}$	-0.026	-39.912	0.023	0.071***	-57.62**	-0.014
	(0.040)	(96.296)	(0.040)	(0.01)	(25.46)	(0.013)
$*\mathbb{1}\{Full\ SS\ Elig\}$	-0.112***	131.562*	0.040	0.042***	-64.00**	-0.02
	(0.025)	(73.223)	(0.033)	(0.014)	(27.18)	(0.016)
N	130,948	130,948	129,478	180,405	180,405	180,405

Table 3.13 shows the second-stage regression estimates of grandparents' labor force characteristics with including age as a fourth order polynomial. The coefficients for $\mathbb{1}\{Grandparent\}$ measure the effect being a grandparent. The coefficients for *Child Count* measure the marginal effect of an additional grandchild on each outcome. Second stage estimates are from Equation (3.4) and the grandparent flag from Equation (3.5). *Retired* is 1 if the grandmother reports being retired, 0 else. *Hours Worked* is the annual number of hours worked. *In Labor Force* is an indicator for whether the grandfather is in the labor force, and *Non-Zero Working Hrs* is an indicator for whether the grandmother reported non-zero working hours in year t. All regressions are weighted with the core family sampling weights provided by the PSID, adjusted for the number of adult children each grandparent has where necessary.

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

Figures

Figure 3.1: Civilian Labor Force Participation Rate for Workers 55 and Over, January 1948-April 2017

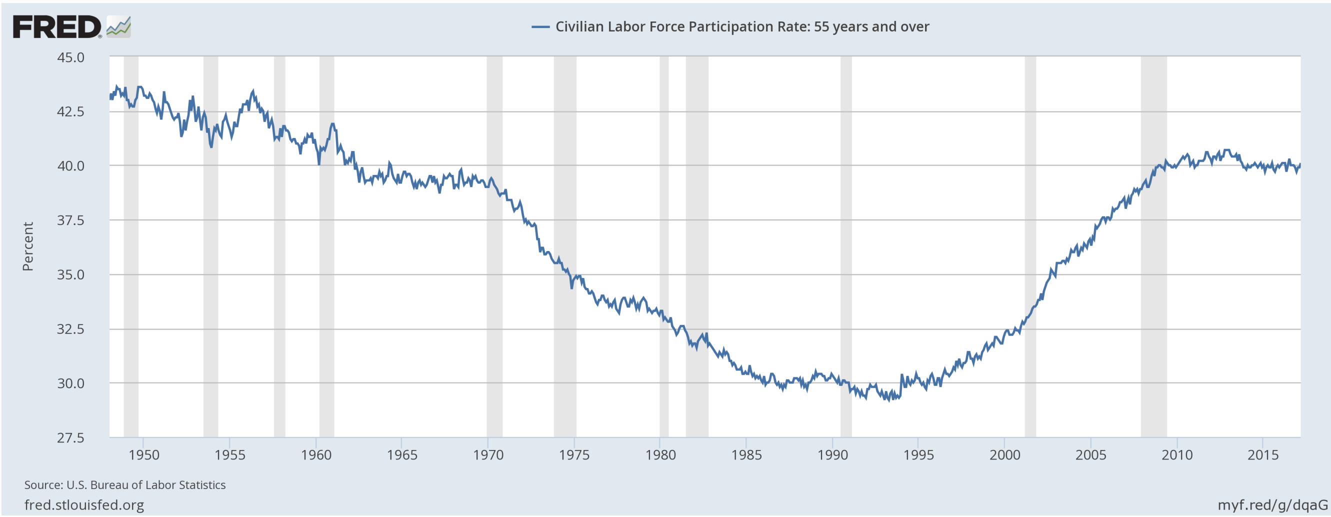

Figure 3.2: United States Birth Rate and Birth Counts, 1920 to 2014.

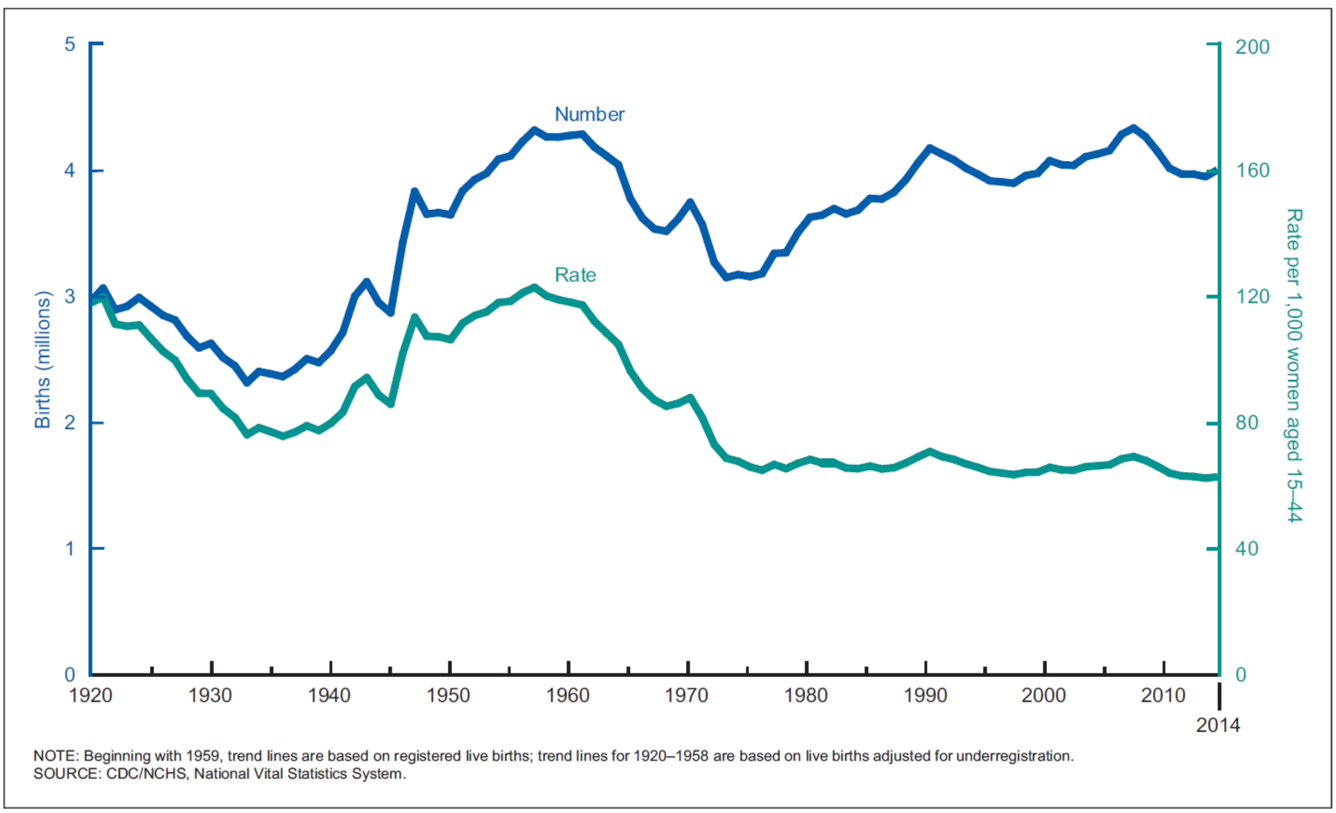

Figure 3.3: Median age of first marriage for men and for women from 1890 to 2016.

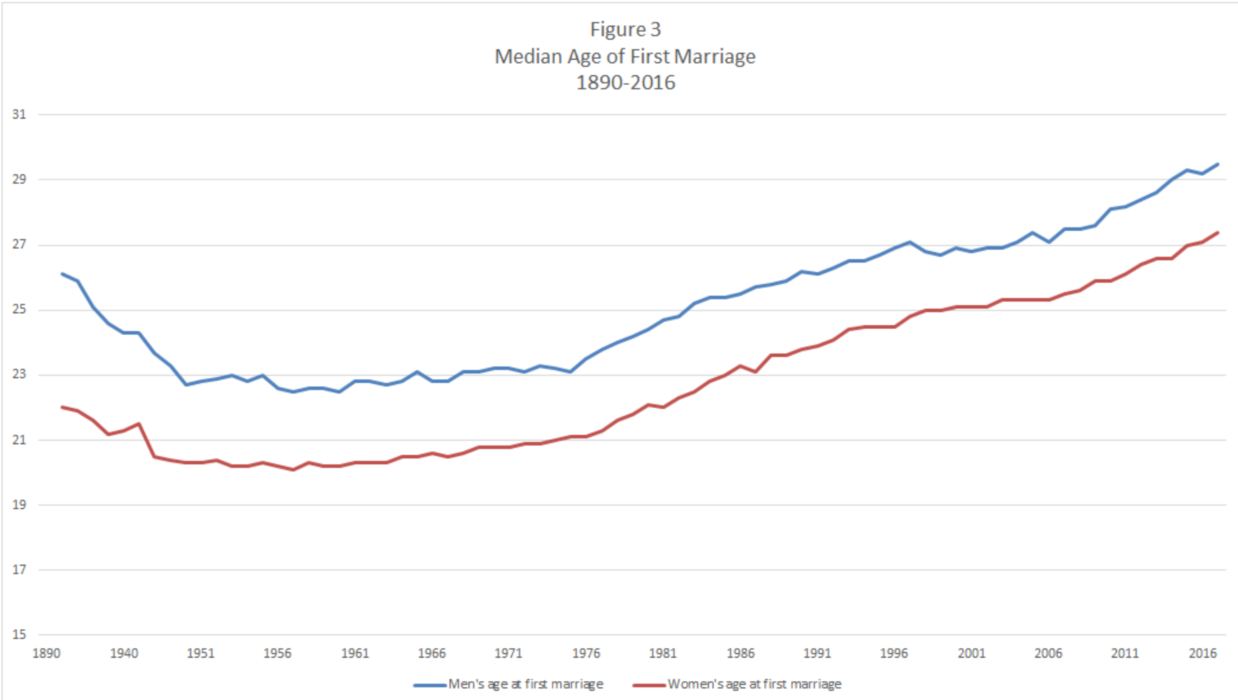

Sources: U.S. Census Bureau, Current Population Survey, March and Annual Social and Economic Supplements.

Figure 3.4: Older Men's Labor Force Participation Rates and Grandfather Characteristics, 1935-2015

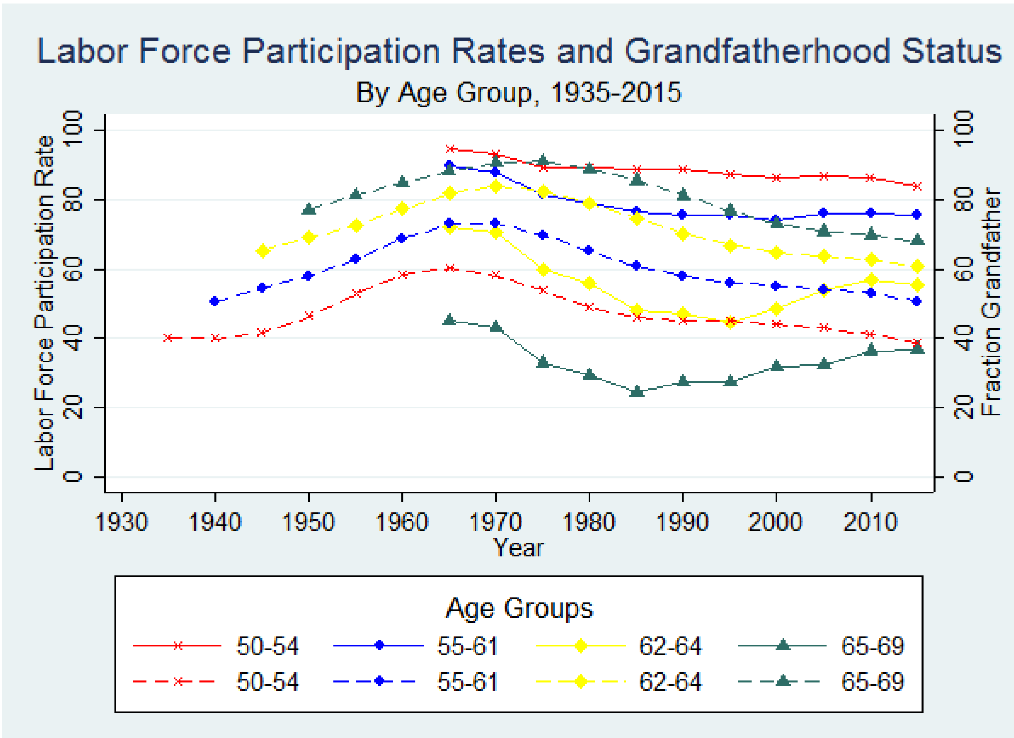

(a) Fraction of men who are grandparents in dashed lines and their labor force participation (LFP) rates in solid lines.

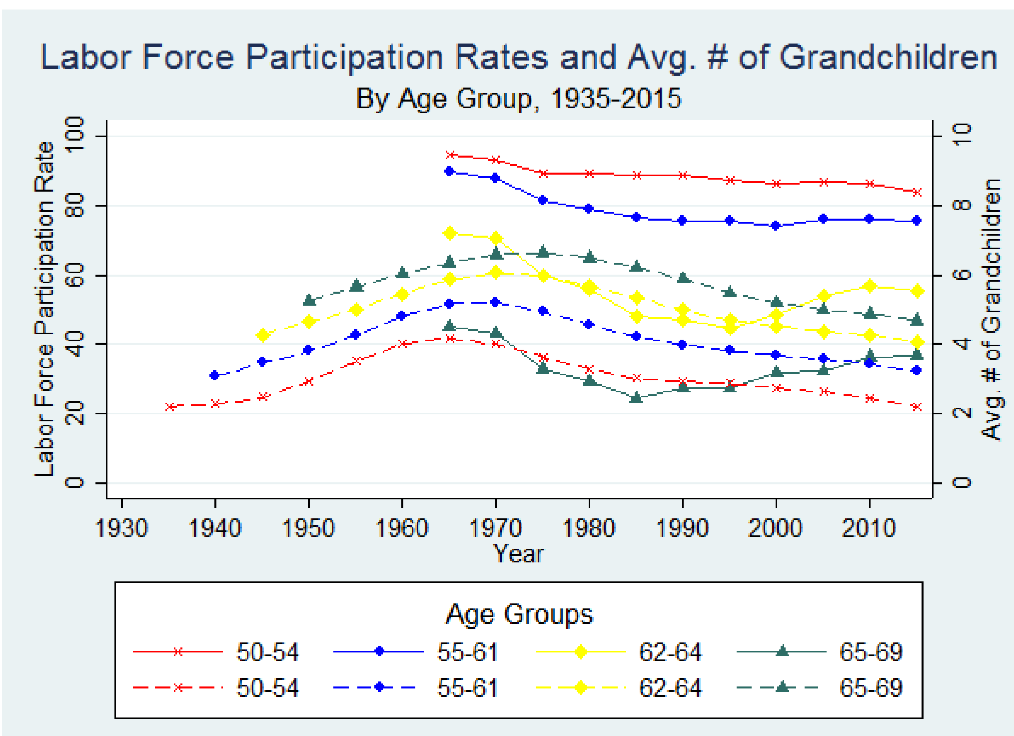

(b) Average number of grandchildren in dashed lines and older men's labor force participation (LFP) rates in solid lines.

