UCLA
UCLA Electronic Theses and Dissertations

Title
Essays in Urban and Labor Economics

Permalink
https://escholarship.org/uc/item/31h3b0mp

Author
Miller, Matthew Michael

Publication Date
2017

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA

Los Angeles

Essays in Urban and Labor Economics

A dissertation submitted in partial satisfaction of the

requirements for the degree of Doctor of Philosophy

in Economics

by

Matthew Michael Miller

2017

ABSTRACT OF THE DISSERTATION

Essays in Urban and Labor Economics by

Matthew Michael Miller

Doctor of Philosophy in Economics

University of California, Los Angeles, 2017

Professor Leah Michelle Boustan, Co-Chair

Professor Adriana Lleras-Muney, Co-Chair

American cities, after decades of decline, are regaining affluent and highly educated residents. This dissertation examines these trends at the neighborhood level and documents that resurgence, in the form of gentrification, is prevalent in cities, specifically in downtown neighborhoods near employment centers. My results indicate a fundamental shift from city center decline to growth around 1990, which motivates my focus on exploring two novel causes of gentrification.

In Chapter 1, I demonstrate that declining gender wage gaps since 1980, and the associated influence on female labor force participation and marriage decisions, are one cause of gentrification. As women increasingly invest in human capital and delay marriage they are more likely to move to, and remain in, urban neighborhoods close to employments areas, allowing these neighborhoods to develop high-quality amenities which facilitate further redevelopment. I document that as the gender wage gap declines so too does the probability of marriage and that, in turn, marital status factors heavily into family residential location decisions, with singles systematically opting to live closer to employment centers. Overall, I find that falling gender wage gaps had a significant but heterogeneous effect on neighborhood prosperity that benefited those neighborhoods nearest the

city center. Specifically, I find that the drop in the gender wage gap from 1970 to 2010 can explain 40% of the average national change in city center income over the same period. One potential factor that influenced the decline in gender wage gaps was a shift in the labor market returns to social skills, a shift that disproportionately aligned with female skill endowments relative to men. I return to this topic in Chapter 3.

In Chapter 2, I document the role that condominium development played in gentrification. The advent of condominiums offered high income individuals a legal form to own, rather than rent, high-density real estate close to employment centers. I use condominium conversion ordinances at the municipal level as a source of exogenous variation in condominium development. Using a differences-in-differences set-up, I find that the passage of ordinances limited the development of condominiums in cities subject to regulations that made it more costly to convert housing stock to condominiums. With this approach, I am able to establish a causal effect of condominium development on certain gentrification outcomes, including income and education.

In Chapter 3, I introduce a framework for estimating the labor market returns to social capital and explore related mechanisms. I find that the wage return to increasing one's high school network by approximately five friends is equivalent to the return to one additional year of schooling. To better understand the mechanisms that underlie this return, I introduce a game theory model wherein students optimize their time allocation between studying and socializing. Empirical results are consistent with model predictions, specifically in that students make social investments in activities, such as drinking alcohol, that generate friendships at the expense of academic achievement. As an application, I demonstrate baseline estimates that suggest there are positive returns to attending a so-called "party" school for college.

The dissertation of Matthew Michael Miller is approved.

Edward Chih Yuen Kung

Michael C. Lens

Adriana Lleras-Muney, Committee Co-Chair

Leah Michelle Boustan, Committee Co-Chair

University of California, Los Angeles

2017

Contents

List of Figures

List of Tables

Acknowledgments

I would like to thank Leah Boustan and Adriana Lleras-Muney for many instances of guidance and support, both for direction on our joint work and advice on my solo projects. Many other faculty have also offered productive commentary, including Edward Kung, Walker Hanlon, Till von Wachter, and Michael Lens, as well as participants in the Labor and Applied Microeconomics seminars series at UCLA.

I owe thanks to a number of peers who participated in research groups including Chad Stecher, Owen Hearey, Richard Domurat, Ben Smith, Carolina Arteaga, and Lucia Yanguas, as well as others who participated in our Applied Microeconomics summer reading group. The resources, and financial support in the form of an NICHD traineeship, provided by the California Center for Population Research (CCPR) have been critical to the completion of this research, as has financial support in the form of a Dissertation Year Fellowship from the UCLA Graduate Division.

Some of the work included is co-authored, and I would like to thank Leah Boustan, Robert Margo, and Justin Steil, as well as Adriana Lleras-Muney and Shuyang Sheng, for allowing me to reproduce drafts of joint work as Chapters 2 and 3, respectively.

Finally, I am incredibly thankful for the support of my family: Jen, Mom, Dad, and Brianna.

MATTHEW MICHAEL MILLER

Education

University of California, Los Angeles   Ph.D., June 2017 (expected)

Los Angeles, CA   M.A., June 2014

Economics

Dartmouth College   A.B., June 2009

Hanover, NH   Mathematics, Economics

The Hun School of Princeton   June 2005

Princeton, NJ

Employment

Oct. 2009 – Aug. 2012   Edgeworth Economics   Principal Economist

Washington, DC

June 2009 – Oct. 2009   NERA Economic Consulting   Research Associate

Boston, MA

In Preparation for Publication

"Equal Pay Saves the City: How Women Initiated an Urban Renaissance."

"Urban Gentrification: The Role of Condominiums" with Leah Boustan, Robert Margo, and Justin Steil.

"Party On: The Labor Market Returns to Socializing" with Adriana Lleras-Muney and Shuyang Sheng.

Fellowships and Awards

UCLA, Dissertation Year Fellowship (2016-17)

UCLA California Center for Population Research, NICHD Traineeship (2015-16)

UCLA Economics Department, Departmental Teaching Award (Fall 2014)

Dartmouth College, James O. Freedman Presidential Scholar (Spring 2009)

Dartmouth College, High Honors in Mathematics (Spring 2009)

Dartmouth College, Honors in Economics (Spring 2009)

Research Assistance

2014   Adriana Lleras-Muney   UCLA and NBER

2013   Leonardo Bursztyn   UCLA Anderson

Teaching Experience

Fall 2015	Introduction to Econometrics	Instructor
Spring 2015, Spring 2014, Fall 2013	Microeconomic Theory I	Teaching Assistant
Winter 2015	Economics of EntrepreneurshipI	Teaching Assistant
Fall 2014	Public Economics	Teaching Assistant
Summer 2014	Introduction to Econometrics	Teaching Assistant
Winter 2014	Microeconomic Theory II	Teaching Assistant

Professional Activities

All-California Labor Economics Conference	Poster Session (2016)
UCLA Graduate Economics Association	Representative (2014-16)
California Center for Population Research	Affiliate (2013-17)

Chapter 1

Equal Pay Saves the City: How Women Initiated an Urban Renaissance

1.1 Introduction

At the geographical center of Washington, DC is a neighborhood adjacent to the White House named Penn Quarter. This neighborhood lies within the economic center of the city, as measured by employment density.[1] In 1990, it had an average household income of just $25,800 (in 2010 dollars), a value in the first percentile of the broader metropolitan area's income distribution. By 2010 the same neighborhood had transformed, with its constituent households earning incomes of $143,900 on average, placing the neighborhood in the 82nd percentile of the metropolitan area income distribution. This transition from poverty to affluence has initiated in neighborhoods within virtually every major city in the United States. At the neighborhood level, this upward movement in the regional income distribution is what I define as *gentrification*. Economists have recently begun to document the trends of urban resurgence that arise when numerous neighborhoods in a city gentrify and have started to explore the causes and consequences. In this paper, I propose a new explanation for why urban neighborhoods have started to gentrify: reduced gender wage inequality. As labor markets have grown increasingly favorable for high-skilled women they have chosen to delay marriage and childbearing until later in life. The postponement of family formation has resulted in a growing population of young, educated, single women *and* men who prefer living in the city over living in suburban areas.

The late twentieth century, stretching into the current millennium, was a transformative era for both urban areas and the labor market. Over this period, labor market prospects for women changed dramatically, with the most remarkable changes occurring among highly educated women. The proportion of women in the first-year cohorts of professional schools climbed steadily from

---

[1]The city centers used in the paper are those identified in Miller (2015), are based on employment density, and are frequently referred to as the "economic center" of their respective cities. They correspond to the notion of a downtown or central business district and I use the terms interchangeably throughout.

approximately 10% in 1970 to 50% by 2005.[2] A well-documented impact of this increase in female labor force participation, and the shift from female "jobs" to "careers", is its effect on the traditional life-cycle of family structure. The "quiet" revolution in female labor force participation manifest importantly in the narrowing of the wage gap starting in 1980, related to a collective anticipation for more promising careers for women that would follow the pursuit of formal education (Goldin (2006)). The time series evidence suggests that the mean age of marriage increased steadily from 1970 to 2010, while the gender wage gap has undergone a remarkable decline in tandem (Goldin and Katz (2002), Goldin (2006), Polachek (2004), Weichselbaumer and Winter-Ebmer (2005)).[3]

Over this same period, urban neighborhoods began to move up their regional income distributions, mimicking the pattern of the Penn Quarter neighborhood in Washington, DC at a much broader scale. In 1970, the average city center neighborhood in the US ranked in the 30th percentile within its metropolitan area income distribution, increasing to the 38th percentile by 2010. In New York, the average city center neighborhood ranked in the 47th percentile in 1970, which improved by 2010 to the 80th percentile. These changes occurring in and around city center neighborhoods represent an extraordinary turn of fortune from the decline that plagued cities during the era of suburbanization in the twentieth century. These gentrification patterns are most pronounced in downtown neighborhoods, and manifest in increases in income, rent, and the share of educated individuals (Miller (2015), Edlund, Machado, and Sviatschi (2015), Couture and Handbury (2015)).