Figure 3.5: Simulated labor force participation rates under different assumptions about fertility for ages 55-69.

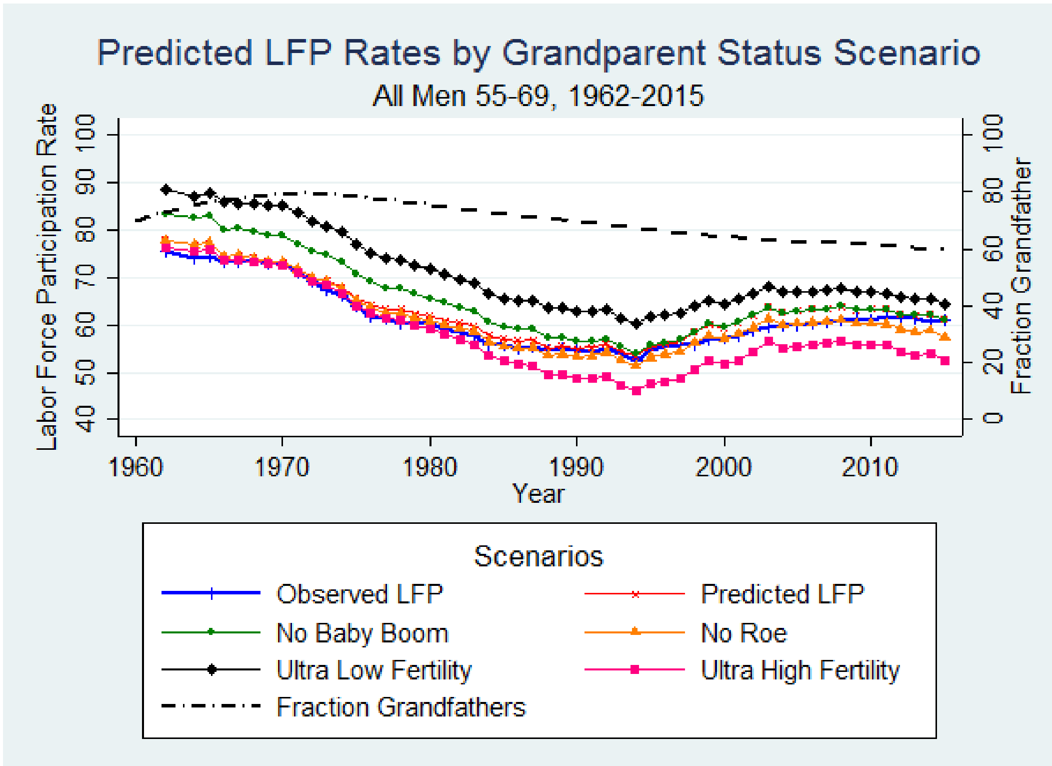

(a)

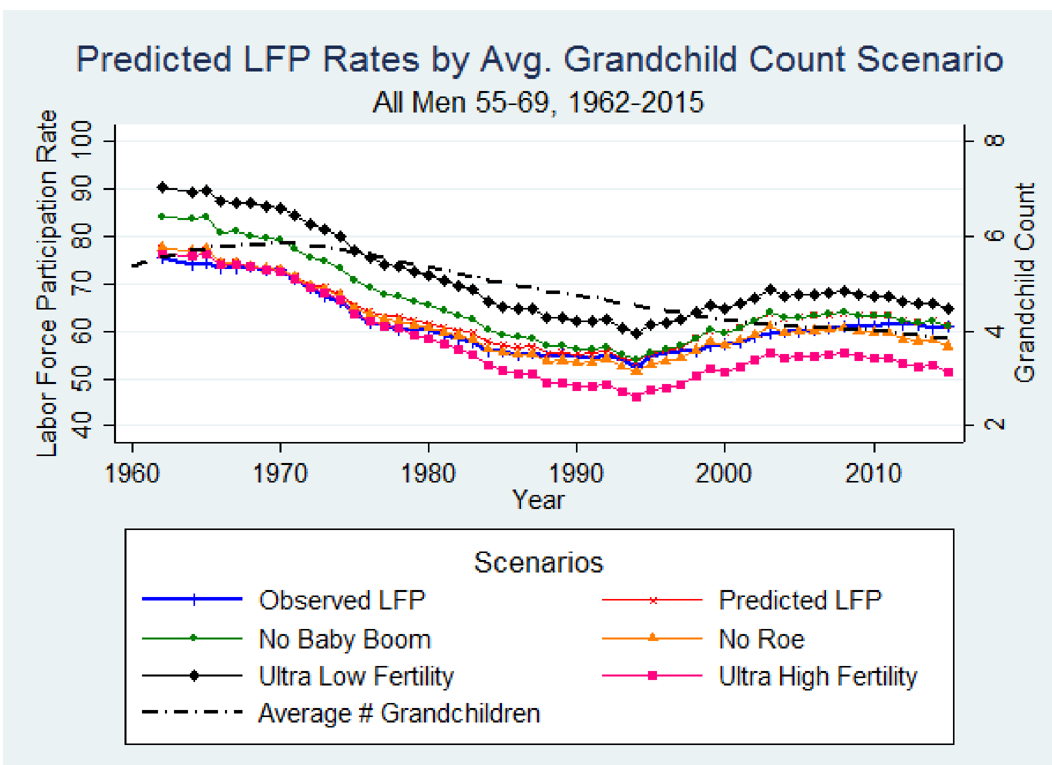

(b)

Figure 3.6: Simulated labor force participation rates under different assumptions about fertility for ages 50-54.

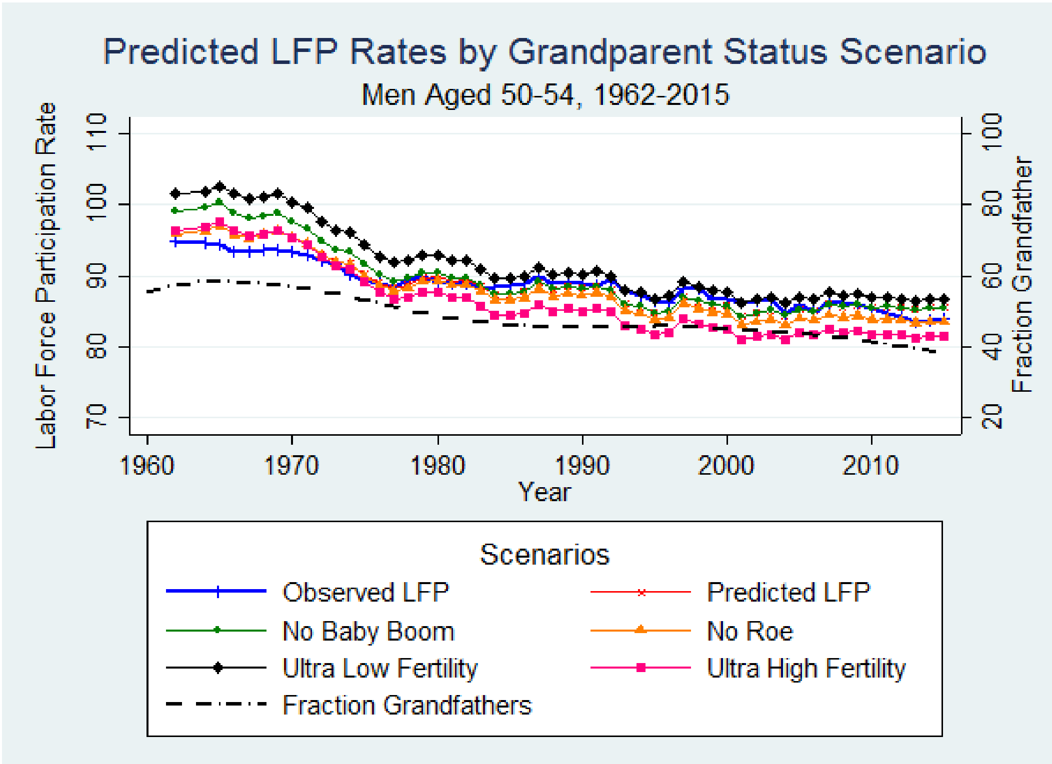

(a)

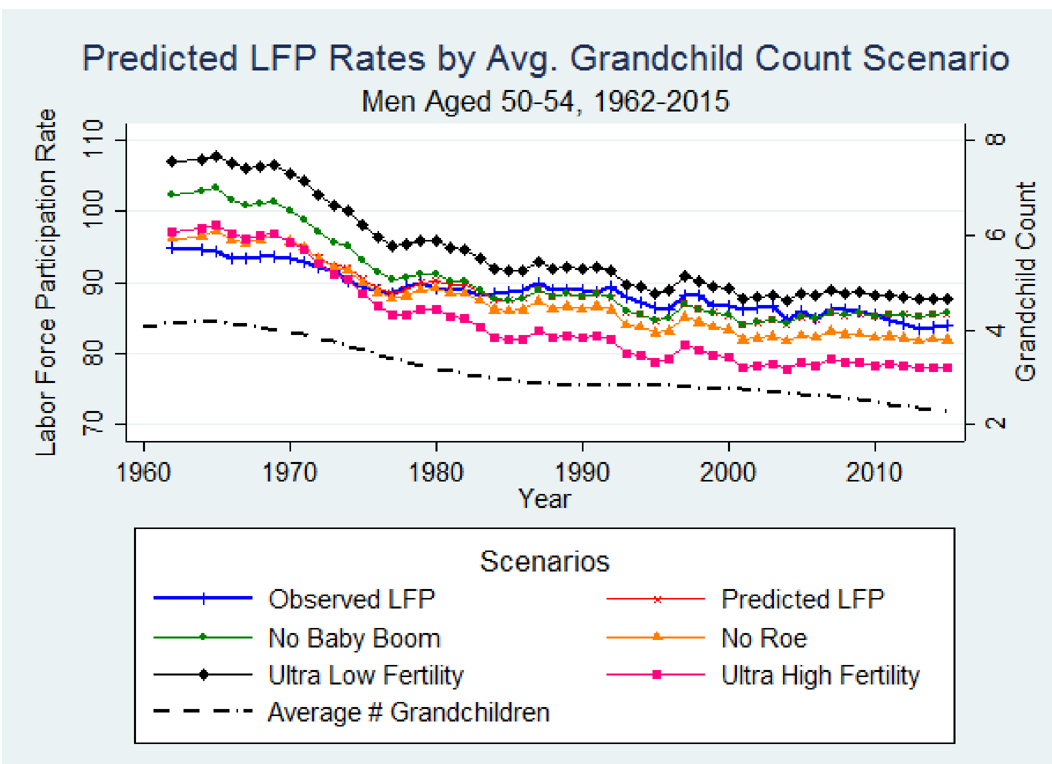

(b)

Figure 3.7: Simulated labor force participation rates under different assumptions about fertility for ages 55-61.

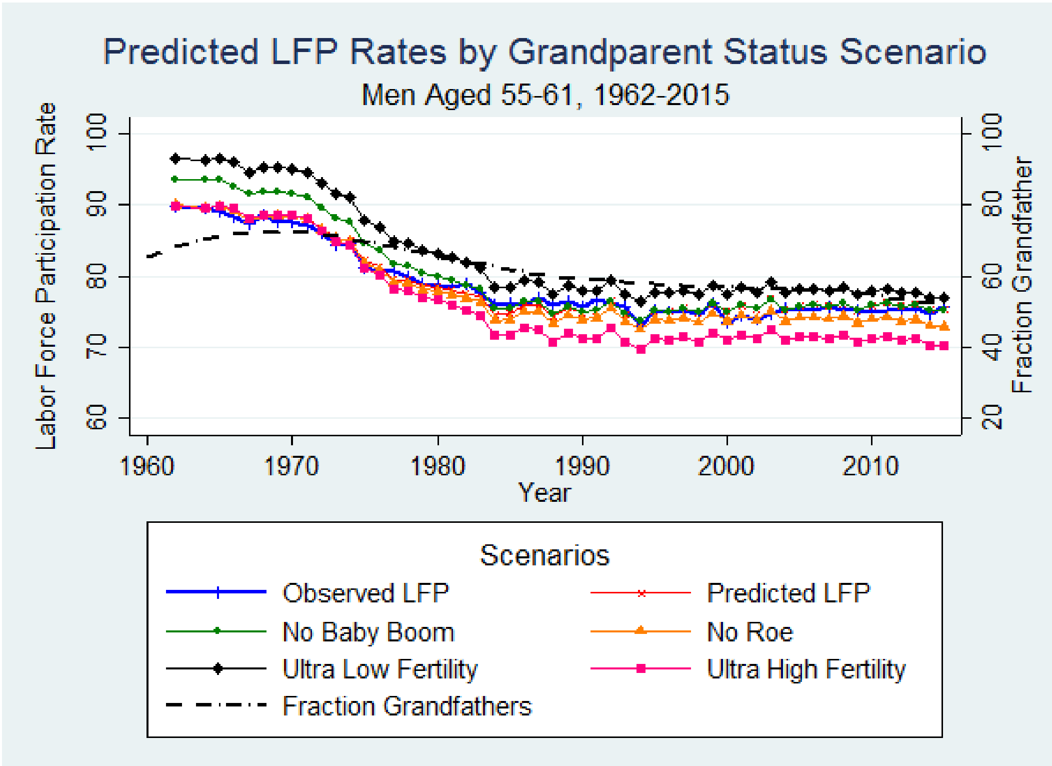

(a)