I introduce a simple model that highlights the mechanisms that tie these two trends together. Men and women meet and decide each period whether or not to marry. Once the marriage decision is made they decide where to live relative to an employment center. Married households differ from single households in two important ways: they make joint decisions based on an aggregation of

---

[2]These numbers reflect medicine. Law school and MBA programs were approximately 5% female in 1970, rising to 50% and 40%, respectively, by the aughts.

[3]Some of the referenced work does not extend as recent as 2010. I provide my own estimates of the evolution of these measures below to supplement existing work.

their individual, unequal earnings, and they have a non-zero probability of having children.[4] The probability of having children inflates the household preference for housing relative to the cost of commuting, inducing them to move further from where they work where space for housing is cheaper. This has parallels to Fujita (1989), Beckmann (1973) and other urban models with family formation. Using individual data, I demonstrate in support of the model that from 1970-2010 married couples are two to three percentage points less likely than singles to live in the central city and that couples who married within the last year are six percentage points more likely to move from the central city to the suburbs than singles in that same year. The percentage of households in a neighborhood that are single is highest at neighborhoods located exactly at the city center (approximately 62% in 1980, 68% in 2010), and decreases monotonically with distance to the city center (five miles from the city center this drops to approximately 30% in 1980 and 38% in 2010, reflecting the overall higher number of singles in the later period).

The gender wage gap has broad influence in residential location through its influence on the marriage decisions of single men and women. This relationship was originally discussed in Becker (1973), stating that "an increase in the wage rate of women relative to men would decrease the incentive to marry." The basis for the argument follows from the reallocation in marriage of the wife's production from labor market activities to home production that is rational when men earn more than women. Shenhav (2015) explores the empirical validity of this theory, and finds that marriage, as well as other family formation decisions, depends causally on the gender wage gap. A narrowing gender wage gap represents a reduction in the "monetary incentive" of marriage, increased financial independence for women, and existing estimates demonstrate the between 1980 and 2010 20% of the marriage decline can be attributed to changes in the gender wage gap. Similarly, I find

---

[4]For the remainder of this paper there is no childbearing decision separate from marriage. The marriage decision accounts for a shift from no potential of having children to a non-zero probability of having children, and therefore condenses childbearing and marriage essentially into one decision.

that a 10% change in the gender wage gap among singles is associated with approximately four percentage point decline in the probability of marriage.[5]

As the gender wage gap declines individuals are less likely to marry and, based on how residential decisions depend on marriage, have a higher propensity to live in the city. Both cross-sectional and time series evidence indicate that across US cities gender wage gaps are positively associated with higher average income in city neighborhoods relative to suburban areas. By delaying marriage individuals spend a larger portion of their working lives living in the city rather than in suburban areas. As wages grow with experience and age, this delay in city exit is a boon to centrally located neighborhoods. Moreover, college educated individuals are more likely to *ever* live in the city center when marriage is delayed. For example, in 1980 the average age of marriage among the college educated is 22 years old, the modal age for receiving a bachelor's degree. In 2010, the average age of marriage increased to 27 years old. In the earlier period, a larger fraction of college educated individuals are married during or immediately after their college graduation, and therefore more likely to move to suburban areas at the exact moment they enter the labor market. In recent years, with women more active in the high-skilled labor market and postponing marriage, there has been an influx of single, college-educated individuals to city neighborhoods.

To explore if gender wage gap changes have impacted the gentrification of urban neighborhoods I estimate a set of region and year-specific gender wage gaps. For clarity in the analysis, I focus on estimates of the gender wage gap *while single*, which I denote as γ.[6] With this wage structure

---

[5]This is in line with Shenhav (2015), who finds that a 10% change in the gender wage gap "leads to a 5.3 percentage point decline in the probability that a woman in married."

[6]Decomposing the the overall gender wage gap into components while single and while married reflects an empirical finding common to a growing literature, including Kleven, Landais, and Sogaard (2015), Bertrand, Goldin, and Katz (2010), Goldin and Katz (2008), and Angelov, Johansson, and Lindahl (2016), that suggests that while gender wage gaps have fallen dramatically, the remaining gender gap in earnings is attributable to changes in female earnings upon marriage and childbirth. Changes to a "marriage penalty" in wages for women would have a different impact on marriage decisions than the gap while single. However, the stylized fact in the literature, which I find consistent with my own estimates, is that the gap while single has been falling dramatically while the marriage gap has remained stable; hence my focus on the gender wage gap while single. While I leave a complete assessment of this topic for future work, sensitivities of the analysis that follows which include an estimate of a marriage penalty generate nearly

in mind I estimate both gender wage gap while single for each metropolitan area and year in the data. To accomplish this efficiently without imposing identical wage specifications in each region I employ the double selection LASSO procedure outlined in Belloni, Chernozhukov, and Hansen (2014). I find that these region specific gender wage gaps, in addition to a broad set of similar metrics, significantly and robustly influence marriage decisions by decreasing the probability that individuals marry.

Finally, I test whether or not the local gender wage gaps can help to explain the gentrification patterns observed in American cities. The results indicate that an erasure of the average 1970 single gender wage gap of 20% between male and female earnings is associated with a city center neighborhood improving by three percentiles in the income distribution, on average. This accounts for 40% of the average national change in neighborhood income over this period. I control for a broad set of competing theories that are believed to underlie urban migration patterns, including transportation (LeRoy and Sonstelie (1983), Baum-Snow (2007)), housing characteristics (Brueckner and Rosenthal (2009)), and racial composition (Boustan (2010), Card, Mas, and Rothstein (2008)). Additional theories which are more challenging to control for, based on data availability, include declining crime rates and improving urban amenities. I approach this by extending the "endogenous gentrification" work in Guerrieri, Hartley, and Hurst (2013), where neighborhood income is used as a proxy for neighborhood amenities, positive or negative, that spillover from adjacent areas. I estimate the model in a spatial econometric framework to account for simultaneity bias in these neighborhood level externalities, including negative amenities such as crime and positive amenities like restaurants and museums. To put the magnitude of these effects in perspective, the influence of a standard deviation change in the gender wage gap on gentrification in city center neighborhoods is comparable to adjacent neighborhoods improving by ten percentiles in the metropolitan area income

---

identical results with respect to the single gender wage gap and gentrification.

distribution. To address endogeneity concerns I adopt an approach similar to Shenhav (2015) and Beaudry and Lewis (2012) who also estimate region and marriage-market specific gender wage gaps. This approach is inspired by a growing literature pertaining to skill endowment differences between men and women, wherein men are better endowed with "brawn", or "hard" skill, and women with "brains", or "soft" skill (Welch (2000)). Technological innovations have increasingly substituted for tasks facilitated by "brawn", including physically demanding manual labor tasks. Those same innovations have complemented the "brains" skill set, including more cognitive and social tasks. I use a measure of the soft skill to high skill mix in a region, adjusted to reflect national changes in the relative returns to these skills, as an instrument for a region's gender wage gap. Major technological shocks, such as the adoption of the PC in the 1980's and 1990's, were most easily adopted in regions where the soft to hard skill mix was higher, and this adoption disproportionately benefited the "soft" skill set of women. Similar in spirit, I also construct a variant of the "wage proxy" from Shenhav (2015), which relies on occupational task-demand differences in that some occupations, especially those better suited for women, grow faster due to skill-biased technological change. The wage proxy and skill mix measures are sufficiently correlated with the gender wage gap predictions to satisfy typical relevance assumptions, and the results above hold using the skill mix measure as an instrument for the gender wage gap while single.

In addition to the gentrification and gender wage gap literature cited above, this paper contributes to work that explores the migration patterns of educated men and women. Costa and Kahn (2000) document the decisions of power couples, wherein both spouses are at least college educated, to migrate to larger metropolitan areas. In more recent work, Diamond (2016) describes how high skilled, college-educated individuals have sorted increasingly into high amenity metropolitan areas. The present work uses this premise as a point of departure by asking how, once a given population of educated, employed individuals has sorted into certain metropolitan areas, do they

choose to locate relative to their jobs? I assume those jobs are concentrated in a single employment center. While in many cities employment has partially suburbanized, and polycentricity may be more appropriate at a general level, the high skilled jobs held by the college educated population of focus are much more centralized (Glaeser and Kahn (2001)). Other work has considered how delayed marriage and childbearing has impacted urban areas. Boustan and Shertzer (2013), for example, document that central city decline in the twentieth century would have been even more pronounced absent population trends related to women and delayed childbearing. Literature that explores the location decisions of the gay population, such as Black et al. (2002), parallels this work, as gay couples are more constrained in having children, thus lowering their housing demand relative to heterosexual couples. Additionally, by definition gay couples face no purely gender-based pay gaps among partners, making this margin of comparison an interesting example of the mechanism of this paper. Madden (1980) found that in the 1970s singles chose to live closer to their jobs and that households with children, typically married, chose larger housing units located further from employment centers, and predicted that "if the increased labor force participation of women leads to a higher first age of marriage... and lower fertility, then it will decrease household size and increase demand for smaller, more centrally located housing." To my knowledge, this work is the first to explore how increased female labor market participation and delays in marriage, which have continued in earnest since the 1980's, help to explain the gentrification patterns presently occurring throughout the United States.[7]

The remainder of this paper is organized as follows. Section 1.2 presents motivating evidence of the links between gender wage gaps, delayed marriage, and gentrification, as well as a description of data sources; Section 1.3 introduces a model of that delivers comparative statics and a framework for stepping into the empirical analysis. Section 1.4 introduces the empirical set up, including estimates

---

[7]Some work, including Boustan and Shertzer (2013), looks at the impact of these changes at the broader city level.

of regional gender wage gap components and presents estimates in a econometric framework for how gender wage gaps influence gentrification. Section 1.5 concludes.