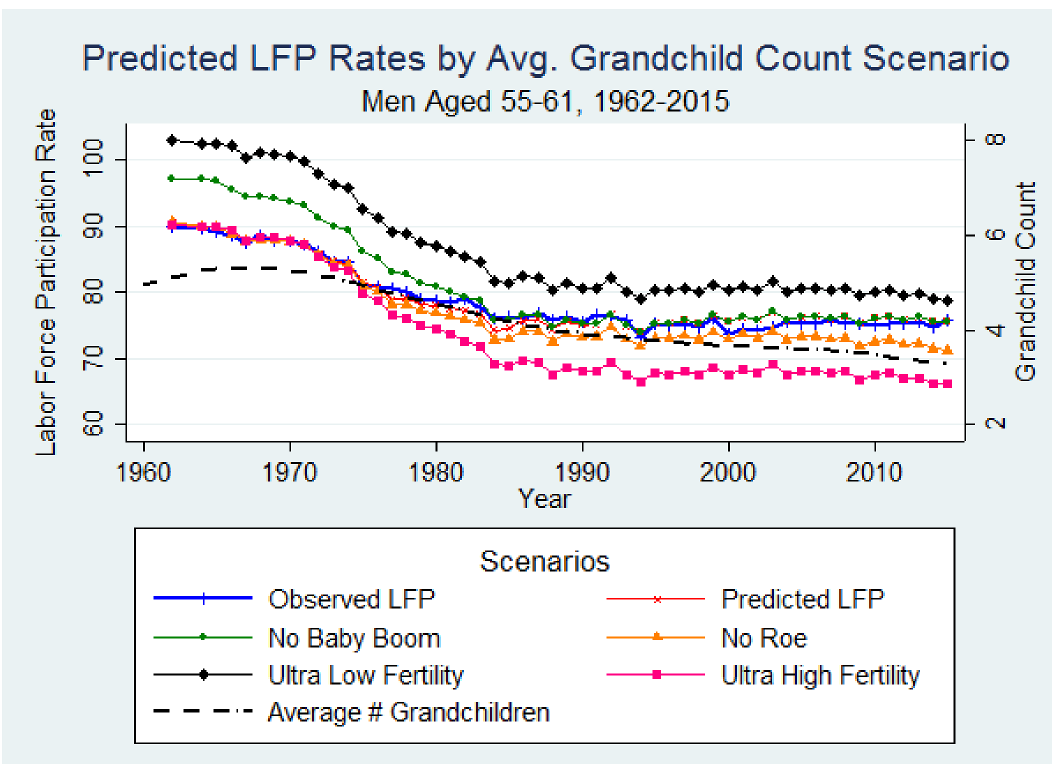

(b)

Figure 3.8: Simulated labor force participation rates under different assumptions about fertility for ages 62-64.

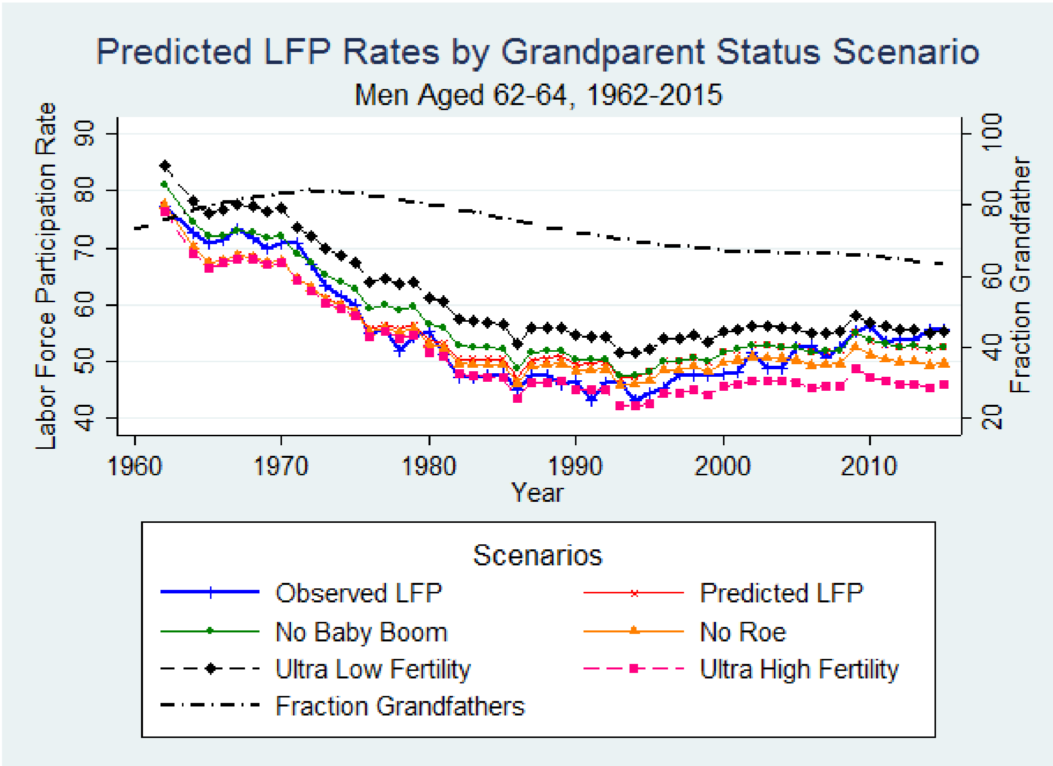

(a)

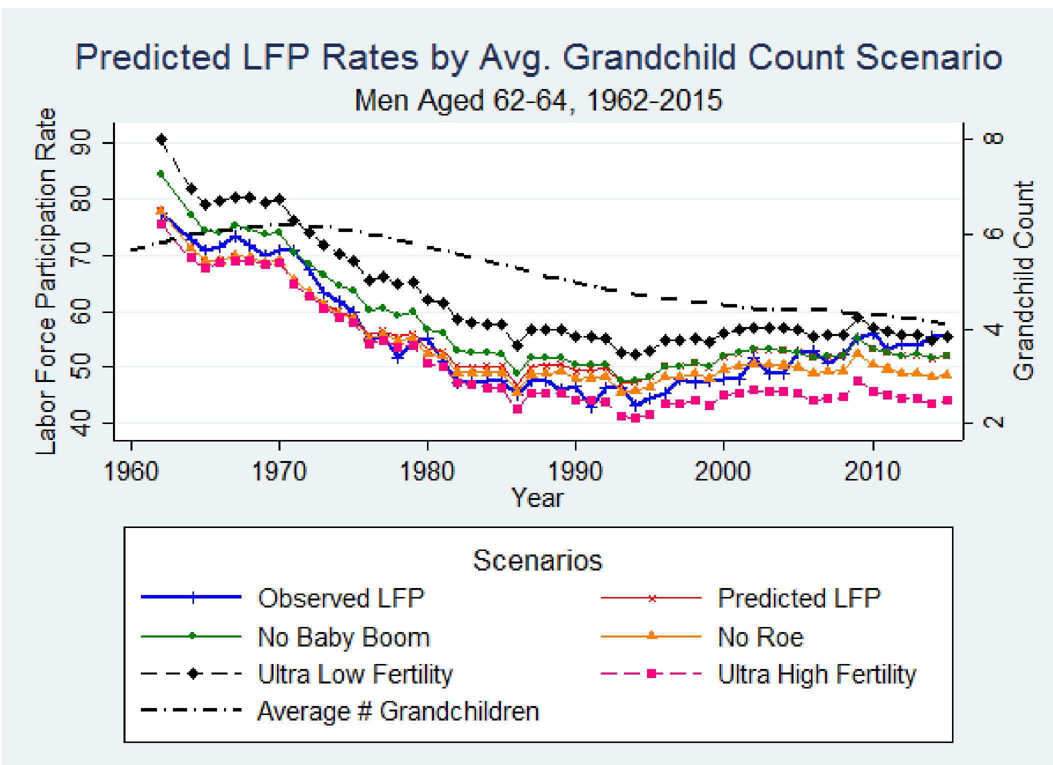

(b)

Figure 3.9: Simulated labor force participation rates under different assumptions about fertility for ages 65-69.

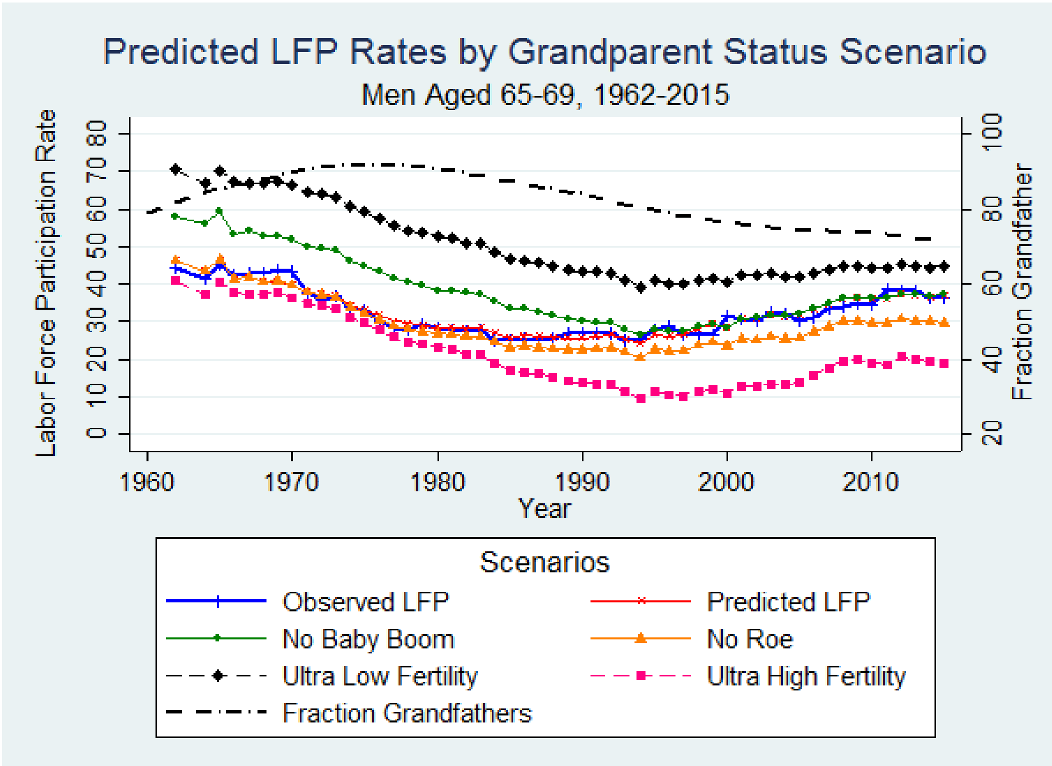

(a)

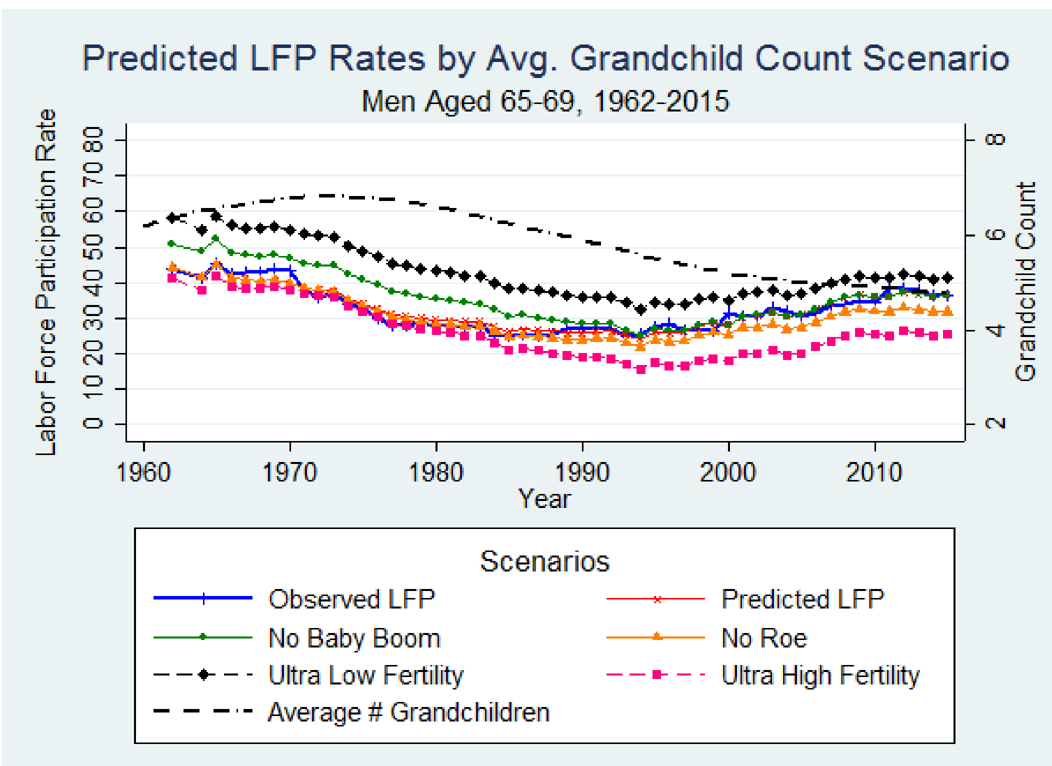

(b)

Figure 3.10: Simulated labor force participation rates, including interactions with Social Security benefit levels, under different assumptions about fertility for ages 62-64.

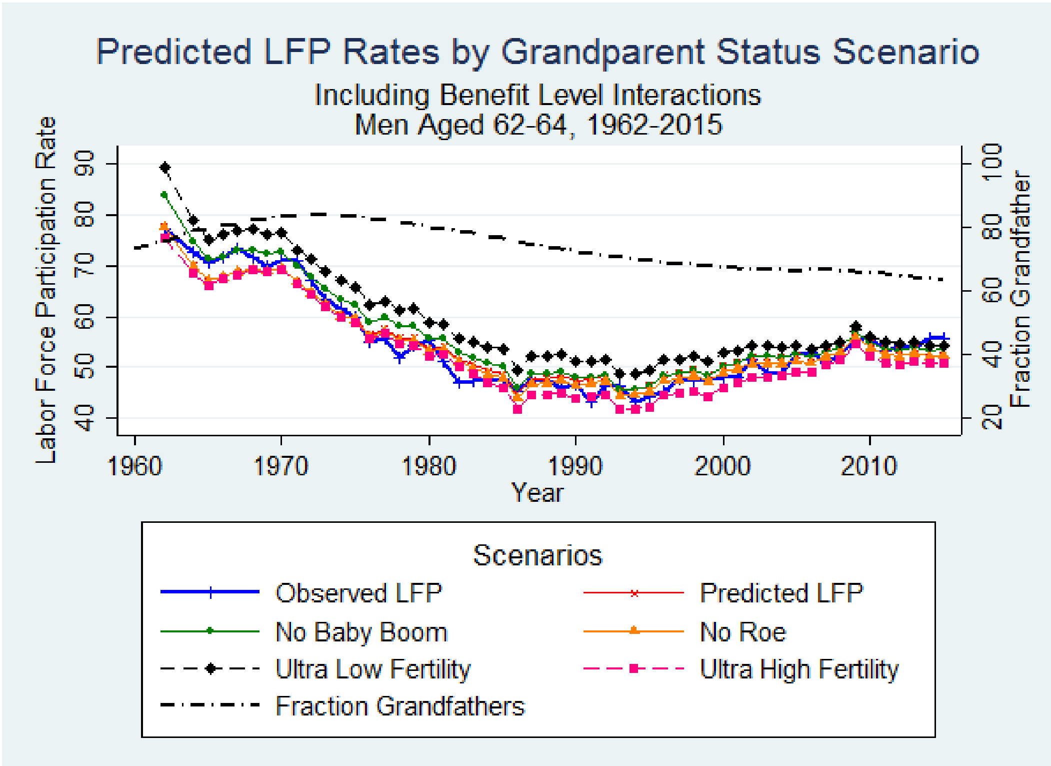

(a)

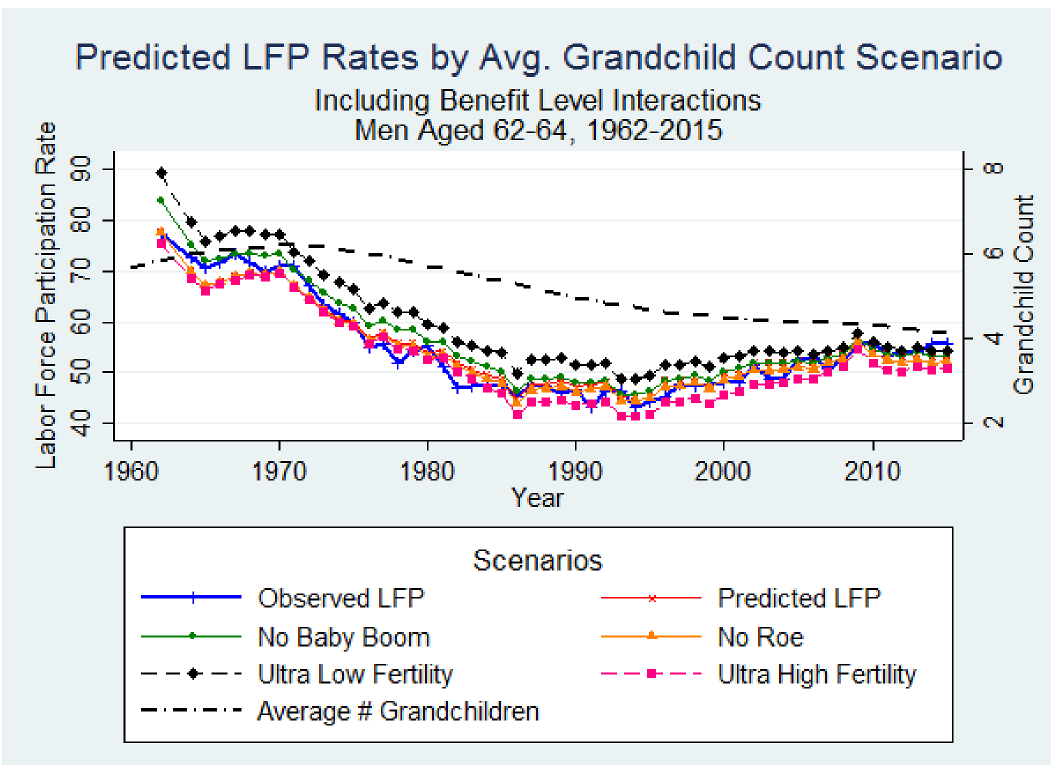

(b)

Figure 3.11: Simulated labor force participation rates, including interactions with Social Security benefit levels, under different assumptions about fertility for ages 62-64.

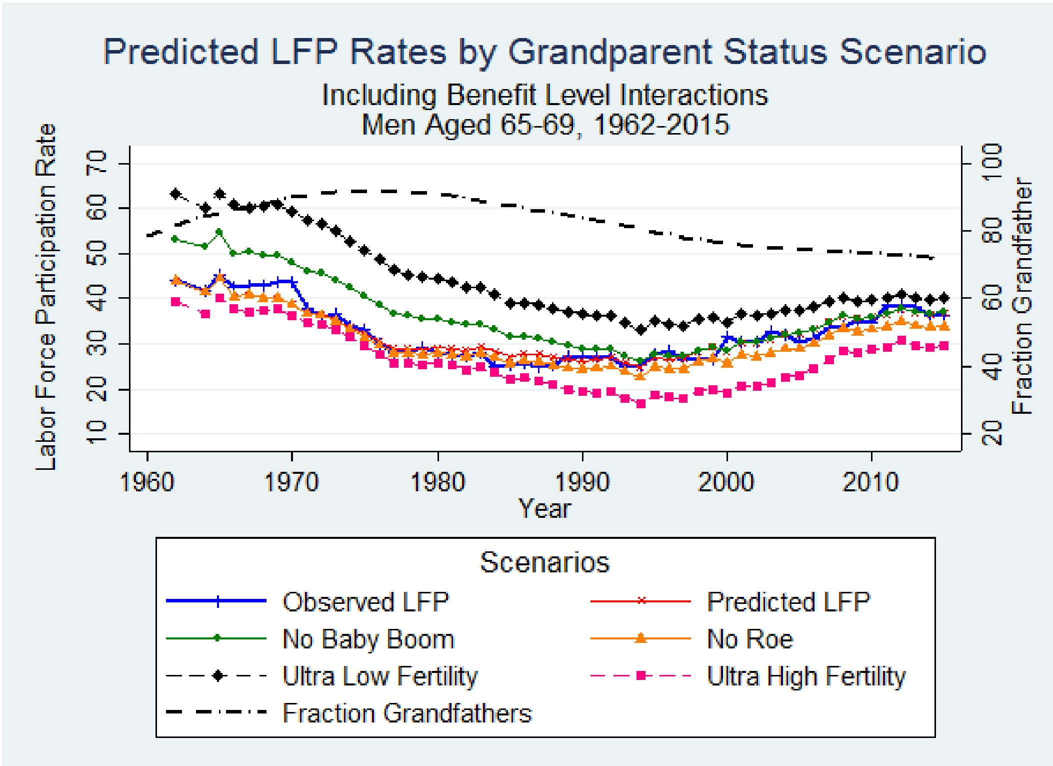

(a)

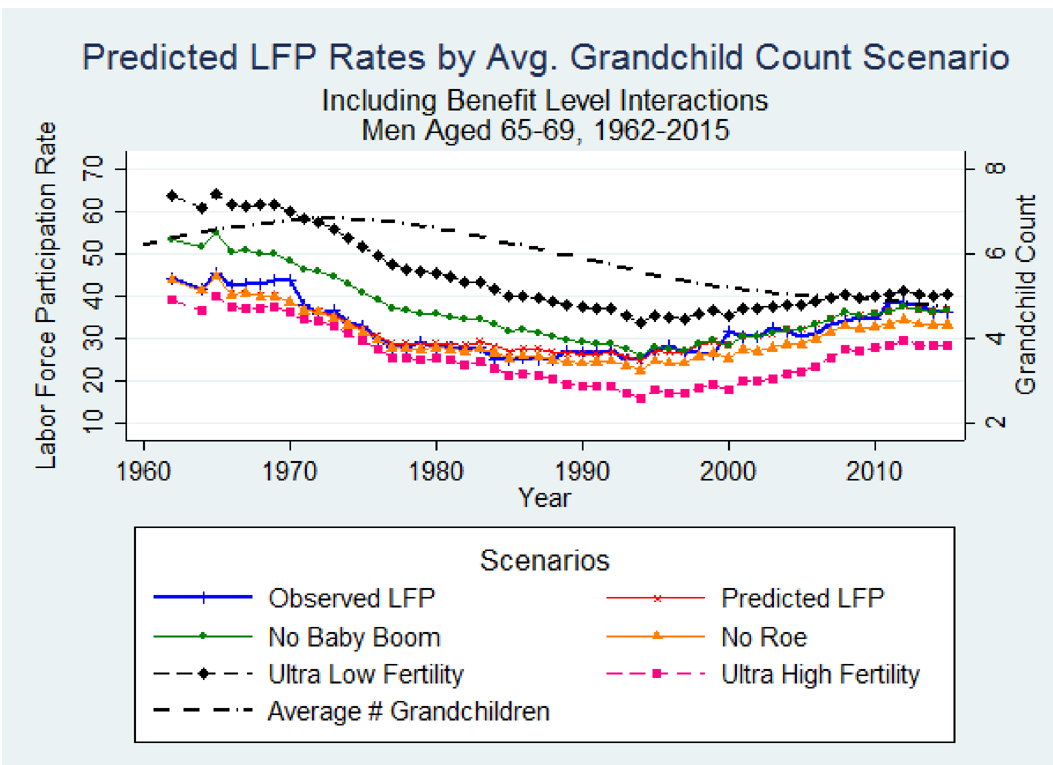

(b)

Bibliography

[1] Aaronson, D. Using Sibling Data to Estimate the Impact of Neighborhoods on Children's Educational Outcomes. *Journal of Human Resources*, 33:915–946, 1998.

[2] Amialchuk, A. The Effect off Husband's Job Displacement on the Timing and Spacing of Births in the United States. *Contemporary Economic Policy*, 31:73–93, 2011.

[3] Ananat, E. O., Gruber, J., and Levine, P. Abortion Legalization and Life-Cycle Fertility. *Journal of Human Resources*, 42:375–397, 2007.

[4] Aparicio-Fenoll, A. and Vidal-Fernandez, M. Working Women and Fertility: The Role of Grandmothers' Labor Force Participation. *CESifo Economic Studies*, 61:123–147, 2007.

[5] Arnott, R. Time for Revisionism on Rent Control? *Journal of Economic Perspectives*, 9:99–120, 1995.

[6] Arnott, R. and Igarashi, M. Rent Control, Mismatch Costs, and Search Efficiency. *Regional Science and Urban Economics*, 30:249–288, 2000.

[7] Arnott, R. and Shevyakhova, E. Tenancy Rent Control and Credible Commitment in Maintenance. *Regional Science and Urban Economics*, 47:72–85, 2014.

[8] Ault, R. W., Jackson, J. D., Saba, R. P. The Effect of Long-Term Rent Control on Tenant Mobility. *Regional Science and Urban Economics*, 47:42–55, 1994.

[9] Autor, D. H. The Unsustainable Rise of the Disability Rolls in the United States: Causes, Consequences, and Policy Options. In J. K. Scholz, H. Moon, and S.-H. Lee, editors, *Social Policies in an Age of Austerity: A Comparative Analysis of the US and Korea.* Edward Elgar Publishing, Northampton, MA, 2015.

[10] Autor, D. H. Women Working Longer: Facts and Some Explanations. In C. Goldin and L. F. Katz, editors, *Women Working Longer*. University of Chicago Press, Chicago, IL, 2017.

[11] Autor, D. H., Duggan, M. G. The Growth in the Social Security Disability Rolls: A Fiscal Crisis Unfolding. *Journal of Economic Perspectives*, 20:71–96, 2006.

[12] Bailey, M. J. More Power to the Pill: The Impact of Contraceptive Freedom on Women's Life Cycle Labor Supply. *The Quarterly Journal of Economics*, 121:289–320, 2006.

[13] Bailey, M. J. "Momma's Got the Pill": How Anthony Comstock and *Griswold v Connecticut* Shaped US Childbearing. *American Economic Review*, 100:98–129, 2010.

[14] Bailey, M. J., Guldi, M., Davido, A., and Buzuvis, E. Early Legal Access: Laws and Policies Governing Contraceptive Access, 1960-1980. *Mimeo*, 2011.

[15] Basu, K. Technological Stagnation, Tenurial Laws, and Adverse Selection. *American Economic Review*, 79:251–255, 1989.

[16] Basu, K. and Emerson, P. M. The Economics of Tenancy Rent Control. *Economic Journal*, 110:939–962, 2000.

[17] Basu, K. and Emerson, P. M. Efficiency Pricing, Tenancy Rent Control, and Monopolistic Landlords. *Economica*, 70:223–232, 2003.

[18] Behaghel, L., and Blau, D. M. Framing Social Security Reform: Behavioral Responses to Changes in the Full Retirement Age. *American Economic Journal: Economic Policy*, 4:41–67, 2012.

[19] Benabou, R. Workings of a City: Location, Education, and Production. *The Quarterly Journal of Economics*, 108:619–652, 1993.

[20] Bitler, M. P and Zavodny, M. The Effect of Abortion Restrictions on the Timing of Abortions. *Journal of Health Economics*, 20:1011–1032, 2001.

[21] Black, D. A., Kolesnikova, N., Sanders, S. G., and Taylor, L. J. Are Children "Normal"? *The Review of Economics and Statistics*, 95:21–33, 2013.

[22] Blau, D. M. and Goodstein, R. M. Can Social Security Explain Trends in Labor Force Participation of Older Men in the United States? *Journal of Human Resources*, 45: 328–363, 2010.

[23] Bloom, D. E., Canning, D., Fink, G., and Finlay, J. E. Fertility, Female Labor Force Participation, and the Demographic Dividend. *Journal of Economic Growth*, 14:79–101, 2009.

[24] Bloom, D. E., Canning, D., Fink, G., and Finlay, J. E. A New Look at Technical Progress and Early Retirement. *IZA Journal of Labor Policy*, 5:1–39, 2016.

[25] Bollinger, C. R. and Chandra, A. Iatrogenic Specification Error: A Cautionary Tale of Cleaning Data. *Journal of Labor Economics*, 23:235–257, 2005.

[26] Brueckner, J. K. Government land-use interventions: An economic analysis. In Lall S. V., Freire M., Yuen B., Rajack R., and Helluin J.-J., editors, *Urban Land Markets: Improving Land for Successful Urbanization.* Springer, 2009.

[27] Brueckner J. K. and Lai F. C. Urban Growth Controls with Resident Landowners. *Regional Science and Urban Economics*, 26:125–134, 1996.

[28] Bunten, D. Is the Rent Too High? Aggregate Implications of Local Land-Use Regulations. *Working Paper*, 2017.

[29] Burtless, G. Can Educational Attainment Explain the Rise in Labor Force Participation at Older Ages? Technical report, Center for Retirement Research, Chestnut Hill, MA, 2013.