1.2 Motivation and Stylized Facts

It is only recently that a set of empirical facts about gentrification have been established in the literature. In this section, I present my contribution in documenting this type of urban redevelopment with a focus on a measure I refer to as the *income percentile*. For each year and metropolitan area I assign each neighborhood a percentile within the metropolitan area income distribution. This serves as a ranking of neighborhoods each year within their region. Positive changes in the income percentile over time indicate a neighborhood is improving relative to other areas in the vicinity, or equivalently, is *gentrifying*.

To explore how these measures vary in urban space, see Figure 1.2, which plots for 1980 and 2010 the income percentile of neighborhoods as a function of distance from the city center, estimated using standard non-parametric, local polynomial methods. The most noticeable changes occur at locations near the city center. At distance zero there is an increase from the 20th to approximately the 30th income percentile. The pattern is most pronounced at the center but occurs for other centrally located areas. The curve pivots around a point approximately six miles from the city center, with neighborhoods close to the city center improving and outlying neighborhoods losing ground in the income distribution.[8]

Figure 1.3 plots an alternative gentrification measure, the percentage of the population that is college educated at each distance, standardized in each region and year by subtracting the mean

---

[8]It is clear also from this figure that the curves are increasing in distance, meaning that on average suburban neighborhoods are higher income than urban neighborhoods. This is consistent with the age-extension version of the model below in that while young, single, educated individuals are increasingly likely to now live in city neighborhoods, once they marry they are still likely to move to the suburbs. Almost mechanically, the wages of the older cohort that move to the suburbs are higher based on wage increases with age and experience.

college educated share and dividing by the standard deviation.[9] City neighborhoods very close

to the center are increasing in their share of college-educated residents, and the percentages are

comparable to the levels in the outlying, higher income areas. On average, regions are increasingly

developing what is colloquially referred to as a "donut" shape, with a well-educated center, encircled

by a region that is relatively less educated, which itself is surrounded by another highly educated

ring.

1.2.1   Observations About Gender Wage Gaps and Gentrification

I estimate the gender wage gap at the metropolitan area level in each year. I describe in detail how I

estimate the gender wage gaps in the econometric section below, however, for this section it suffices

to clarify that is it a measure of male earnings relative to females and it, therefore, is nearly always

positive; the larger the gap, the less equitable the local labor market across gender. See Figures 1.4

for a time series comparison of the gender wage gap and gentrification measures. Plotted in this

figure is the average income percentile of city center neighborhoods and the gender wage gap for

the metropolitan area that contains that city center. In Figure 1.4 the city center income percentile

is increasing over time, with the income percentile in city centers rising from the 30th percentile in

1970 to the 38th percentile in 2010. In contrast, there is a downward trend in the gender wage gap.

In 1970 men on average earned nearly 65% more than women. By 2010 the gap had shrunk to less

than 20%. [10]

    There is also cross-sectional correlation across cities in these measures. In Figure 1.6 I plot the

---

[9]This is to better emphasize spatial trends; the share of college educated increases at all distances over time and makes the intra-year differences by distance difficult to see.

[10]In Figure 1.5 I plot the same time series trends with a focus on New York, NY and Chicago, IL. These two cities are the prototypical examples of "monocentric" cities, the urban structure that underlies the theoretical portions of this work. Moreover, New York and Chicago are two of the first cities to show signs of gentrification, with their city centers increasing from the 37th and 46th income percentiles to the 73rd and 74th, respectively, by 2010. Over the same period, both metropolitan areas experienced dramatic declines in the gender wage gap, from 60% to 14% for Chicago and from 52% to 12% for New York.

city center income percentiles against the gender wage gaps for both 1980 and 2010, in addition to year-specific, population weighted linear fits. In Figure 1.6a I include all of the top 100 cities. Figure 1.6b highlights a selection of large cities. In 1980 there is a negative relationship between income percentile and the gender wage gap, with a large range of gender wage gaps across different cities. Notice that by 2010 the points have shifted left and up, indicating that cities tend to be more equitable by male-to-female pay, and higher in the metropolitan area income distribution, in 2010.

1.3   Conceptual Framework: Why Do Gender Wage Gaps Influence Residential Decisions?

The intention of this section is to clarify the channels through which the gender wage gap influences residential decisions. There are three key theoretical components of this mechanism. For each, I present brief empirical evidence to support the claim:

1. The *first* component is an assumption that marriage decisions depend on the size of the relative wage differences between potential spouses. This insight was first proposed in Becker (1973), with the prediction that the incentive to marry declines as the hourly wages of women improved relative to men. I elaborate on this prediction based on recent literature and show evidence consistent with recent empirical work by Shenhav (2015) that demonstrates individual marriage probabilities decrease significantly as region-specific estimates of the gender wage gap falls.

2. The *second* component is a prediction that family structure influences residential location. Specifically, married households prefer outlying areas where housing space is more affordable but commutes are longer. This follows from the assumption that married households face the prospect of having children, while singles do not, which inflates their preference for housing

relative to commuting costs, and is an adaptation of monocentric city models with family structure as in Beckmann (1973) and Fujita (1989) using a probability of children (following Low (2014)).

3. The *third* component is a prediction that among singles, higher income individuals will choose to live closer to the city center. This is due to commuting costs that are a function of wages foregone while commuting, a standard component of urban residential location models.

1.3.1   Model Set-up For Marriage

The setting for the model is a single, monocentric region, with all locations identified by a distance d from the city center, wherein all employment is located. Within the region is a population of women, indexed by i, equal in size to the population of men, indexed by j. Each individual has an education level $s \in \{L, H\}$, a low or high skill-type, where wages are increasing in education $w(H) > w(L)$.[11] Couples (i, j) form within the region's dating market so that there is positive assortative matching, meaning couples form among individuals with common s. Men earn a premium relative to women in the form of the gender wage gap γ, such that $w_j(e) = (1 + \gamma)w_i(e)$. Each couple (i, j) decides whether or not to marry, and given a marriage decision chooses a location d to live.

1.3.2   Theoretical Assumption and Empirical Support: Probability of Marriage Falls with the Gender Wage Gap

A smaller gender wage gap γ lowers the pecuniary gains from marriage; this is a prominent result from Becker (1973), and follows from the optimal gains that may accrue to both parties from marriage. If γ is large, the household upon marriage can optimally reallocate lower-earning hours worked by women to home production and higher-earning hours worked by men to the labor market.

---

[11]For example, L could denote high school or lower and H could college or higher.

If the wage difference is zero, there is no possible reallocation of home versus labor production that can improve household utility beyond what is achievable while single, and hence marriage incentives are lower. If I let M indicate marriage, I assume the probability of marriage is increasing with the gender wage gap, so that $\partial P(M)/\partial \gamma > 0$.

To test this hypothesis, I use individual-level data from IPUMS for each decade from 1970-2010 to regress marital status on the a measure of the gender wage gap while single.[12] I estimate this relationship in a logistic regression and find that the probability of marriage decreases with the gender wage gap, with a 10% decline in the gender wage gap associated with a 3.8 percentage point drop in the probability of marriage. This result is comparable to both OLS and IV estimates by Shenhav (2015), who finds a 10% drop in the gender wage gap leads to a 5.3 percentage point drop in the probability of marriage. This evidence supports the prediction that γ is a key determinant of the share of the population that is married.

1.3.3 Theoretical Prediction about Single versus Married Residential Decision

Consider a single region that has an equal number of women, indexed by i, and men, indexed by j. Each receives, via an exogenous, fixed matching function, a match (i,j).[13] As outlined above, couples decides whether or not to marry based on the gender wage gap between them while single (γ), and conditional on a marriage decision, where to live relative to their places of work. The resulting commuting distance is denoted by d, relative to a common, unique center of employment. Each household has a structure (n_i, n_j, π), where n_i is the number of female wage earners, n_j is the number of male wage earners, and π is the probability of having children, with $n = n_i + n_j$. The household decision problem is:

---

[12]As with the motivational section, a full data description follows in Section 1.4 below.

[13]This is only for simplicity, as it is certainly of interest for their to be same-sex couples in the sense of (j, j') or (i, i'). This type of case is excluded due to the complication of keeping track of varying cohort sizes. There are interesting patterns observed with respect to gentrification and the gay population, which are discussed later on.

$$\max_{d,c,Q} \alpha c + (n + \pi)\beta \ln Q \quad s.t. \quad c + R(d)Q = (n + n_j\gamma)w(s)(T - b_t(d)) \tag{1.1}$$

Where c is a numeraire consumption good that is split evenly among household members, and Q is a quantity of housing that is public. d is the distance each wage-earner must commute for work, $R(d)$ is the price per square foot of housing at that distance, $b_t(d)$ is the time cost of commuting incurred by all wage earners as a function of distance, and T is the number of hours, less fixed leisure time, during which wage earners can work.[14] All male wage earners earn $(1+\gamma)w$, but female wage earners earn w, the male wage less the gender wage gap.