[30] Calabrese, S., Epple, D., and Romano, R. On the Political Economy of Zoning. *Journal of Public Economics*, 91:25–49, 2007.

[31] Caldwell, J. C. Demographic Theory: A Long View. *Population and Development Review*, 30:297–316, 2004.

[32] Center for Budget and Policy Priorities. Where Do Our Federal Tax Dollars Go? Technical report, The Center on Budget and Policy Priorities, Washington, DC, 2016.

[33] Chetty, R. and Hendren, N. The Impacts of Neighborhoods on Intergenerational Mobility: Childhood Exposure Effects and County-Level Estimates. *Working Paper*, 2015.

[34] Cintina, I. The Effect of Minimum Drinking Age of Laws of Pregnancy, Fertility, and Alcohol Consumption. *Review of Economics of the Household*, 13:1003–1022, 2015.

[35] Colman, S., Dee, T., and Joyce, T. Do Parental Involvement Laws Deter Risky Sex? *Journal Of Health Economics*, 32:873–880, 2013.

[36] Compton, J. and Pollak, R. A. Family Proximity, Childcare, and Women's Labor Force Attachment. *Journal Of Health Economics*, 79:72–90, 2014.

[37] Dai, D. and Weinzimmer, D. Riding First Class: Impacts of Silicon Valley Shuttles on Commute & Residential Location Choice. *Working Paper, UCB-ITS-WP-2014-01*, 2014.

[38] Dee, T. S. The Effects of Minimum Legal Drinking Ages on Teen Childbearing. *Journal of Human Resources*, 36:823–838, 2001.

[39] Desmond, M. Eviction and the Reproduction of Poverty. *American Journal of Sociology*, 118:88–133, 2012.

[40] Dietz, R. The Estimation of Neighborhood Effects in the Social Sciences: An Interdisciplinary Approach. *Social Science Research*, 31:539–575, 2002.

[41] Downs, A. *Residential Rent Controls: An Evaluation.* Urban Land Institute, Washington, DC, 1988.

[42] Early, D. W. Rent Control, Rental Housing Supply, and the Distribution of Tenant Benefits. *Journal of Urban Economics*, 48:185–204, 2000.

[43] Eeckhout, J., Pinheiro, R., and Schmidheiny, K. Spatial Sorting. *Journal of Political Economy*, 122:554–620, 2014.

[44] Evans, W. N., Oates, W. E., and Schwab, R. M. Measuring Peer Group Effects: A Study of Teenage Behavior. *Journal of Political Economy*, 100:966–991, 1992.

[45] Fauth, R., Leventhal, T., and Brooks-Gunn, J. Early Impacts of Moving from Poor to Middle-Class Neighborhoods on Low-Income Youth. *Applied Developmental Psychology*, 26:415–439, 2005.

[46] Fernandez, R. and Rogerson, R. Keeping People Out: Income Distribution, Zoning, and the Quality of Public Education. *International Economic Review*, 38:23–42, 1997.

[47] Fishback, P. V., Haines, M. R., Kantor, S. Births, Deaths, and New Deal Relief during the Great Depression. *The Review of Economics and Statistics*, 89:1–14, 2007.

[48] National Center for Health Statistics. Health, United States, 2015: With Special Feature on Racial and Ethnic Health Disparities. Technical report, National Center for Health Statistics, Hyattsville, MD., 2016.

[49] Glaeser, E. L. Does Rent Control Reduce Segregation? *Swedish Economic Policy Review*, 10:179–202, 2003.

[50] Glaeser, E. L., Ward, B. A. The Causes and Consequences of Land-Use Regulation: Evidence from Greater Boston. *Journal of Urban Economics*, 65:265–278, 2009.

[51] Goldin, C. and Katz, L. F. The Power of the Pill: Oral Contraceptives and Women's Career and Marriage Decisions. *Journal of Political Economy*, 110:730–770, 2002.

[52] Goldin, C. and Katz, L. F. Long-Run Changes in the Wage Structure: Narrowing, Widening, Polarizing. *Brookings Papers on Economic Activity*, 2:135–167, 2007.

[53] Gruber, J., Levine, P., and Staiger, D. Fertility Effects of Abortion and Birth Control Pill Access for Minors. *Demography*, 45:817–827, 2008.

[54] Gustman, A. and Steinmeier, T. Retirement in Dual-Career Families: A Structural Model. *Journal of Labor Economics*, 18:503–545, 2000.

[55] Gustman, A. and Steinmeier, T. How Changes in Social Security Affect Recent Retirement Trends. *Research on Aging*, 31:261–290, 2008.

[56] Gyourko, J. and Linneman, P. Equity and Efficiency Aspects of Rent Control: An Empirical Study of New York City. *Journal of Urban Economics*, 26:54–74, 1989.

[57] Gyourko, J. and Linneman, P. Rent Controls and Rental Housing Quality: A Note on the Effects of New York City's Old Controls. *Journal of Urban Economics*, 27:389–409, 1990.

[58] Häckner, J. and Nyberg, S. Rent Control and Prices of Owner-Occupied Housing. *The Scandinavian Journal of Economics*, 102:311–324, 2000.

[59] Heiland, F. W. and Li, Z. Changes in Labor Force Participation of Older Americans and Their Pension Structures: A Policy Perspective. Technical report, Center for Retirement Research, Chestnut Hill, MA, 2012.

[60] Ho, C. Grandchild Care, Intergenerational Transfers, and Grandparents' Labor Supply. *Review of Economics of the Household*, 13:359–384, 2015.

[61] Hochman, O. and Lewin-Epstein, N. Determinants of Early Retirement Preferences in Europe: The Role of Grandparenthood. *International Journal of Comparative Sociology*, 0:1–19, 2013.

[62] Hsieh, C.-T. and Moretti, E. Why Do Cities Matter? Local Growth and Aggregate Growth. *NBER Working Paper 21154*, 2015.

[63] Hubert, F. Contracting with Costly Tenants. *Regional Science and Urban Economics*, 25:631–654, 1995.

[64] Hughes, W. and Turnbull, G. Uncertain Neighborhood Effects and Restrictive Covenants. *Journal of Urban Economics*, 39:160–172, 1996.

[65] Hurd, M. and Rohwedder, S. Trends in Labor Force Participation: How Much is Due to Changes in Pensions? *Journal of Population Ageing*, 4:81–96, 2011.

[66] Iwata, S. The Japanese Tenant Protection Law and Asymmetric Information on Tenure Length. *Journal of Housing Economics*, 11:125–151, 2002.

[67] Jacks, D. S., Krishna, P., and Shigeoka, H. Infant Mortality and the Repeal of Federal Prohibition. *NBER Working Paper 23372*, 2017.

[68] Johnsen, J. V. Retirement, Grandparental Childcare, and Maternal Employment. *Working Paper*, 2015.

[69] Joyce, T., Tan, R., and Zhang, Y. Abortion Before & After Roe. *Journal of Health Economics*, 32:804–815, 2013.

[70] Karney, C. F. F. Transverse Mercator with an Accuracy of a Few Nanometers. *Journal of Geodesy*, 85:475–485, 2011.

[71] Katz, L., Kling, J., Liebman, J. Moving To Opportunity in Boston: Early Results of a Randomized Mobility Experiment. *The Quarterly Journal of Economics*, 116:607–654, 2001.

[72] Kendall, T. D. and Tamura, R. Unmarried Fertility, Crime, and Social Stigma. *Journal of Law and Economics*, 53:185–221, 2010.

[73] Kirmeyer, S. E. and Hamilton, B. E. Transitions Between Childlessness and First Birth: Three Generations of U.S. Women. Technical report, National Center for Health Statistics., 2011.

[74] Kleibergen, F. Generalizing Weak Instrument Robust IV Statistics Towards Multiple Parameters, Unrestricted Covariance Matrices, and Identification Statistics. *Journal of Econometrics*, 139:181–216, 2007.

[75] Klerman, J. A. U.S. Abortion Policy and Fertility. *American Economic Review*, 89: 261–264, 1999.

[76] Kreuger, A. B. and Pischke, J.-S. The Effect of Social Security on Labor Supply: A Cohort Analysis of the Notch Generation. *Journal of Labor Economics*, 10:412–437, 1992.

[77] Krol, R. and Svorny, S. The Effect of Rent Control on Commute Times. *Journal of Urban Economics*, 58:421–436, 2005.

[78] Kutty, N. K. The Impact of Rent Control on Housing Maintenance: A Dynamic Analysis Incorporating European and North American Rent Regulations. *Housing Studies*, 11:69–89, 1996.

[79] Levine, P. Parental Involvement Laws and Fertility Behavior. *Journal of Health Economics*, 22:861–878, 2003.

[80] Levine, P., Staiger, D., Kane, T. J., and Zimmemrman, D. J. *Roe v Wade* and American Fertility. *American Journal of Public Health*, 89:199–203, 1999.

[81] Linneman, P. The Effect of Rent Control on the Distribution of Income Among New York City Renters. *Journal of Urban Economics*, 22:14–34, 1987.

[82] Maestas, N. and Zissimopoulos, J. How Longer Work Lives Ease the Crunch of Population Aging. *Journal of Economic Perspectives*, 24:139–160, 2010.

[83] McFarlane, A. Rent Stabilization and the Long-Run Supply of Housing. *Regional Science and Urban Economics*, 33:305–333, 2003.

[84] Miron, J. R. Rent Stabilization and the Long-Run Supply of Housing. *Urban Studies*, 27:167–184, 1990.

[85] Moon, C.-G. and Stotsky, J. G. The Effect of Rent Control on Housing Quality Change: A Longitudinal Analysis. *Journal of Political Economy*, 101:1114–1148, 1993.

[86] Morgan, S. P. Is Low Fertility a Twenty-First-Century Demographic Crisis? *Demography*, 40:589–603, 2003.

[87] Munch, J. R. and Svarer, M. Rent Control and Tenancy Duration. *Journal of Urban Economics*, 52:542–560, 2002.

[88] Myers, C. K. Young Women's Access to Contraception and Abortion, 1960-present. *Mimeo*, 2016.

[89] Myers, C. K. The Power of Abortion Policy: Re-Examining the Effects of Young Women's Access to Reproductive Control. *Journal of Political Economy*, Forthcoming, 2017.

[90] Nagy, J. Increased Duration and Sample Attrition in New York City's Rent Controlled Sector. *Journal of Urban Economics*, 38:127–137, 1995.

[91] Nagy, J. Do Vacancy Decontrol Provisions Undo Rent Control? *Journal of Urban Economics*, 42:64–78, 1997.

[92] Navarro, P. Rent Control in Cambridge, Massachusetts. *Public Interest*, 78:83–100, 1985.

[93] Neumark, D., Zhang, J., and Ciccarella, S. The Effects of Wal-Mart on Local Labor Markets. *Journal of Urban Economics*, 63:405–430, 2008.

[94] Nilsson, J. P. Alcohol Availability, Prenatal Conditions, and Long-Term Economic Outcomes. *Journal of Political Economy*, Forthcoming, 2017.

[95] Okrent, D. *Last Call: The Rise and Fall of Prohibition.* Scribner, New York, 2010.

[96] Oreopolous, P. The Long-Run Consequences of Living in a Poor Neighborhood. *The Quarterly Journal of Economics*, 118:1533–1575, 2003.

[97] Organisation for Economic Development and Cooperation. Live Longer, Work Longer: A Synthesis Report. Technical report, 2006.

[98] Posadas, J. and Vidal-Fernandez, M. Grandparents' Childcare and Female Labor Force Participation. *IZA Journal of Labor Policy*, 2:1–20, 2013.

[99] Raess, P. and von Ungern-Sternberg, T. A Model of Regulation in the Rental Housing Market. *Regional Science and Urban Economics*, 32:475–500, 2002.

[100] Rosenberg, J. Measuring Income for Distributional Analysis. Technical report, Urban-Brookings Tax Policy Center, Washington, DC, 2013.

[101] Rupert, P. and Zanella, G. Grandchildren and Their Grandparents' Labor Supply. *Working Paper*, 2016.

[102] Sabia, J. J. and Anderson, M. D. The Effect of Parental Involvement Laws on Teen Birth Control Use. *Journal of Health Economics*, 45:55–62, 2016.

[103] Sacerdote, B. Peer Effects with Random Assignment: Results for Dartmouth Roommates. *The Quarterly Journal of Economics*, 116:681–704, 2001.

[104] Schaller, J. Booms, Busts, and Fertility: Testing the Becker Model Using Gender-Specific Labor Demand. *Journal of Human Resources*, 51:1–29, 2016.

[105] Schirle, T. Why Have the Labor Force Participation Rates of Older Men Increased Since the Mid-1990s? *Journal of Labor Economics*, 26:549–594, 2008.

[106] Sims, D. P. Out of Control: What Can We Learn from the End of Massachusetts Rent Control? *Journal of Urban Economics*, 61:129–151, 2007.

[107] Sinai, T. Spatial Variation in the Risk of Home Owning. In E. Glaeser and J. Quigley, editors, *Housing Markets and the Economy, Risk Regulation, and Policy*, pages 83–112. Lincoln Institute of Land Policy, 2009.

[108] Skelley, C. Rent Control and Complete Contract Equilibria. *Regional Science and Urban Economics*, 28:711–743, 1998.

[109] Staiger, D. and Stock, J. H. Instrumental variables regression with weak instruments. *Econometrica*, 65:557–586, 1997.

[110] Stock, J. H. and Yogo, M. Testing for Weak Instruments in Linear IV Regression. In D. W. K. Andrews, editor, *Identification and Inference for Econometric Models.*, page 557–586. Cambridge University Press, New York, 2005.

[111] Svarer, M., Rosholm, M., and Munch, J. R. Rent Control and Unemployment Duration. *Journal of Public Economics*, 89:2165–2181, 2005.

[112] Taylor, P., Livingston, G., Cohn, D., Wang, W., and Dockterman, D. The New Demography of American Motherhood. Technical report, Pew Research Center., Washington, DC, 2010.

[113] Taylor, P., Morin, R., Parker, K., Cohn, D., and Wang, W. Growing Old in America: Expectations vs. Reality. Technical report, Pew Research Center., Washington, DC, 2009.

[114] Turner, B., and Malpezzi, S. A Review of Empirical Evidence on the Costs and Benefits of Rent Control. *Swedish Economic Policy*, 10:11–56, 2003.

[115] Wang, Y. and Marcotte, D. E. Golden Years? The Labor Market Effects of Caring for Grandchildren. *Journal of Marriage and Family*, 69:1283–1296, 2007.