This set up differs from other urban models with family structure in two ways. First, instead of making decisions based on actual children, households make decisions here based on the prospect of having children; a child is a random variable and never actually realized. Second, as is outlined with some emphasis in this paper, wages for spouses are historically and notoriously *not* equal, with men very often earning more than women. Therefore, I incorporate wage heterogeneity between potential spouses in this model. Following the standard urban literature, I replace c in the utility function with the budget constraint and fix d to solve for the indirect utility, solve for the bid-rent function by invoking the usual spatial equilibrium condition. For simplicity, consider the case where $\gamma \geq 0$ and households are either single and female ($n_i = 1$, $n_j = 0$), single and male ($n_i = 0$, $n_j = 1$), or married ($n_i = n_j = 1$). The following proposition explains how these households will choose in equilibrium to locate in the urban space:

Proposition 1.3.1. *Let all households be characterized as either female single, male single, or married, and let $\pi^s = 0 < \pi^m$ denote the probability of children for singles and married couples, respectively. Then for $\gamma > 0$, equilibrium single men will locate closest to the city center, followed by*

---

[14]This is an assumption of monocentricity that maintains tractability for the model, fitting into the standard Alonso-Muth-Mills framework. I emphasize the t to highlight the distinction between pecuniary and time costs of commuting.

single women, and married couples will locate furthest. Moreover, if $\gamma = 0$ single male and female households choose locations identically, still locating closer to the city center than married couples.

Proof. Follows from standard argument in Fujita (1989). □

The intuition for this result is as follows. There are two channels that induce households to prefer suburban living. First is the probability of having children, which increases the desire for larger housing while keeping the commuting cost fixed, motivating married households to move further from employment centers than those who face no prospect of children. This result establishes the importance of marriage decisions in affecting residential location. Second, there is a time cost to commuting that depends on wage. Therefore when $\gamma > 0$ and women earn less than men they face a lower foregone cost of commuting and in equilibrium will locate further from the city center. When there is gender wage parity, male and female singles have identical time cost of commuting and hence behave identically and live closer to the city center than married couples. This is the first channel through which declining gender wage gaps theoretically influence urban living: when gender wage gaps while single are small, the city population will have a higher proportion of women. Moreover, these women are higher income relative to when $\gamma > 0$.

1.3.4   Where do Singles Choose to Live?

These family structure predictions regarding residential decisions are largely supported by the data. Family structure more generally refers to the composition of a household and the relations between the household members. The focus here is on two types: single versus married households. By single, I refer to "non-family" households, as defined by the Census, which include people who live alone or with others who are unrelated. Any household that includes at least two individuals who are related by birth, marriage, or adoption is called a family.[15] As described in Section 1.3.3, all

---

[15]https://www.census.gov/population/methods/p251129.txt

members of a household derive utility from housing, but not all members need to commute. Children especially do not bear any commuting costs but do enjoy the quality and size of the family's housing. Therefore any disparity between the number of family members and the number of wage earners, who must bear the cost of commuting to their employer, results in a higher willingness to pay for housing space relative to single households who by definition have no dependents.

In Figure 1.8 I plot curves fitted by local polynomial regression of the percentage of households who are single (solid line) and married (dashed line) for both 1980 (red) and 2010 (blue) by distance to the city center. For both years singles are more likely to live close to the city center, with the percentage single dropping precipitously from 60% to 30% within five miles of the city center. Similarly, married households are minimally represented in neighborhoods near the city center, increasing in their share for areas further and further away. Married couples make up nearly 70% of the households in the average neighborhood 18 miles from the city center. The curves for 2010 are similar, and additionally, reflect the overall increased proportions of single versus married households. Singles make up a higher proportion of households at every distance from the city center in 2010 relative to 1980.

In Tables 1.1 and 1.2 I extend this analysis by breaking down family types and running estimates at the individual level. In Table 1.1, using IPUMS samples for 1980, 1990, 2000 and 2010, I estimate the difference in the probability of living in the city between three types of family structures: single versus married (no children), lone earner versus dual-earner, and no children versus children. For both general and high prestige employment married couples without children are less likely to live in the central city relative to singles. However, neither dual-earner versus lone-earner, nor children versus no children, show any significant difference. This is one reason, in part, the focus of this paper is on the margin of single versus married, with the assumption that married people account for the expectation of children in their residential decisions. In Table 1.2, I look at whether couples who

16

are married within the last year are more likely to leave the central city for the suburbs than other central city households. Married couples, both those with and without children, are six percentage points more likely to leave the city upon marriage than other types of households.

1.3.5 How Does a Change in the Gender Wage Gap Influence Gentrification?

In the static residential decision, the ordering of households when $\gamma > 0$ was, from nearest to farthest from the city center: single men, single women, and married couples. The only type of wage heterogeneity introduced was between men and women. It is true that as $\gamma \to 0$ women are both increasing in income and centralizing; this is indeed one source of gentrification. But now, to pull together this stylized model, consider two skill levels, H (perhaps college or higher) and L (high school or lower), for which $w(H) > w(L)$. Again, I assume matches only occur within education type. This institutes another level of preference for distance: all else equal, those with higher wage will prefer to live closer to the city center due to having a higher time cost of commuting. Ignoring gender differences momentarily, this implies that the spatial equilibrium relative to the city center regarding income will be, from nearest to farthest from downtown: single H, single L, followed by married (both H and L).[16]

Now, consider what happens when the gender wage gap declines (see Figure 1.1). Let $d^s(\gamma)$ denote the distance that marks the boundary between single and married households. Falling γ increases the share of the population that is single, and $d^s(0)$ denotes this when there is no gender wage gap, at a value larger than when $\gamma > 0$. Alternatively, this can be thought of as a portion of the married population under $\gamma > 0$ – the red area in the figure – shifting to single (the two blue areas), who locate so that the "new" high income single households are closer to the center than the low income single households. This discussion is informal and stylized, but the dynamics illustrate

---

[16]This basic argument for how this equlibrium fills out follows from a discussion of bid-rent curve orderings and fixed land-use and density, explained in more detail in Fujita (1989).
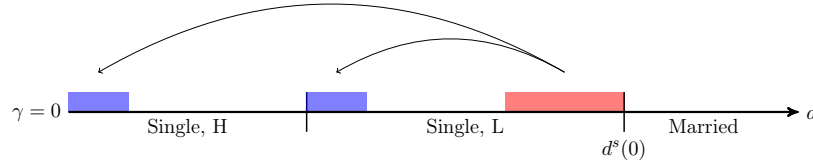
Figure 1.1: Influx of High Income Singles to City Center with Fall in Gender Wage Gap

a key component of the gentrification story.

Do Women Make Ip a Larger Proportion of the Urban Population?

Recall that the static residential decision not only implied that singles would live closer to the city center than married couples, it also suggested that among singles female households would locate further from the center relative to male households if the gender wage gap γ was large. As the gender wage gap narrows over time, the theory predicts that under large γ areas directly adjacent to the city center should be disproportionately male, while further out the areas should be disproportionately female, eventually at some distance giving way to areas where married households locate and hence male to female ratios should approach one. In later periods, where γ is small, this ratio should be more consistently flat, potentially even reversing in areas where $\gamma > 0$. In Figure 1.9 I plot a local polynomial fit of the male to female ratio among 16-34-year-olds across the entire sample of cities. In the NCDB I cannot jointly observe marital status and gender, so I focus on areas within five miles from the city center, which Figure 1.8 indicates is the area where singles are concentrated. In 1980, when the average gender wage gap while single is 19%, the areas directly at the city center are predominantly male and decrease steadily, with an increasingly higher proportion of women, until the five miles distance. However, in 2010 when the average gender wage gap while single is nearly zero, the line is much flatter, with more women in areas close to the city center and more men in areas further out. The national figure can be difficult to interpret given that distances are not uniformly relevant across cities. Areas five miles from the city center may be urban in a large

city, such as New York, and suburban in a smaller city, such as Washington, DC. Figures 1.10 and 1.11 include the same analysis but for New York, NY and Washington, DC, respectively.[17]

1.4 Empirical Evidence of the Gender Wage Gap Influence on Gentrification

In this section, I estimate the influence of gender wage gaps on gentrification. The outcome of interest is y_{mkt}, which is the income percentile of neighborhood m, in metropolitan area k, year t. Below is the main specification:

$$y_{mkt} = \alpha_k + \mu_t + \beta_1 g_{kt} + \beta_2 d_m + \beta_3 g_{kt} d_m + X_{mkt}\theta + \zeta_{mkt} \qquad (1.2)$$

From the prediction of the model, the main effect to focus on is the interaction between the estimates of the gender wage gap, g and d, the distance to the city center, which is predicted to be positive. As the gender wage gap decreases neighborhoods closer to downtown, with smaller d, should be improving as individuals, on average, delay marriage longer. From the econometric model above, the effect of interest is (assuming momentarily that $\mathbb{E}(\zeta|g, d, X) = 0$ holds):

$$\frac{\partial \mathbb{E}(y|g, d, X)}{\partial g} = \underbrace{\beta_1}_{<0} + \underbrace{\beta_3}_{>0} d$$

Before presenting the results for the estimation of this effect, I explain the data sources used, the estimation of the gender wage gaps, the various controls embedded in X, and endogeneity concerns about g.

---

[17]City center neighborhoods in both cities are actually majority female by 2010, which is consistent with the model. While in 1980 both cities have gender wage gaps among singles equal to approximately 6%, by 2010 they are -3% and -4% for New York and Washington, respectively, indicating that single women are making more than men on average, and hence have a higher commuting time cost and live closer to the center.

1.4.1 Data

At the individual level, I use Integrated Public Use Microdata (IPUMS) USA data, available each decade from 1970 to 2010.[18] This data includes detailed labor market variables, of which I use income and wages, usual hours worked, occupation and industry categories, native-born status, and full-time and part-time employment status, along with information on the highest level of education obtained. Additionally, demographic information is available that allows for the tracking of marital status, as well as the number of children. I limit the samples to individuals who live in identifiable metropolitan areas, and whose residential status can be categorized, within the metropolitan area, as either within or not within the central city. The prestige of an occupation is also tracked in these data sets, and in some robustness exercises, I focus on high-prestige careers, which I identify using the Siegel occupational prestige score.