[116] Whelan, S. Work or Care? The Labour Market Activity of Grandparents in Australia. *Working Paper*, 2013.

[117] Wickersham, G. W. *Enforcement of Prohibition Laws, Vol. 5.* U. S. Government Printing Office., Washington, DC, 1931.

Appendix A

Rent Control Data Appendix

A.1 Commuter Shuttle Stop Data Appendix

This Appendix has a list of citations for the sources used for shuttle stops and routes, and lists any assumptions made on routes. Information on shuttle routes comes chiefly from Google,[1] Stamen Design,[2] and the publicly available sources detailed below.

There are some gaps in the publicly-available information on where the shuttle stops are, so some assumptions were made about timing and placement. These are detailed in Section A.1.4.

A.1.1 Websites

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

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

[1]Interview with Brendon Harrington conducted on May 30, 2014

[2]`http://content.stamen.com/zero1`

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

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

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

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

Map of Google Shuttle Stops as of July 2013 Maintained by Anonymous Google Maps User, https://www.google.com/maps/d/viewer?mid=z6pdx1V0R3mU.kAHGFAmzMOVo&hl=en_US.

Map of Apple Shuttle Stops as of September 2013 Maintained by Google Maps User Gtok, https://maps.google.com/maps/ms?gl=US&ptab=1&ie=UTF8&oe=UTF8&msa=0&msid=2158120776704329 0004e5e26e576e99b6e0b&dg=feature.

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

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

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

A.1.2 News Stories and Publications

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

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

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

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

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

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

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

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

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

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

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

A.1.3 Miscellaneous Sources

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

A.1.4 Assumptions

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

A.2 Public Bus Stop Data Appendix

A panel of all public bus, train, and light rail stops was assembled to complement the dataset on the commuter shuttles. While this was a straightforward exercise for the Caltrain, the BART, and the MUNI light rail, SFMTA records on public bus stop information are only available in four waves of information. The agency directly provided a series of PDFs on all the SFMTA bus routes as of February 2008 and then a supplementary dataset of stop characteristics as of 2015. From the 2008 PDFs, a list of stops was assembled and some basic information such as stop length was collected. Data prior to 2008 was not forthcoming, but representatives of the agency claim that public bus stops were virtually unchanged between 2003 and 2008.

From `data.sfgov.org`, the main repository for any publicly available information published by the City and County of San Francisco, a March 2012 list of stops with their characteristics was made available. This was merged on to the 2008 routes data to fill in more information about each stop. While it necessarily misses any changes in stop characteristics between February 2008 and March 2012, it is likely that most stops changed very little over time.[3]

Starting from February 2009 onwards, the SFMTA makes publicly available a refreshed list of their public bus stops on a semi-regular basis (updated usually every 2-3 months).[4] In total, there were 43 waves of public bus stop names released by the SFMTA between February 2009 and December 2013. However, the stop naming system used between 2008 and 2009 changed, so that it was impossible to link the two datasets directly. Namely, the 2008 data uses special stop name abbreviations that are used also in the 2012 and 2015 stop characteristics spreadsheet, but the 2009 and onward spreadsheets report only the full address.[5] A merge of the 2008 names onto the 2009-2013 datasets on the basis via the March

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

[4]The data are available at a repository at this web address: `http://www.gtfs-data-exchange.com/agency/san-francisco-municipal-transportation-agency/`

[5]*e.g.*, "Powell St and Market St" has the special abbreviation of "POWLMRKT" or "MRKTPOWL", depending which street the actual stop is on. The full address and the abbreviation can be linked with some difficulty.

2012 spreadsheet, which has both fields, left about 14% of the dataset unmatched on average. A careful review of the names of the unmatched street addresses left only 31 stops (or about 0.7%) of 2009-2013 stops not assigned one of the SFMTA's abbreviations. When narrowed down only to stops that were eligible to host a commuter shuttle, only 7 out of 880 were dropped from an inability to consistently identify them.

The result is a panel of 880 eligible bus stops. The vast majority are concentrated around the central business district, but also throughout the city clustered around significant arteries. Figure 2.3 displays the locations of the eligible public bus stops, Caltrain, BART, and SFMTA light rail stops as of September 2004. The most significant change that occurs over the study's time period is the opening of the SFMTA's T Line on 3rd Street on the far east side of the city in January 2007.

Appendix B

Grandparents Appendix

B.1 Legal History and Policy Coding Detail

I am indebted to the detailed research conducted by Middlebury's Caitlin Knowles Myers, without which, this paper would not be possible. For the instrument, access is coded as the fraction of the year that a conception could be blocked or aborted. Abortion was legalized by *Roe* and most other state statutes through the first trimester, so eligibility is backdated 93 days (or the equivalent for other eligibility periods) prior to the legalization date. For the pill and Prohibition, the policy is coded as is.

In the PSID, the first wife's age, year of birth, and other characteristics were used to code access for adult sons. Access is coded based on the state where the daughter or daughter-in-law was living in 1968, to avoid introducing potential endogeneity from women moving to states where contraception or abortion was legalized. Unless directly observed, the daughter-in-law's 1968 state is assumed to be the same as her husband's. Following the exhaustive reviews of the state laws on abortion and contraception given in Myers [88], [89] and Bailey et al. [14] each daughter or daughter-in-law is coded as having access to abortion or access to contraception if there were no barriers to her access, such as spousal or parental consent requirements.[1]

The variation in state laws I use for the instrument is given in four tables. Tables B1-B4 shows the state-by-year policy variation. Table B1 shows the month and year of access to oral contraceptives for unmarried women between ages 18-20 and under 18. Table B2 shows the month and year of access to abortion on-demand for women 21 and over and between 18-20. Table B3 shows the month and year of access to abortion by minors, and shows that there is substantial variation between and within states when minors did and did not have free access to abortion. Where exemptions existed for married individuals or minors who graduated from high school, I used PSID data to code these minors as having access for the individual-level estimates. Table B4 shows the dates alcohol sales were banned, inclusive of the 18th and 21st Amendments to the Constitution where applicable.

Table B5 shows the frequency distribution of birth years for the daughters and daughters-in-law. For adult sons who never married, the birth year of the mother of their oldest child is used. For adult sons who never married and never had children, their birth year minus 2 is used, reflecting the fact that on average, men marry women 2 years younger than themselves. Including them this way reflects their potential to provide grandchildren, even if it is never realized.

[1]There are some cases where a fair reading of the law prohibited or potentially allowed access, but either the provision was not enforced or contemporary sources indicated that physicians did not perform reproductive services until the laws were clarified. I am indebted to Myers [88], [89]) for identifying which provisions were likely enforced, and which were not, beyond a fair reading of the plain text. In those cases where Myers [88], [89] indicated that a law was not enforced or not followed, I have deferred to her work. Additional ambiguities were resolved with supplemental information from Joyce, Tan, and Zhang (2013), Levine [79], Sabia and Anderson [102], Bailey [12], Levine et al. (1999), and Bitler and Zavodny [20].

B.1.1 Access to Oral Contraceptives

The legislative history of access to oral contraceptives begins in 1960 when Enovid was approved by the Food and Drug Administration for the prevention of pregnancy (Junod and Marks (2002)). At this time, legal minors were largely defined as being under 21 and could not freely obtain hormonal birth control.[2] In fact, many states had complete or partial bans on contraceptive sales through a series of laws known as "Comstock Laws".[3] In 1965, the Supreme Court ruled in *Griswold v Connecticut* that Connecticut's Comstock law banning the sale of contraceptives to married couples was unconstitutional, holding that the Constitution ensured a right to privacy.[4] In practice, in every state except for Massachusetts, this meant that all women of the age of the legal majority could freely buy oral contraceptives. The right to privacy for unmarried woman of legal majority was formally established by the 1972 Supreme Court ruling *Eisenstadt v Baird*.[5] Here, the Supreme Court struck down a Massachusetts law that prohibited the sale of contraceptives to unmarried individuals. In between these rulings, states and courts directly or indirectly reshaped the laws governing access to contraception. Many studies have exploited this variation, including Goldin and Katz [51], Bailey [12], Bailey et al. [14], Guldi [53], Myers [89], among others, to find that access to the pill allowed women to increase their labor supply, chiefly by delaying births.

For women between 18-20, states lowered the age of legal majority to 18 or 19 in waves, culminating in unimpeded access to oral contraceptives for all 18-20 year olds by 1976. Bailey [12] establishes that the laws that permitted young women to purchase oral contraceptives were passed for reasons mostly orthogonal to expanding access to reproductive technology. Commonly, states in this period lowered their age of legal majority from 21 to 18, which incidentally allowed women of those ages to buy birth control. These legislative actions were

[2] Only in Arkansas and Alaska was 18 the age of full legal majority for women in 1960.

[3] For a history and discussion of the influence of the state Comstock Laws, see Bailey [13]

[4] 405 U.S. 438 (1965)

[5] 405 U.S. 438 (1972)

taking place in the context of the debate over the draft, voting rights of soldiers, and the Vietnam War, and so had little connection to greater demands for reproductive freedom of choice. Other states had mature minor statutes, which hold that minors can consent to medical procedures and services if the minor clearly demonstrates they understand the implications. Often, these predated the introduction of the pill or court rulings that established the doctrine, usually for reasons have nothing to do with access to oral contraceptives. For example, the Ohio Supreme Court established a mature minor doctrine in 1956 (four years before the introduction of the pill) following *Lacey v Laird*, which was litigated over a nose surgery performed on a minor.[6] While the age of majority laws would open up access for women aged 18-20, the mature minor doctrines would often allow all minors to obtain contraceptives. The same rulings that would grant access to 18-20 year olds were often extended to all minors, creating further exploitable variation in access. There is thus substantial state-by-year variation in who could freely buy birth control pills that forms the basis for the empirical strategy used in this paper.

B.1.2 Access to Abortion

Prior to the January 1973 Supreme Court *Roe v Wade* decision legalizing abortion on-demand through the first trimester, 6 states had already done so: California in Sept 1969,[7] followed by Hawaii,[8] Alaska,[9] New York,[10] and Washington State[11] in 1970, and Washington D.C. in 1971, with *de facto* legalization occurring there in the wake of *United States v.*

[6]Lacey v. Laird 166 Ohio St. 12, 139 N.E. 2d 25 (1956).

[7]People v. Belous 71 Cal. 2d 954 (September 5, 1969)

[8]Haw. Rev. Stat. § 453-16 (2010)

[9]Alaska Stat. § 18.16.010 (2010)

[10]Klerman [75]

[11]Wash. Stat. § 9.02.100 et seq. Washington's statute permitted abortion through the first four months instead of just the first trimester

Vuitch.[12,13] Ananat, Gruber, and Levine [3], Levine et al. (1999), Gruber, Levine, and Staiger (1999), Joyce, Tan, and Zhang (2013), and others have shown that live births declined for women in their prime childbearing years in the early repeal states compared to the non-repeal states in a manner consistent with a response to the change in policy.

Within the early repeal states, abortion on-demand was legalized inconsistently by age. California initially required minors (20 and younger) to obtain parental consent for an abortion, whereas hospitals in New York City announced they would perform them on minors between 17 and 20 without it. Further, Joyce, Tan, and Zhang (2013) demonstrate that the residency requirements (or lack thereof) acted as an exogenous shock on neighboring states, inducing women to travel to have an abortion, and lowering the birth rate of neighboring states, a finding also corroborated in Klerman [75], Ananat, Gruber, and Levine [3], and Levine et al. (1999) among others.

After *Roe*, some states acted to impose restrictions on abortion access, mostly requiring minors to obtain parental consent or to notify their parents before an abortion. These laws have been shown to effectively reduce access to abortion, such that variation in access continues after 1973. They are included as a source of variation although their ultimate impact on pregnancy incidence is unclear.[14] This can also be exploited to identify changes in the likelihood to have a child exogenous to the labor force characteristics of either the potential grandparents or parents. Thus, a key innovation in this paper is to instrument for

[12] 402 U.S. 62 (April 1, 1971)

[13] Myers [88] and Klerman [75] point out that in addition to the full-repeal states, 11 states had adopted the American Law Institutes' Model Penal Code (MPC) statutes on abortion, which permitted it if the progression of the pregnancy would cause mental or physical harm to the mother. The convention I use in this paper is to code access as being only those states that granted abortion on-demand, which the MPC statutes did not. Myers [88], [89] duly shows that while abortion rates in the MPC states were somewhat higher than in the non-reform states, they were significantly lower than the full-repeal states.

[14] Bitler and Zavodny (2001) showed that requiring parental notification or consent did in fact lower the abortion rate among teens in the states that passed these laws. Levine [79] also found that parental involvement laws lower the abortion rate but did not find a statistically significant reduction in the overall birth rate. The mechanism is itself unclear: Sabia and Anderson [102] extend Levine's finding by testing specifically for the effect of the parental involvement laws on teen birth control use. Their findings suggest that parental involvement laws do increase the probability that sexually active minors use birth control, but Colman, Dee, and Joyce [35] examine the same question and do not.

timing and number of grandchildren by using state-by-year differences in access to abortion and oral contraceptives.

B.1.3 Prohibition

Although federal Prohibition began on January 16, 1920 after the 18th Amendment had been ratified the year prior, 33 states (including D.C.) had already enacted bans on the sale of alcohol. The earliest of these being Kansas in 1881. While temperance's champions were optimistic that America's "Noble Experiment" would improve social outcomes, it became clear over the course of the 1920's that Prohibition itself created several ill effects. Although alcohol consumption did steeply decline, these benefits were offset by criminal gangs formed to run bootlegging operations, undermined respect for the rule of law due to private consumption from the well-off's private stashes of alcohol, and a spike in alcohol poisonings from moonshine and other homemade brews (Okrent [95]).

The advent of the Great Depression heralded the end of Prohibition, because governments badly needed the revenues that alcohol provided. Franklin Roosevelt ran on a platform of repealing Prohibition, and throughout 1933, state conventions ratified the 21st until Utah became the 36th state to do so on December 5th, 1933, whereby Prohibition was immediately revoked (Okrent [95]). However, several states had statutes or constitutional provisions outlawing the sale or manufacture of alcohol. The majority of these lingering state Prohibitions were repealed shortly thereafter, but Oklahoma, Kansas, and Mississippi retained statewide Prohibition through World War II, and many Southern and Midwestern states have “dry" counties to this day. The last of the statewide Prohibitions was repealed by Mississippi in 1966.[15]

The net effect of Prohibition on births rates is unclear, since a positive effect could exist

[15]See Table B4 for more information.

due to less violence against women (a chief driver of the temperance movement), and improved maternal health and marital quality. However, inasmuch as alcohol is a facilitator of risky behavior, then the birth rate could decline due to fewer unintended pregnancies. Evidence shows that infant mortality increased after the ratification of the 21st Amendment, so that by 1939 an excess of 13,665 infant deaths could be attributable to the end of Prohibition (Jacks, Pendakur, and Shigeoka (2017)). Further, an examination of Figure 3.2's birthrates from 1920 to 2014 shows that the downward trend in the birth rate was arrested around 1932-1933. Fishback et al. [47] attributes this in part to New Deal relief programs, but the role of the repeal of Prohibition played remains largely unexplored.