At the neighborhood level, I use data from the Neighborhood Change Database (NCDB) which is curated by Geolytics, Inc.[19] I use 2010 Census tracts boundaries to delineate neighborhoods for all years. Since the Census redraws tract boundaries over time according to population changes, it is necessary for this analysis to have to be based on a standard, time invariant boundary definition. NCDB revises Census data from the native year boundaries, tracks the changes over time to the tract boundaries, whether they merged with another, split, etc., and standardizes all years from 1970 - 2010 to the 2010 boundaries. At the neighborhood level, the NCDB for each year contains information on average household income, education levels, and age, race and gender compositions. It also contains information on the housing stock characteristics of the neighborhood, including the percentage of housing structures that are multi-unit or single dwellings, the number of bedrooms,

---

[18]IPUMS-USA, University of Minnesota, www.ipums.org. 2010 data is from the American Community Survey (ACS), 2000 data is from the 5% sample, 1990 data is from the 1% sample; 1980 data is drawn from the 5% state sample, and 1970 data is drawn from the 1% Form 1 Metro sample.

[19]http://geolytics.com/USCensus,Neighborhood-Change-Database-1970-2000,Products.asp

the age of the housing structures, along with vacancy rates, ownership rates, and data on the predominant modes of commuting transportation such as by automobile, public transit, or walking. In many instances, a schematic is needed to arrange neighborhoods as they are spatially ordered, and for this, I use the standard Census Tiger/Line®shapefiles for 2010.[20] The set of Census tracts spanned by the metropolitan areas consistently identifiable in the IPUMS samples contains 47,000 neighborhoods within 330 metropolitan areas. I use the CBSA codes available in NCDB, which are assigned uniquely to a 2010 Census tract, to identify a given metropolitan area.

1.4.2   Estimating Regional Gender Wage Gap Components

The gender wage gap constitutes the main independent variables of interest and must be estimated. Importantly, I want to estimate these gaps within a region, and I want to do so flexibly. With a standard Mincer regression that adds a female dummy variable there are potential omitted factors that, beyond experience and education, may influence wages, be correlated with gender, and therefore bias the wage gap estimates. Bias in these estimates translates into measurement error in regressions where the gaps are used as regressors. In order to minimize this issue, I include a large set of potential controls, including number of hours worked, race and native-born dummies, a set of occupation and industry characteristics at the individual and regional level, sector dummies, and a set of second order and interaction terms.[21] However, instead of using this comprehensive set of variables identically for each year and region, I use the double selection procedure from Belloni, Chernozhukov, and Hansen (2014) to select a subset of these variables, and therefore a specification, specific to each region. Without this regularization procedure my estimates would face high dimensional problems that arise from varying region sizes (and the corresponding variation in sample

---

[20]https://www.census.gov/geo/maps-data/data/tiger-line.html.

[21]A nice meta-analysis of the literature on estimating gender wage gaps and the types of variables and approaches typically used can be found in Weichselbaumer and Winter-Ebmer (2005).

size). The double selection procedure automates the specification search across regions, mitigating the high dimensionality issue.

Let the full set of potential individual controls be denoted by Z_{it}, where now i denotes an individual, rather than the female in a couple. F_i denotes gender, equal to one for female, and M_{it} denotes marital status, equal to one if ever married and zero otherwise. I estimate γ using the following approach:

$$\ln w_{ikt} = \mu_{kt} + \gamma_{kt}(1 - F_i) + f_{kt}(Z_{it}) + \varepsilon_{ikt} \quad \text{if } M_{it} = 0 \tag{1.3}$$

Here, $f_{kt}(\cdot)$ represents a set of coefficients on variables selected by a double selection LASSO procedure across the set of variables Z, specific to a region-year (k, t). LASSO is a regularization procedure that adds an additional constraint to the minimization problem of OLS that penalizes the magnitude of the coefficients. The double selection runs LASSO of $\ln w$ on all Z, then F on all Z, and takes the intersection of the two chosen sets and runs the full model with these variables (which constitute f_{kt}).

Equation (1.3) is estimated for singles only, allowing the male and female wages to differ only by γ. Denote the estimated γ as g_{kt}. Figure 1.7 plots g_{kt} over time, along with confidence intervals that display the variation across cities in a given year. Common to most estimates of the gender wage gaps, the estimates decline significantly over time. However, different from many estimates, these gender wage gap estimates are much lower than less robust, non-econometric estimates, and begin in 1970 on average at 15% and decrease to near zero by 2010, with some metropolitans areas having single gender wage gaps that are negative by 2010, meaning that women earn a premium relative to men. The magnitude of these estimates is influenced in part by the flexibility of the LASSO procedure, but also heavily by the limitation to single women. The finding that the gender wage gap among singles in nearing zero in recent years is consistent with other work that decompses the

gender wage gap by marital and motherhood status such as Kleven, Landais, and Sogaard (2015).

1.4.3 Alternative Gentrification Theories

There are a number of theories of urban migration patterns that are necessary to control for. These stem from both the suburbanization literature and the more recent literature on urban revitalization and include:

- *Transportation* - The prevailing means of transportation among household types likely influences residential location decisions if the cost per unit of distance is different across types. LeRoy and Sonstelie (1983) outline a plausible, transportation based story for gentrification through the evolution of transportation technology, available earlier to the rich than to the poor, allowing them to commute further at lower or equal cost, with a return to central areas once the technology proliferates and their advantage is nullified. Baum-Snow (2007) explains how highways constructed to access suburban areas opened up the true potential of time-saving commutes by automobile, benefiting suburban areas. At the neighborhood level, I control for the share of residents using different types of transportation to commute, including walking, public transportation, and automobile.

- *Housing Stock* - Housing stock characteristics may influence gentrification in that certain types of housing structures are more desired by higher income individuals. Besides the standard desire for larger and detached housing units, research suggests that the age of the housing stock and the number of units in a structure play an important role in residential decisions, especially as related to gentrification. Brueckner and Rosenthal (2009) explain how the affluent are attracted to new housing, which initially arises in suburban areas as cities expand outward, and eventually, once the metropolitan area has reached its peak expansion, shifts to development in the declined city center, and hence new housing and wealth theoretically relocate downtown.

Recent empirical work by Boustan et al. (2016) explores how the increased ability to own housing in multi-unit buildings, in the form of condominiums, has influenced gentrification patterns. To account for these potentially confounding influences I include controls for the age of the housing stock, house size, and the number of units in a structure.

- *Racial Composition* - The minority share of a neighborhood can have a dramatic effect on its income level and the types of in and out-migration it experiences. Tipping point analyses, originating largely from Schelling and tested recently by Card, Mas, and Rothstein (2008) in a regression discontinuity design setting demonstrate that neighborhood racial composition can "tip" based on the share of minorities. In related work in the suburbanization literature, Boustan (2010) documents that the each black arrival following postwar migration from the south led to 2.7 white departures from cities. To account for these racial influences I control for the minority share of residents in a neighborhood.

- *Amenities* - Positive amenities, including restaurants, bars, varied topography, a coastline, and negative amenities, including crime and pollution, have been widely proposed as influencing gentrification patterns and urban decline and resurgence. Crime is an alternative explanation that is still undergoing study with respect to gentrification patterns. The drop in urban crime rates from 1990 - 2010 is of comparable magnitude and cultural importance to the decline in the gender wage gap. While Cullen and Levitt (1999) and others documented that city decline prior to 1990 was in part attributable to differences in crime rates between the suburbs and cities, the current research on the influence of declining crime rates on gentrification is far from definitive. Edlund, Machado, and Sviatschi (2015) have a brief discussion of their results as being robust to splitting the sample by crime statistics, and Couture and Handbury (2015) find mixed results when assessing how the young, college-educated respond to changes in the

urban-suburban crime gap. Positive amenities have also quickly developed in gentrifying areas, including restaurants, bars, and music venues. These amenity changes, in either direction, can potentially confound the estimates. One limitation to the study of these issues is the data availability. The primary source of crime data is the FBI's Uniform Crime Reporting (UCR), which is only available back to the late 1980's and at the municipal, rather than neighborhood, level. In recent years, newly available data on amenities from sources like Yelp has been used to explore gentrification since 2000, but this information again is not available over a longer historical period.

To overcome this issue, I adopt an approach similar to Guerrieri, Hartley, and Hurst (2013), who propose a theory of "endogenous" gentrification, meaning that the areas most likely to gentrify are those adjacent to high-income areas, and the spillover amenities effects from these high-income areas benefit the circumstances in nearby neighborhoods. In this case surrounding neighborhood income is a proxy for the amenities, both positive and negative, that are associated with gentrification. Since neighborhood income data and the "network" structure of neighborhoods is available for all years, I adapt the analysis into a spatial econometric framework, where the income percentile of the surrounding neighborhoods is included as a spillover effect: $\sum_{l=1}^{L_k} s_{ml} y_{lt}$, where l denotes a neighborhood, L_k the total number of neighborhoods in region k. s_{ml} allows for spatial and time dependence between the neighborhood income percentiles of m and l. To account for the spatial dependencies embedded in this neighborhood "spillover" effect, I estimate this using spatial econometric techniques that use distant neighbors as instruments for these terms.[22]

I include the variables that account for these alternative theories in the X matrix when estimating Equation (1.2) above.

---

[22]These account for the obvious simultaneity bias in the spillover term.