Research on other changes to alcohol policy also shows convincing evidence that these laws can change natality outcomes. Shocks to alcohol consumption during pregnancy differentially effect male fetuses, so that the sex ratios at birth are closer to even, implying that mothers pregnant with boys suffered higher miscarriage rates (Nilsson [94]). Decreases in state minimum legal age drinking (MLDA) laws has been shown to decrease fertility among non-poor white teenagers (Cintina [34]), but increases in MLDA's have been shown to decrease fertility among black teenagers (Dee [38]). Taken as a whole, this evidence is strong enough to warrant including exposure to Prohibition as an instrument for grandparenthood.

Month and Year of Unhindered Access to Oral Contraception for Women Under 21

TABLE B1

State	18-20	Under 18
Alabama[a]	10/1971	10/1971 (14)
Arizona	5/1972	10/1977
Arkansas	7/1873	3/1973
California	3/1972	1/1976
Colorado	4/1971	4/1971
Connecticut	10/1971	
District of Columbia	8/1971	8/1971
Florida[a]	7/1973	
Georgia	4/1971	7/1972
Illinois	10/1961	
Indiana	9/1973	
Iowa[b]	7/1973	
Kansas	5/1970	5/1970
Kentucky	6/1968	7/1972
Louisiana	8/1972	7/1975
Maine[a,b]	6/1972	
Maryland	7/1971	7/1971
Massachusetts[c]	1/1974	1/1977
Michigan	1/1972	2/1980
Minnesota[c]	6/1973	1/1976
Mississippi	5/1966	5/1966
Missouri[d]	7/1977	
Nebraska[b]	7/1972 (19)	
New Jersey[d]	1/1973	
New York[b]	9/1973	7/1975
North Carolina	7/1971	7/1977
Ohio	6/1965	6/1965
Oregona	9/1971	9/1971
Pennsylvania	4/1970	9/1997
South Carolina	6/1972	6/1972 (16)
South Dakota	7/1972	
Tennessee	5/1971	7/1971
Texas[d]	8/1973	
Utah	7/1960	7/1975
Virginia	11/1971	11/1971
Washington	7/1968	7/1968
West Virginia	7/1972	7/1992
Wisconsin[c]	3/1972	7/1978

Table B1 shows the first year/month of legal, unhindered access to contraception for unmarried, childless women under 21.

[a] Access for minors under certain exemptions: being married, already being a parent, being a high school graduate, or the physician believes there is harm to the minor by not providing service.

[b] IA lowered its age of majority first to 19 in July 1972; ME lowered it to 20 first in 10/1969. NE lowered it to 20 first in 3/1969. NY first lowered age of access to 16 in 1973.

[c] Granted access to married minors before granting it to all: MA (1965), MN (1971), and WI (1960).

[d] Married minors can get access, year effective in parenthesis: NJ (1965), TX (1974).

Sources: Author's coding using the state statutes, Myers [88], [89], Bailey [12], Bailey et al. [14].

Month and Year of Unhindered Access to Abortion On-Demand for Women 18 and Over

TABLE B2

State	**21 and Over**	**18-20**
Alabama	1/1973	1/1973
Arizona	1/1973	1/1973
Arizona	1/1973	1/1973
California	9/1974	5/1971
Colorado	1/1973	7/1973
Connecticut	1/1973	1/1973
District of Columbia	4/1971	8/1974
Florida	1/1973	7/1973
Georgia	1/1973	1/1973
Illinois	1/1973	1/1973
Indiana	1/1973	1/1973
Iowa	1/1973	1/1973
Kansas	1/1973	1/1973
Kentucky	1/1973	1/1973
Louisiana	1/1973	1/1973
Maine	1/1973	1/1973
Maryland	1/1973	1/1973
Massachusetts	1/1973	1/1974
Michigan	1/1973	1/1973
Minnesota	1/1973	1/1973
Mississippi	1/1973	1/1973
Missouri[a]	7/1976	7/1976
Nebraska[b]	1/1973	1/1973
New Jersey[c]	1/1973	1/1973
New York	7/1970	7/1970
North Carolina	1/1973	1/1973
Ohio	1/1973	1/1973
Oregon	1/1973	1/1973
Pennsylvania	1/1973	1/1973
South Carolina	1/1973	1/1973
South Dakota	1/1973	1/1973
Tennessee	1/1973	1/1973
Texas	1/1973	1/1973
Utah	1/1973	1/1973
Virginia	1/1973	1/1973
Washington	12/1970	12/1970
West Virginia	1/1973	1/1973
Wisconsin[c]	1/1973	1/1973

Table B2 shows the first year/month of legal, unhindered access to abortion for unmarried, childless women 18 and over.

[a] Prior to the Supreme Court's *Danforth* decision, Missouri had a spousal consent requirement for married women seeking abortions.

[b] Minors in NE are 18 and under.

[c] New Jersey and Wisconsin had pending court cases challenging the validity of anti-abortion statutes and the legality of abortion on-demand prior to Roe is unclear. Most studies do not treat these as repeal states.

Sources: Author's coding using the state statutes, Myers [88], [89], Ananat, Gruber, and Levine [3], Levine et al. (1999), Joyce, Tam, and Zhang (2013).

Parental Involvement Laws for Legal Minors, Date Enjoined or Enforced, 1968-2013

TABLE B3

State	Enjoined or Explicit Access	Enforced
Alabama	1/1973-9/1987	9/1987
Arizona	1/1973-7/1982	7/1982-10/1985
	10/1985-5/1986	5/1986-8/1987
	8/1987-2/2003	3/2003-9/2009
	10/2009-8/2011	2011-present
Arkansas[a]		1/1973-2/1976
	2/1976-2/1989	3/1989-present
California		9/1969-5/1971
Colorado		1/1973-2/1975
	2/1975-6/2003	6/2003-present
Connecticut[a]		1/1973-11/1998
District of Columbia		4/1971-8/1974
Florida		1/1973-1/1978
	1/1978-6/2005	7/2005-present
Georgia[a]		9/1991-present
Illinois	1/1973-8/2013	8/2013-present
Indiana	1/1973-4/1973	4/1973-1/1975
	2/1975-8/1984	9/1984-present
Iowa	1/1973-12/1996	1/1997-present
Kansas	1/1973-6/1992	7/1992-present
Kentucky[a]		1/1973-11/1974
	11/1974-3/1989	3/1989-7/1991
	7/1991-7/1994	7/1994-present
Louisiana	1/1973-6/1973	6/1973-1/1976
	1/1976-9/1978	9/1978-3/1980

Parental Involvement Laws for Legal Minors, Date Enjoined or Enforced, 1968-2013

TABLE B3

State	Enjoined or Explicit Access	Enforced
	3/1980-7/1980	7/1980-present
Maryland	1/1973-5/1977	5/1977-12/1985
	1/1986-present	
Massachusetts		8/1974-6/1976
Michigan	1/1973-3/1991	3/1991-8/1992
	8/1992-3/1993	4/1993-present
Minnesota	1/1973-7/1981	8/1981-11/1986
	11/1986-8/1988	8/1988-present
Mississippi	1/1973-7/1993	7/1993-present
Missouri[b]	11/1973-6/1974	6/1974-2/1975
	2/1975-6/1983	6/1983-11/1983
	11/1983-8/1985	8/1985-present
Nebraska	1/1973-5/1973	5/1973-11/1975
	11/1975-6/1977	7/1977-12/1978
	1/1979-5/1981	5/1981-9/1983
	9/1983-9/1991	9/1991-present
North Carolina	5/1973-10/1995	10/1995-present
Ohio	1/1973-9/1974	3/1976-8/1976
	8/1976-10/1990	10/1990-present
Oklahoma	2/1973-5/1975	5/1975-6/1976
	7/1976-6/2001	7/2001-6/2002
	6/2002-11/2004	11/2004-present
Pennsylvanina	1/1973-3/1994	3/1994-present
South Carolina[c]		7/1973-11/1974
	11/1974-5/1990	5/1990-present

Parental Involvement Laws for Legal Minors, Date Enjoined or Enforced, 1968-2013

TABLE B3

State	**Enjoined or Explicit Access**	**Enforced**
South Dakota	1/1973-3/1973	3/1973-6/1976
	7/1976-6/1997	7/1997-present
Tennessee	1/1973-11/1992	11/1992-7/1996
	7/1996-1/2000	1/2000-present
Texas	1/1973-12/1999	1/2000-present
Utah	1/1973-3/1973	3/1973-9/1973
	9/1973-4/2006	5/2006-present
Virginia[a]		1/1973-6/1976
	7/1976-6/1997	7/1997-present
Washington		11/1970-1/1975
West Virginia	1/1973-5/1984	5/1984-present
Wisconsin	1/1973-6/1992	6/1992-present

Table B3 gives the effect dates of free or conditional access to women under the age of 18 for 1968 PSID states that had changes in the law. Statute dates should be read left to right on down.

[a] Preexisting parental consent or notification law or attorneys general ruling whose legality was left unclear after *Roe* and *Danforth*.

[b] Spousal consent law in effect between 6/1974 and 2/1975.

[b] Abortions without parental involvement permitted for women 16 and over between 7/1973-11/1974, and 17 and over currently.

Sources: the author's coding based on state statutes, Myers (2012, 2014), Sabia and Anderson (2016), Levine (2003), and Bitler and Zavodny (2001).

Date of Prohibition Enactment and Repeal by State

TABLE B4

State	**Effective Enactment Date**	**Effective Repeal Date**
Alabama[a]	07/01/1915	03/10/1937
Arizona	01/01/1915	12/05/1933
Arkansas	01/01/1916	12/05/1933
California	01/16/1920	12/05/1933
Colorado	01/01/1916	12/05/1933
Connecticut	01/16/1920	12/05/1933
Delaware	01/16/1920	12/05/1933
District of Columbia	03/03/1917	03/01/1934
Florida	01/01/1919	11/06/1934
Georgia[b]	01/01/1908	03/22/1935
Idaho	01/01/1916	11/06/1934
Illinois	01/16/1920	12/05/1933
Indiana	04/01/1918	12/05/1933
Iowa	01/01/1916	06/19/1934
Kansas	05/01/1881	07/08/1949
Kentucky	11/05/1919	11/10/1935
Louisiana	01/16/1920	12/05/1933
Maine	01/07/1885	10/1/1934
Maryland	01/16/1920	12/05/1933
Massachusetts	01/16/1920	12/05/1933
Michigan	04/30/1918	12/05/1933
Minnesota	01/16/1920	01/06/1934
Mississippi	12/31/1908	07/01/1966
Missouri	01/16/1920	01/1934
Montana	12/31/1918	12/05/1933
Nebraska	05/01/1917	11/06/1934
Nevada	12/18/1918	12/05/1933
New Hampshire	05/01/1918	1934
New Jersey	01/16/1920	12/05/1933
New Mexico	10/01/1918	12/05/1933
New York	01/16/1920	12/05/1933
North Dakota	11/02/1889	11/3/1936
North Carolina	01/01/1909	1937
Ohio	05/27/1919	12/05/1933
Oklahoma	09/17/1907	04/07/1959
Oregon	01/01/1916	12/05/1933
Pennsylvania	01/16/1920	12/05/1933
Rhode Island	01/16/1920	12/05/1933
South Carolina	12/31/1915	1934
South Dakota	07/01/1917	12/05/1933
Tennessee	07/01/1909	1937
Texas	05/1919	Early 1936
Utah	08/01/1917	Jan 1934
Vermont	01/16/1920	12/05/1933
Virginia	11/01/1916	03/07/1934
Washington	01/01/1916	12/05/1933
West Virginia	07/01/1914	11/06/1934
Wisconsin	01/16/1920	12/05/1933
Wyoming	07/01/1919	11/06/1934

Table B4 shows when states first enacted a Prohibition law at least outlawing the sale of hard liquors and when it was repealed.
Sources: Author's coding using the state statutes, Jacks et al. [67], National Association of Distillers and Wholesale Dealers (1918), and Patch (1933). See Section B.1.4 for others.

PSID In-Sample Daughter/Daughter-in-Law Year of Birth Distribution

TABLE B5

	Grandfather Sample			Grandmother Sample		
Year of Birth ⇓	**Frequency**	**Percent**	**Cumulative Percent**	**Frequency**	**Percent**	**Cumulative Percent**
Before 1940	21	0.385	0.385	64	0.806	0.806
1940-1944	60	1.1	1.485	140	1.763	2.569
1945-1949	359	6.582	8.067	716	9.015	11.584
1950-1954	950	17.418	25.486	1,578	19.869	31.453
1955-1959	1,200	22.002	47.488	1,917	24.137	55.591
1960-1964	1,257	23.047	70.535	1,808	22.765	78.356
1965-1969	862	15.805	86.34	978	12.314	90.67
1970-1974	419	7.682	94.023	441	5.553	96.223
After 1974	326	5.977	100	300	3.777	100

Table B5 shows the birth year distribution of four types of adult children: adult daughters, the first wife of an adult son, the mother of the adult son's oldest child if the adult son did not marry, or if the adult son never married and never had a child, his birth year minus 2. The first frequency table is adult children in the grandfather sample and the second table is for the grandmother sample.

B.1.4 Prohibition References

General Sources

Jacks, D. S., Pendakur, K., Shigeoka, H., 2017. Infant Mortality and the Repeal of Federal Prohibition. NBER Working Paper 23372.

National Association of Distillers and Wholesale Dealers. *The Anti-Prohibition Manual 1918: A Summary of Facts and Figures Dealing with Prohibition.* Cincinnati: The Publicity Department of the National Association of Distillers and Wholesale Dealers, 1918.

Patch, B. W. (1933). Liquor control after repeal. Editorial research reports 1933 (Vol. II). Washington, DC: CQ Press. Retrieved from `http://library.cqpress.com/cqresearcher/cqresrre1933100400`

Wickersham, G. W., 1931. Enforcement of Prohibition Laws, vol. 5. Washington: U. S. Government Printing Office.

Alabama

White, R., 2016. Death and Re-Birth of Alabama Beer. *Business History* 58, 785-795.