1.4.4 Endogeneity Concerns in the Single Gender Wage Gap

The flexibility of the LASSO procedure allows for the estimates of the single gender wage gap, g_{kt}, to be robust to the standard array of explanatory variables cast as possible confounding factors which, once accounted for, may explain away the gender gap. Additionally, the set-up is structured to separate out the wage gap while single versus the transition that occurs for women after marriage and childbirth. What remains in g_{kt} is a level of the gender wage gap that simply cannot be recovered by any observable, labor market channel for women: not by schooling, or additional experience, or an alternative selection of industry or occupation. This is precisely why these estimates of the gap are useful in this context. g_{kt}, no matter the underlying source, is a resource attainable only through marriage. This is the empirical counterpart to the theoretical relevance of γ in the marriage decision outlined in Section 1.3.

Given that g_{kt} is the primary focus, a discussion is warranted over what it is capturing, and in turn the potential sources of endogeneity that may arise from testing its role in explaining gentrification patterns. Of particular concern is any story wherein an unobservable factor influences g_{kt} as well as the gentrification outcomes in consideration. The primary candidate alternatives have already been discussed, but to the extent this list is incomplete, or the methodology insufficient, general omitted variable bias remains a concern. Also, g_{kt} is an estimated object, and therefore subject to measurement error issues. Any sources of bias in Equation (1.3) may contribute to this issue, and while the LASSO helps mute this concern, measurement error in g_{kt} cannot be ruled out. Additionally, reverse causality is a potential concern. The most reasonable story is that higher income neighborhoods, which in this empirical setting are high amenity by definition, may be sufficiently desirable so that people prefer to stay single simply to remain urban dwellers, in the causal sense of higher city income resulting in delayed marriage. This set-up accounts for this in two ways. First, marriage is not directly in the specification but rather the mechanism through

26

which the gender wage gap is theorized to influence gentrification. While it's certainly true that the parameters being estimated are reduced-form with regards to the effect of changing the age of marriage, this reverse causality concern extends less obviously to the gap itself.

Cumulatively these concerns amass enough doubt about the exogeneity of g_{kt} that an instrumental variable analysis is constructive. Moreover, beyond these concerns, the thought exercise of what determines g_{kt} is informative. Recent work by Beaudry and Lewis (2012), Shenhav (2015), and Autor, Levy, and Murnane (2003) motivate a valid instrument for the gender wage gap. All are based on shift-share instruments, which is common in recent urban work on gentrification (Couture and Handbury (2015), Edlund, Machado, and Sviatschi (2015)). First, Beaudry and Lewis (2012) remark on the observed negative relationship between the gender wage gap decline (overall, not only for singles) and the rise in the return to education. They find evidence that suggests there is a skill-set common to women and the well-educated, and less prevalent for men and individuals with lower educational attainment. A growing literature points to this skill set as "soft-cognitive" skills related both to one's cognitive capacity and ability to socialize and work with people. As technological innovations have increasingly displaced the tasks historically requiring "hard" skills, returns to "soft" skills have increased (Deming (2015)).

Any major shocks to the relative prices of these skills versus so-called "hard" skills, typically related to manual labor, for example, the adoption of personal computing technology, plausibly induces the gender wage gap to contract if the softer skills are more abundant in women. A straightforward instrument that follows from the endogenous technology adoption literature is the initial skill mix in a region and year with respect to soft versus hard skills, adjusted for the relative returns to these skills at the national level. I construct this *skill mix* measure for years 1980 - 2010 when computer technology was proliferating, and use it to instrument for g_{kt} above.[23]

---

[23]The skill mix accounts for the aggregate labor hours in a city and year relative to the hours worked by a high education cohort, adjusted by parameter estimates at the national level of the relative returns to education. While

The skill mix captures how the composition of the labor market in a given region and year (k, t) with respect to the mix of soft-cognitive skill, a measure of human capital accumulation, and hard, "brawn" manual labor. The intuition is that places that have an initial labor market composition that is higher in soft skills will be more receptive to shocks that shift the relative input prices of these skills for firms. An important technological shock to the relative prices of soft versus hard-skill labor is the adoption of PC technology, which broadly served as a substitute for many hard-skill tasks and a complement for soft skill labor. Therefore a labor market with a higher composition of soft-skill labor will more rapidly adopt technology that complements the soft skills more abundant in women, and this measure will reflect broad, national trends in the relative price changes of these skills rather than region-specific shocks that may also correspond to migration patterns.

The "wage proxy" instrument from Shenhav (2015) is similarly structured and motivated. Her instrument is computed by beginning with the 1970 share of occupation-industry bins in a given region. These shares then grow at national, within-industry rates, and are used as weights in calculating the proxied wage for a region by using national industry wages in each year. This is done for each gender, and then the difference between male and female measures is referred to as the *relative wage proxy*, which I compute for the metropolitan regions in the NCDB. I also use this as an instrument for the gender wage gap in the specification above. See Figure 1.13 for plots of the correlations between these measures and the gender wage gap.

1.4.5   Results

In the tables below I present results from this estimation outlined in the previous section. I do so for three different measures of g_m. In Table 1.3 d_m is the distance in miles of neighborhood m from the city center. It includes results from estimation of Equation (1.2), which uses a lag of all

---

it is standard for education to link to cognitive returns, social skills are also believed to develop meaningfully in educational settings (Lleras-Muney, Miller, and Sheng (2017)).

covariates, fixed effects by year and metropolitan area, and adjustments for spatial dependence in income. First, consider income percentile as the outcome. The coefficient on the gender wage gap standalone term is negative and significant. This effect indicates that as the gender wage gap grows the income percentile of neighborhoods in the city center (i.e. distance equal to zero) decline on average. A one standard deviation decline in the gender wage gap, equal to 0.2, is associated with a two percentile improvement in neighborhood income percentile immediately at the city center. As stated above the main focus is the interaction between the gender wage gap and distance which, as predicted, is positive and significant: when the gap is larger neighborhoods further out improve in the income distribution, with the reverse holding in more equitable circumstances. This finding is robust to year and MSA fixed effects, as seen in specification (2), and the broad set of alternative gentrification controls in (3). Two-stage least squares estimates using both the skill mix and relative wage proxy are also broadly consistent, although both coefficients increase relative to the estimates from other specifications.

To clarify the nature of this effect and the intuition behind the interaction term, see Figure 1.12. This plot demonstrates the heterogeneous effect on neighborhood movement within the income distribution by distance to the city center for a decline in the single gender wage gap of 20 percentage points – approximately the drop in the single wage gap from 1970 to 2010. At each distance the band plots the range of effects estimated in the specifications presented in Table 1.3. Immediately at the city center, neighborhoods are predicted to improve on average anywhere from two to five income percentiles as the gender wage gap drops. In contrast, neighborhoods 15 miles from the city center are predicted to fall two income percentiles. This improvement in income is substantial: two to five percentiles accounts for 25-62.5% of the entire average improvement of city center neighborhood income over the period of interest (see Figure 1.4). Next I discuss set of robustness checks and alternative specifications for these analyses:

- *Addition of Alternative Gentrification Controls.* In Table 1.4 I present a set of estimates that begin with the baseline specification and add each set of gentrification controls one at a time. The results vary minimally with the iterative inclusion of these controls.

- *Skill Mix and Wage Proxy Instruments.* In Table 1.5 I present a set of two stage least squares estimates using the skill mix and wage proxy as instruments alone, and then jointly. These estimates are again largely consistent with the main results presented above, however, the coefficients increase noticeably relative to OLS estimates when using either instrument alone.

- *Other Gender Wage Gap Measures.* In addition to the gender wage gaps while single used above, I also test these estimates with a variety of other measures of the gender wage gap, including a raw male to female wage ratio in a metropolitan area, basic Mincer estimates, and a standard Oaxaca decomposition estimate; all estimates with these variables are comparable to those presented above, as are results using the percentage of single women with a college degree as a proxy for the presence of women in the upper echelons of a region's labor market.

Alternative Measures of Gentrification

While the preferred measure of gentrification for this work has been the income percentile of a neighborhood, I briefly investigate whether similar patterns with respect to the gender wage gap hold for other gentrification metrics. Other work focuses on housing prices and rent, education, and the minority share of a neighborhood. Estimates of the gender wage gap influence on these measures can be found in Table 1.6. The same pattern holds for rent percentile and education, although results are minimally significant for education. The results for minority share move oppositely, which is consistent with either displacement or a dramatic influx of white residents. Minority share directly at the center is not influenced by gender wage equality. However, the coefficient on the distance and gap interaction, in contrast to the other outcomes, is negative, suggesting that decreases in the

30

gender wage gap lower the minority share in neighborhoods closer to the city center.

1.5   Conclusion

Cities in the United States have finally started to experience redevelopment following decades of twentieth century decline. This redevelopment has mainly manifest in the gentrification of centrally located neighborhoods that are increasingly high income, high rent, and well educated, while shrinking in their share of minorities. The transition of some neighborhoods, like the Penn Quarter in Washington, DC, has been swift and dramatic. The costs and benefits of this form of urban transformation are becoming a critical issue for urban policymakers and municipal officials in the cities around the country.

The prolific monocentric city model focuses on the fundamental trade-off between housing and commuting, Starting with a basic distinction between married and single households, I highlight an important driver of residential decisions: marital status of a household and the prospects they face for having children, which pushes them further from the city center relative to singles. As pecuniary incentives for marriage declined, and women gained substantial financial independence from men, the population of urban dwelling singles increased substantially. College educated individuals played a major role in this demographic change. During the later decades of the twentieth century the average age at marriage was nearly identical to the average age a person obtained a bachelor's degree, such that in expectation marriage – and suburban living – were nearly immediate upon college graduation. These labor market changes altered the composition of single people who prefer city living, such that in recent years in cities that have made remarkable strides in remedying certain forms of gender wage inequality, neighborhoods near the city core have experienced tremendous revitalization. As the gradient of development pivots back toward urban living, there are policy issues that persist. While the evidence is strong that affluence is returning to American cities, it is

arriving at a high price that is excluding many long-time residents from partaking in the renaissance of their own neighborhoods. In fact many of these changes negatively impact the share of minority residents in cities, and are associated with increased suburbanization of lower income, minority populations. This could potentially lead to even more severe residential segregation and spatial mismatch among these populations, reinforcing the very issues some hope re-urbanization is poised to address.

Figure 1.2: Income Percentile by Distance to City Center, 1970 - 2010
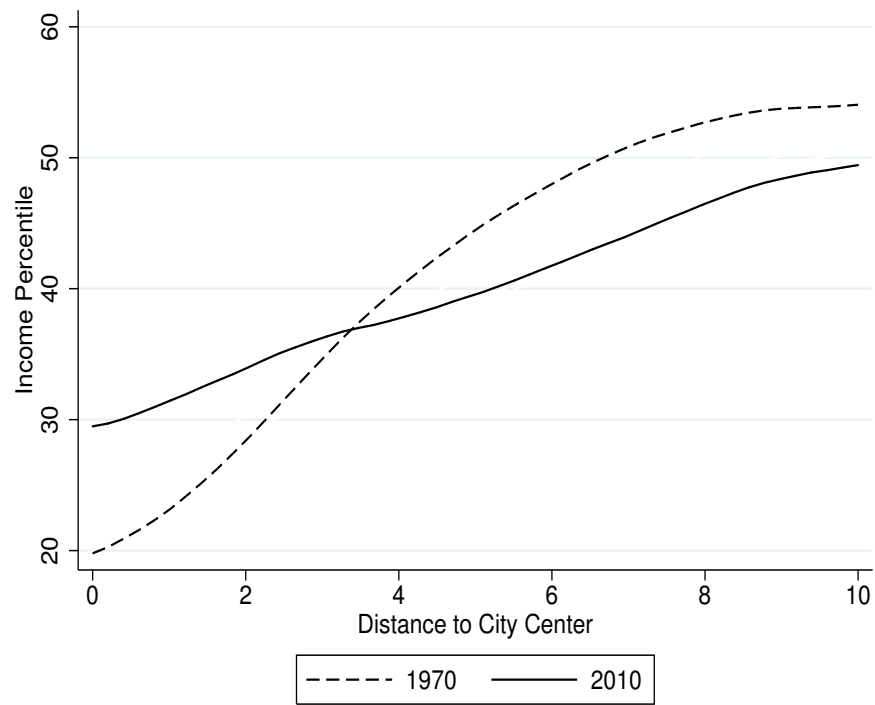
Figure 1.3: College Educated Percentage by Distance to City Center, 1970 - 2010

Share with College Degree is standardized within each metro area and year.
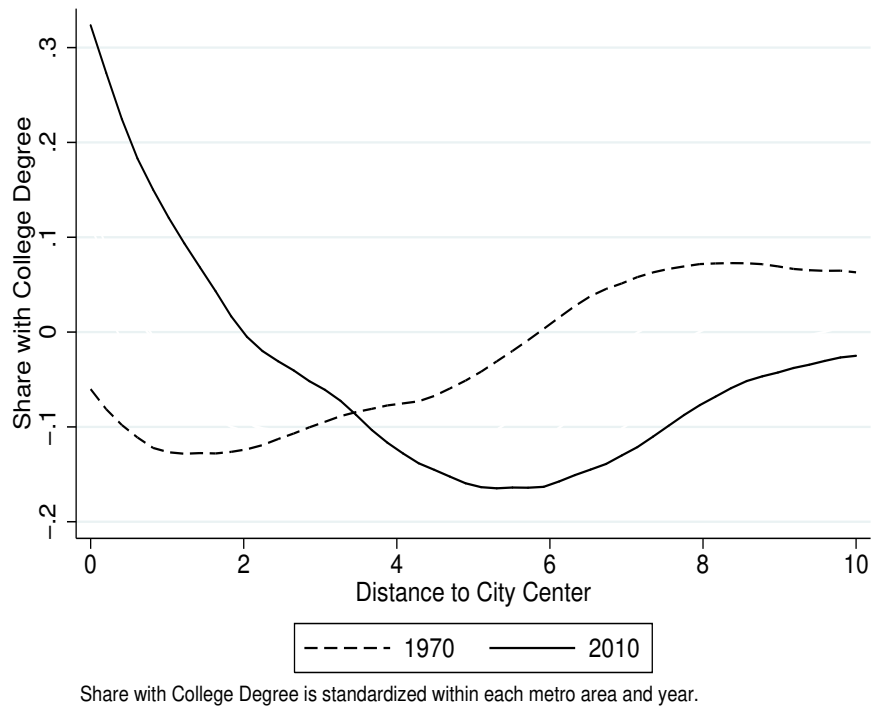
Figure 1.4: Trends in Gentrification and Gender Wage Gap
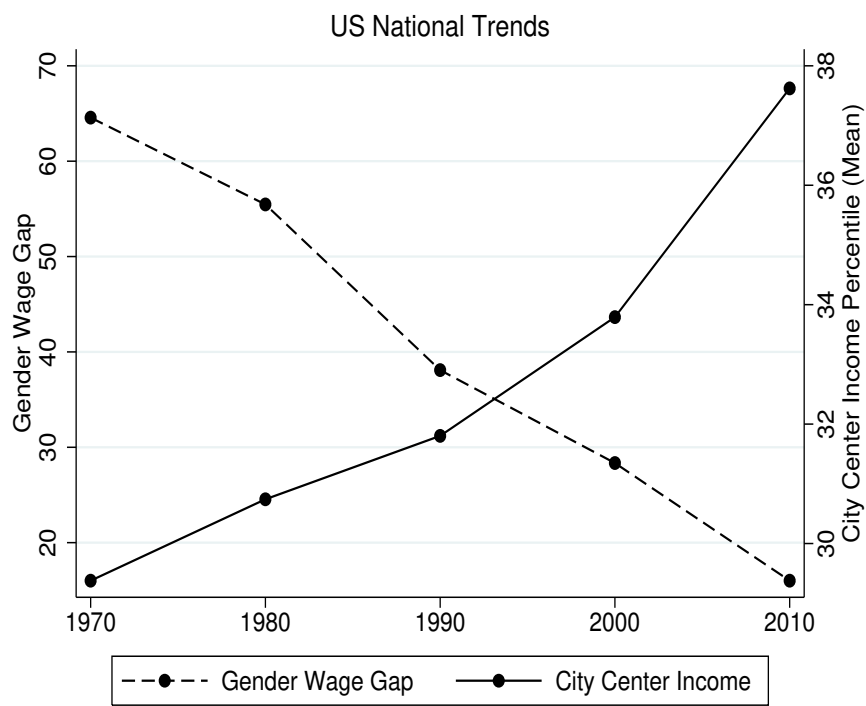
Figure 1.5: Gentrification and Gender Wage Gap - New York, NY and Chicago, IL

New York, NY and Chicago, IL
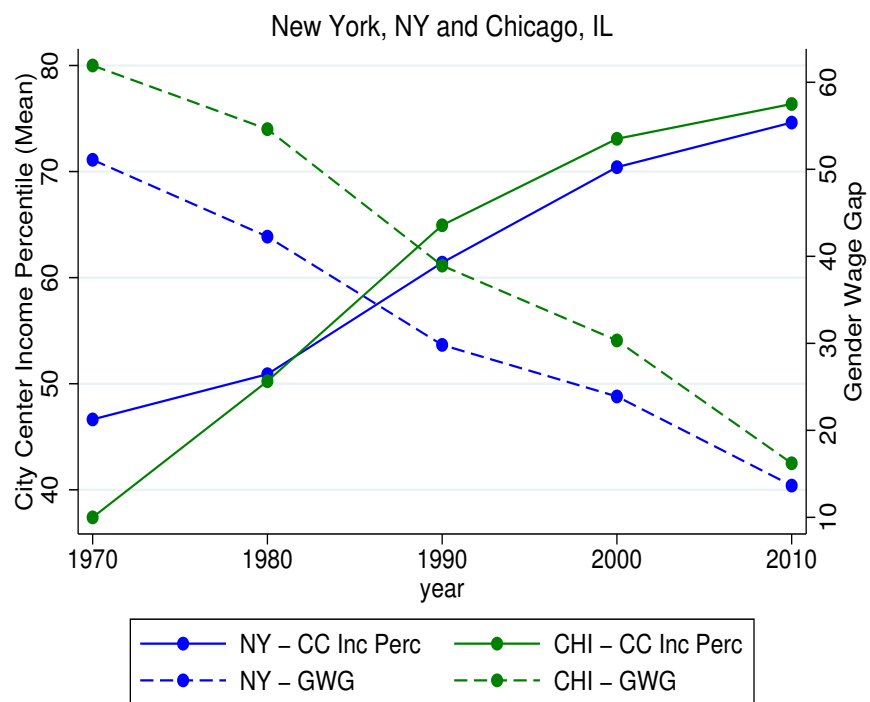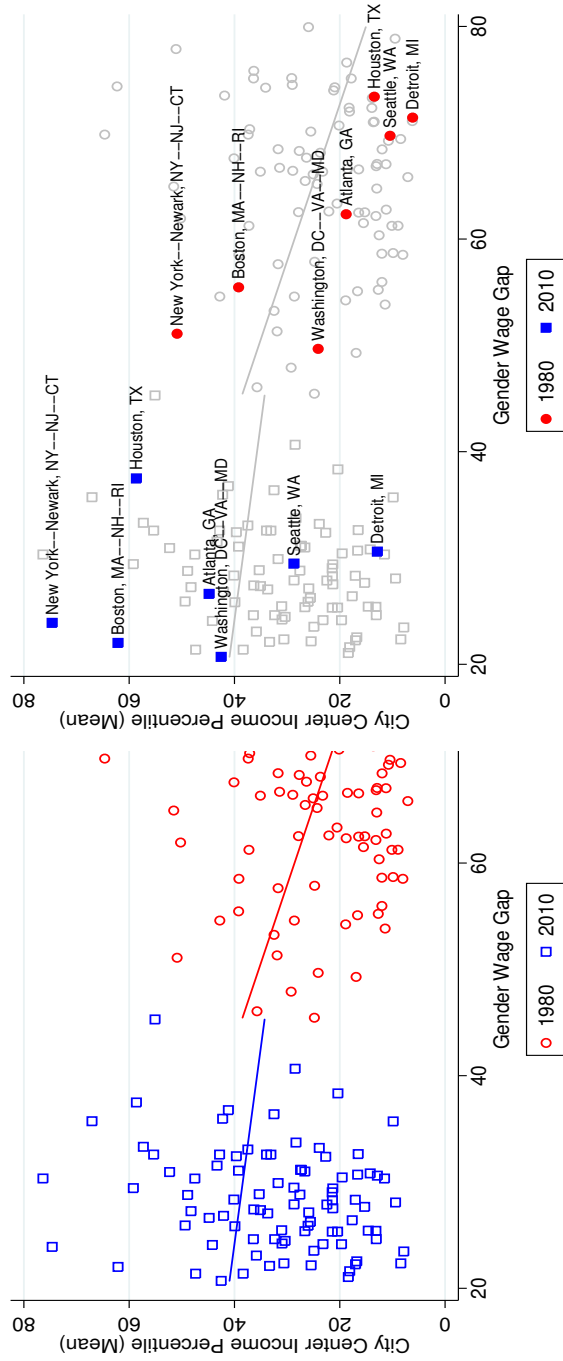
(a) Top 100 Cities by Population