District of Columbia

The District of Columbia Prohibition Act of 1917, Pub. L. No 64-383, 39 Stat. 1123 (1917). Retrieved from `http://legisworks.org/congress/64/publaw-383.pdf`

Jones, M. "Happy Repeal Day, Maryland and Virginia! (Sorry, D.C.)" *Boundary Stones, WETA's Local History Blog.* WETA. 5 Dec, 2012. Last Accessed 23 May, 2017. `http://blogs.weta.org/boundarystones/2012/12/05/happy-repeal-day-maryland-and-virginia-sorry-dc`

Florida

Guthrie, J. J. *Keepers of the Spirits: The Judicial Response to Prohibition Enforcement in Florida, 1885-1935 (Contributions in American History).* Greenwood Publishing Group, 1998. Google Books. Web. Last Accessed 23 May, 2017. https://books.google.com/books?id=FmGyiyiyVGwC&pg=PA126&lpg=PA126&dq=florida+statewide+prohibition+repealed+1934&source=bl&ots=LdECXOVLpw&sig=vhBtMYk-hO1SHeaKjfGJWjk9pVY&hl=en&sa=X&ved=0ahUKEwjfl-vC5YbUAhXl1IMKHX5MA-IQ6AEILjAB#v=onepage&q=florida%20statewide%20prohibition%20repealed%201934&f=false

"This Day in Georgia History, March 22, 1935, Prohibition Ended." *GeorgiaInfo: An Online Georgia Almanac*, 2017. Digital Library of Georgia, An Initiative of the University System of Georgia. 2017. Web. Last Accessed 23 May, 2017. Retrieved from http://georgiainfo.galileo.usg.edu/thisday/gahistory/03/22/prohibition-ended

Idaho

"Idaho Constitutional Amendment History." *Idaho Secretary of State Election Division*. State of Idaho. 2017. Web. Last Accessed 23 May, 2017. Retrieved from http://www.sos.idaho.gov/elect/inits/hst20_30.htm

Iowa

Alcohol Beverages Division, State of Iowa. "Iowa's Alcoholic beverages Laws and You". A Licensee's Guide to Compliance with Iowa's Alcoholic Beverages Laws." 2016 Edition. Web. Retrieved from https://abd.iowa.gov/sites/default/files/documents/2016/04/redbook_4.21.16.pdf

Kansas

Cassity, P. D., 1986. The Waning of Kansas Prohibition: 1933-1948. Masters of Arts Thesis, Emporia State University. Retrieved from https://esirc.emporia.edu/bitstream/handle/

123456789/1972/Cassity%201986.pdf?sequence=1

Kentucky

Brunn, S. D., Appleton Jr., T. H., 1999. Wet-Dry Referenda in Kentucky and the Persistence of Prohibition Forces 39, 172-189.

Maine

"Constitutional Amendments Enacted, 1834-" *Maine State Legislature.*, 27 Feb, 2017. Government of the State of Maine. Web. Last Accessed 23 May, 2017. https://legislature.maine.gov/9203/

Minnesota

"Repeal and Decline" *History on the Web.* Web. Last Accessed 23 May, 2017. Retrieved from http://www.historyontheweb.org/minnbrew/repeal.html

Mississippi

Holder, H. D., Cherpitel, C. J., 1996. The End of U.S. Prohibition: A Case Study of Mississippi. *Contemporary Drug Problems* 23, 301-330.

Missouri

"Prohibition in Missouri." *Missouri Life*, Missouri Life Inc. Web. Last Accessed 23 May, 2017. Retrieved from http://www.missourilife.com/life/prohibition-in-missouri/.

Nebraska

"Statistics on Constitutional Amendments, Initiated and Referred Measures." Nebraska State Government, 2012. Web. Last Accessed 23 May, 2017. Retrieved from `https://web.archive.org/web/20150316183737/http://nebraskalegislature.gov/pdf/bluebook/245-274.pdf`.

Nevada

"In Las Vegas, Prohibition Was Sporadically Enforced", *Prohibition: An Interactive History.* The Mob Museum of Las Vegas. Web. Last Access 23 May, 2017. `http://prohibition.themobmuseum.org/the-history/prohibition-in-las-vegas/las-vegas-and-prohibition/`

North Carolina

Crowell, M., 1979. A History of Liquor-by-the-Drink Legislation in North Carolina. *Campbell Law Review* 1, 61-110.

North Dakota

North Dakota Liquor Control Initiative, State Law 1937, ch. 259 (1936). Retrieved from `http://www.legis.nd.gov/files/resource/library/measuresbeforethevoters.pdf`

Ohio

Ohio Memory. "The Prohibition Era Begins...." *Ohio Memory*, State Library of Ohio, 1 June, 2012. Last Accessed 23 May, 2017. Retrieved from `http://www.ohiohistoryhost.org/ohiomemory/archives/641`.

Oklahoma

Franklin, J. L., "Prohibition," *The Encyclopedia of Oklahoma History and Culture*, Oklahoma Historical Society, 2009. Last accessed 23 May, 2017. Retrieved from `http://www.okhistory.org/publications/enc/entry.php?entry=PR018`.

Texas

"Prohibition Elections in Texas", *Texas Almanac*, Texas State Historical Association. Last Accessed 23 May, 2017. Retrieved from `http://texasalmanac.com/topics/elections/prohibition-elections-te`

Utah

Means, S. P. "Prohibition: 75 Years Ago, Utah's Vote Repealed 'The Nobel Experiment'", *The Salt Lake Tribune*, The Salt Lake Tribune, 1 Dec, 2008. Last Accessed 23 May, 2017. Retrieved from `http://archive.sltrib.com/story.php?ref=/ci_11115941`.

Virginia

Benbow, M. E. "Prohibition in Virginia," *Rusty Cans*, rustycans.com, 2017. Last Accessed 23 May, 2017. Retrieved from `http://www.rustycans.com/HISTORY/virginia.html`.

West Virginia

West Virginia Const. Art. VI. Sec. 46 et seq.

Wyoming

Roberts, P. Regulating Liquor: Prohibition Enforcement, Official Corruption, And State Efforts to Control Alcohol After Prohibition Repeal. *Wyoming Law Review* 12, 389-451.

B.2 Estimating Grandparenthood Measures

As far as I am aware, no one data source tracks longitudinally how many grandchildren respondents have. Thus, the average number of grandchildren and the fraction of each birth cohort that are grandparents has to be estimated from extant sources. Unfortunately, the PSID is not a broad enough sample to credibly estimate this figure at the state level or education group level by birth cohort. I thus used the Health and Retirement Study (HRS) and Retirement History Longitudinal Survey (RHLS) data,[16] which cover information from 1973-1979 (biennially) and from 1992-2014 (biennially).[17] Specifically, I used the RAND HRS files which compress the survey responses into a "wide" dataset of each's respondent's longitudinal responses.[18] These retirement surveys have large samples of older individuals and provide the necessary sampling breadth to credibly calculate grandparent statistics.

I combined the HRS and RHLS responses into a synthetic panel covering bienially 1973-1979 and 1992-2014 that estimated by age, birth cohort, and education group the fraction who are grandfathers and their total number of grandchildren. However, this left many cells with missing information. Thus, the second step fits a simple model of either the fraction grandfather or number of grandchildren by age by birth year to extrapolate these results to missing years, ages, and birth cohorts:

$$\begin{aligned} GP\_Measure_{etab} &= \beta_0 + \beta_1 CumulativeBirthrate_{tab} + \beta_2 \mathbb{1}\{Age_{etab} \geq 33\} \\ &+ \beta_3 BirthYear_{db} + \beta_4 BirthYear^2_{db} + \delta EducationGroup_e + \epsilon_{etab}, \end{aligned} \tag{B.2.1}$$

where $GP\_Measure_{etab}$ is either the fraction who are grandfathers in birth cohort b at age a in year t, and are in education group e or the number of grandchildren each grandfather has.

[16] The HRS (Health and Retirement Study) is sponsored by the National Institute on Aging (grant number NIA U01AG009740) and is conducted by the University of Michigan.

[17] While the RHLS in fact covers 1969-1979 for the 1906-1911 birth cohort, the first three survey years (1969, 1971, and 1973) did not ask about grandchildren.

[18] The RAND HRS Data file is an easy to use longitudinal data set based on the HRS data. It was developed at RAND with funding from the National Institute on Aging and the Social Security Administration.

$EducationGroup_e$ is a vector of education group dummies for the four categories: less than high school, high school, some college, and college degree. $CumulativeBirthrate_{tab}$ is the cumulative sum of the national crude birthrate that starts at age 33 for grandfathers and then monotonically increases until age 84 for each intervening age a in year t.

This model was calibrated only for the very earliest grandparent years. For men under 33, the fraction grandparent and the number of grandchildren was set to zero. For men 33 to 34, the fraction grandparent is set to 0.5% and the grandchild count is set to 0.005. For men aged 35, the fraction grandparent is set to 1% and the grandchild count is set to 0.01. Remaining values for other ages are set by the model.

The reasoning here is that for each birth cohort, higher birthrates in year t represent a higher likelihood of becoming grandparents and a higher likelihood of welcoming a new grandchild, so birthrates are an important control. Yet, simply lagging the birthrates would not work well here, because a 20 year lag on the birth rate for an individual at age 40 is meaningless while being meaningful for a man at age 60. Thus, the running sum of the birthrate for the individual captures both that higher birthrates mean higher chances of being a grandparent and more grandchildren while also accounting for the fact that sustained high birth rates over time should increase these measures monotonically with age.

In order to completely fill the synthetic panel by means of the above model, it was necessary to find information on birthrates then going back to at least 1925, when the members of the oldest in-sample birth cohort (1892) turned 33. For the 1925-1930 period, I used the "Vital Statistical Rates in the United States, 1900-1940" published by the National Office of Vital Statistics, which reports in Table 44 on p. 666-667 the crude birth rate for the birth registration states from 1915-1940. It is important to note that the national crude birth rate is computed just from participating states, which by 1933, included all states.[19]

[19] By 1925, all states were participating except for 13 holdouts: Alabama, Arkansas, Colorado, Georgia, Louisiana, Missouri, Nevada, New Mexico, Oklahoma, South Carolina, South Dakota, Tennessee, and Texas. Alabama, Arkansas, Louisiana, Missouri, and Tennessee joined in 1926. Colorado, Georgia, Oklahoma, and South Carolina joined in 1927. Nevada and New Mexico joined in 1928. South Dakota did not join until 1932 and Texas was the last continental state to join the registry in 1933. Alaska and Hawaii joined upon statehood, with statistics being reported for Alaska in 1959 and Hawaii in 1960.

For the 1941-1967 period, I used the annual National Vital Statistics of the United States reports, which listed the counts of births for each state. To generate the crude birthrates spanning 1941-1967, I then used for the population denominators the Census Bureau's "Annual Estimates of the Population for the U.S. and States, and for Puerto Rico".[20]

For the 1968-2004 period, I used the National Centers for Health Statistics publicly-available natality microdata.[21] These contain either a full or partial sample of all of the birth records down to the county level for all registry states from 1968-2004. For population denominators spanning 1969-2004, I then used the Surveillance, Epidemiology, and End Results (SEER) Program's county by sex by single-age yearly population estimates that I could then aggregate up to the national level. For the 1968 population denominator, I again used the Census Bureau's "Annual Estimates of the Population for the U.S. and States, and for Puerto Rico".

From 2005 to 2015, the publicly available natality microdata suppresses all geographic identifiers, so I used the published birthrates and counts made available in the National Vital Statistics Systems' annual publication on births. These publications include the state and national birth counts, and for population denominators, I again used the SEER population estimates.

Repeating this for the state-level analysis required significantly more steps. First, I recreated the synthetic panel on grandparenthood measures by year/age, birth year, and Census Division. Unfortunately, both the HRS and the RHLS only publicly report respondents location at the Census Division level, so I could not further disaggregate it to the state level. The grandparenthood model then becomes:

$$\begin{aligned} GP\_Measure_{dtab} &= \beta_0 + \beta_1 CumulativeBirthrate_{stab} + \beta_2 \mathbb{1}\{Age_{stab} \geq 33\} \\ &+ \beta_3 BirthYear_{sb} + \beta_4 BirthYear_{sb}^2 + \delta State_s + \epsilon_{stab}, \end{aligned} \tag{B.2.2}$$

[20]Their total population estimate includes the Armed Forces serving overseas, so I instead aggregated their state population estimates to generate the national estimate of people resident in the United States in a given year.

[21]Most readily available courtesy of the National Bureau of Economic Research at `http://www.nber.org/data/vital-statistics-natality-data.html`

The d index in Equation (B.2.2) in $GP\_Measure$ and the error is the state's Census division, and s represents the state.

State-by-year birth rates were computed first by compiling the birth counts (numerators) for each state and then the population counts (denominators) and the calculating accordingly. As described above, for 1925-1940, crude birth rates were drawn from the Vital Statistics of the United States reports. For 1941 to 2015, birth counts are drawn from the Vital Statistics of the United States reports (1941-1967), the National Center for Health Statistics (NCHS) Natality microdata (1968-2004), or the births reports published by the National Center for Health Statistics (2005-2015). For the 1968-2004 period, individual-level birth records were aggregated into state by year cells, but in all other cases, births are reported at the state level. Population denominators are the same as those described above.

B.3 Replicating Blau and Goodstein Social Security Estimates

To calculate average mean earnings and Social Security and disability benefits, I followed the procedures outlined by Blau and Goodstein (2010) in their Appendix 1 (p. 356-361) to extend their estimates to the 2006-2015 period for the education group, national-level analysis. My only modification was to use the earnings levels instead of log earnings, for the estimation that predicts earnings at ages 27, 32, 37, 42, 47, 52, and 57 as reported on p. 357.

For the state-level analysis, I modified their approach by replacing the education group cells with state cells. Doing so presented two challenges. The first is that some state's birth cohort cells were consistently too small even over multiple years to generate plausible results. The second is that the CPS itself does not report individual states in some cases between 1968-1976. This presented particular problems because in the Blau and Goodstein model, generating the education group-to-population ratios is done as a function of birth years. The decade-long stretch with no earnings history for certain states rendered the out-of-sample predictions for the model to be several orders of magnitude off.

To circumvent both problems, I first grouped together the earnings histories of some states that both followed the patterns already established by the CPS and generated more plausible results for state-to-population earnings ratios. After running the Social Security Administration's ANYPIA benefits calculator for the old age and disability insurance amounts at various entitlement ages, I then reestablished the synthetic panel for all 50 states plus D.C., where certain states essentially share the same benefit amounts.