(b) Selection of Top Cities

Figure 1.6: Association between Gender Wage Gap and Gentrification in the Cross-Section

Figure 1.7: Estimates of Gender Wage Gap While Single
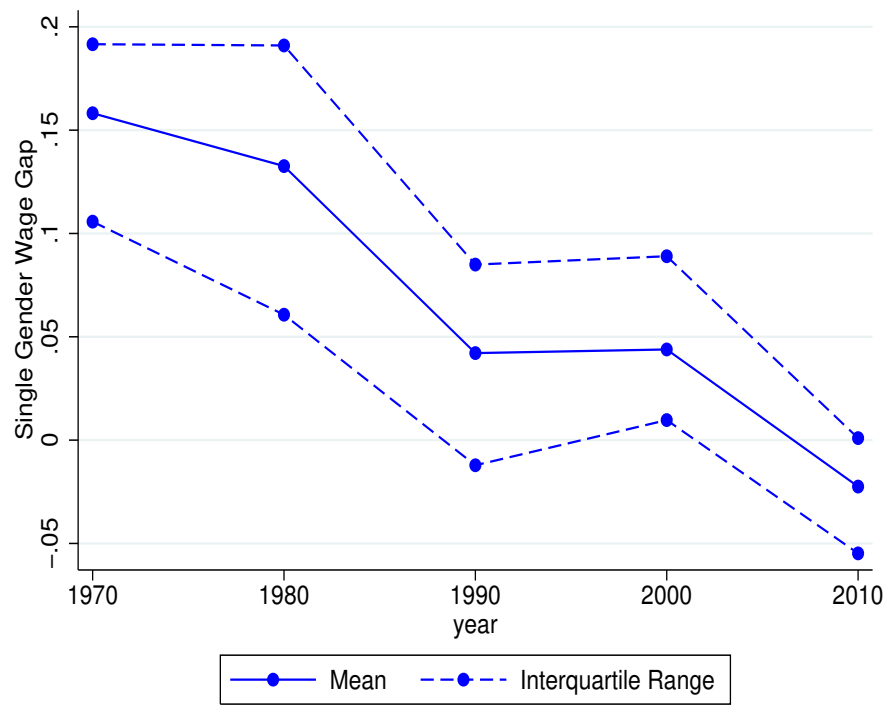
Figure 1.8: Single vs Married Pct of Households (by Distance from City Center)
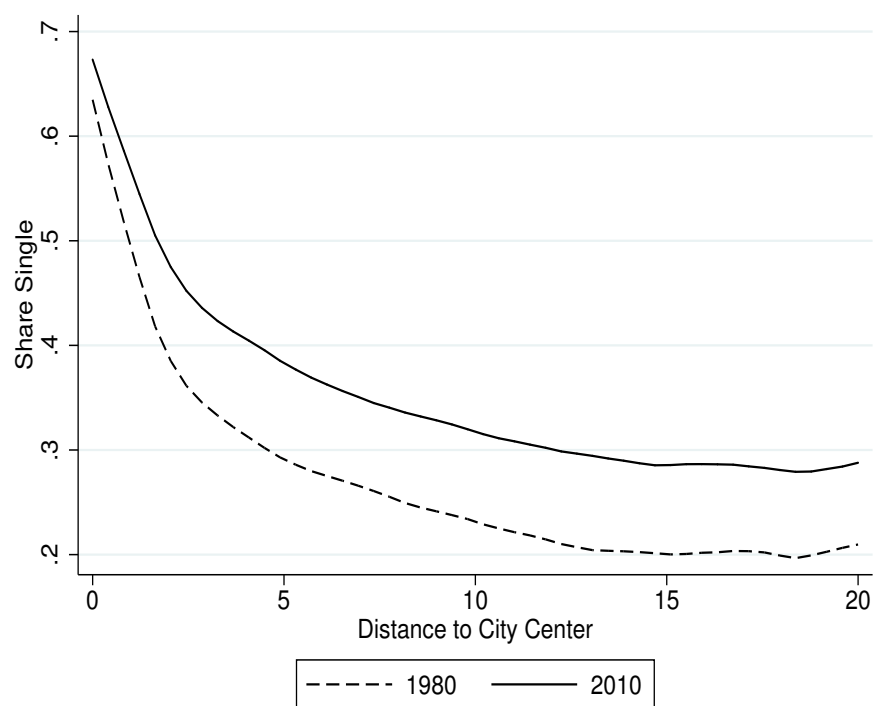
Table 1.1: Difference in Probability of Living in Central City

	(1) General Employment	(2) High Prestige Employment
(1) Married vs. Single	-0.0228***	-0.0374***
	(-3.60)	(-5.27)
(2) Lone vs. Dual-Earner	0.00569*	-0.000150
	(2.58)	(-0.04)
(3) No Children vs. Children	-0.00281	-0.00801
	(-1.21)	(-1.93)
Observations	1545462	1545462
R^2	0.039	0.039

t statistics in parentheses

High Prestige employment based on occupation in highest quartile of Siegel career prestige scale.

Regressions include interactions of employment, marital status and number of children, in addition to educational attainment, age, and income measures.

Specifications also include metropolitan area and year fixed effects, with cluster-robust standard error.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 1.2: City-Suburban Migration and Recent Marriage

	(1) No Children	(2) Children
Married within Last Year	0.0656***	0.0559*
	(4.16)	(2.17)
Observations	12335	3274
R^2	0.018	0.030

t statistics in parentheses

Sample limited to women between 18 and 65 who, in previous year, lived in central city.

Includes years 2008-2011.

Additional controls include employment, income, education, and age.

Includes metropolitan area and year fixed effects, cluster-robust standard errors.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Figure 1.9: Male to Female Ratio (US), Ages 16-34

Male–female ratio for population 16–34.
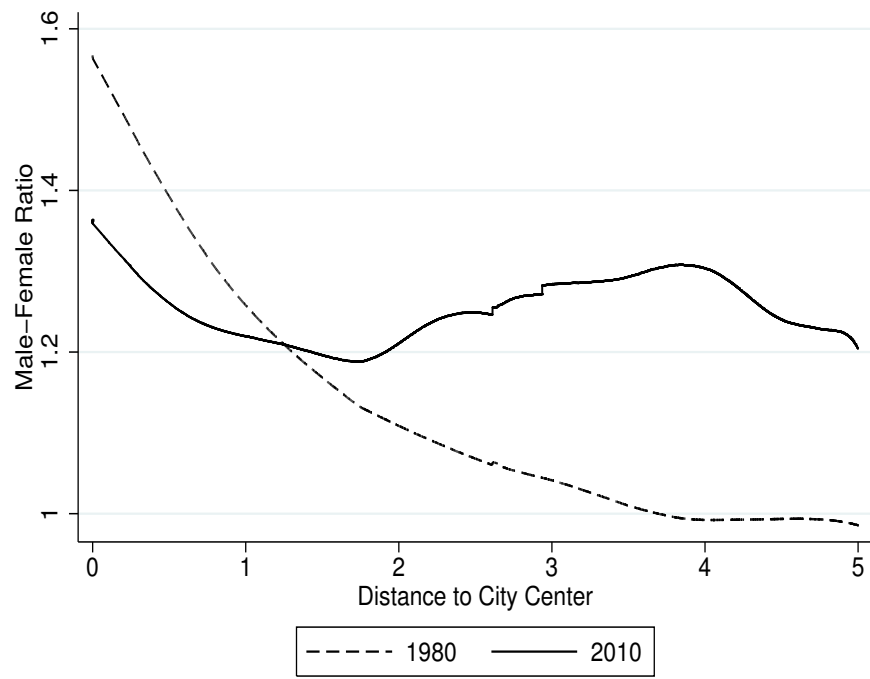
Figure 1.10: Male to Female Ratio (NY), Ages 16-34

New York, NY
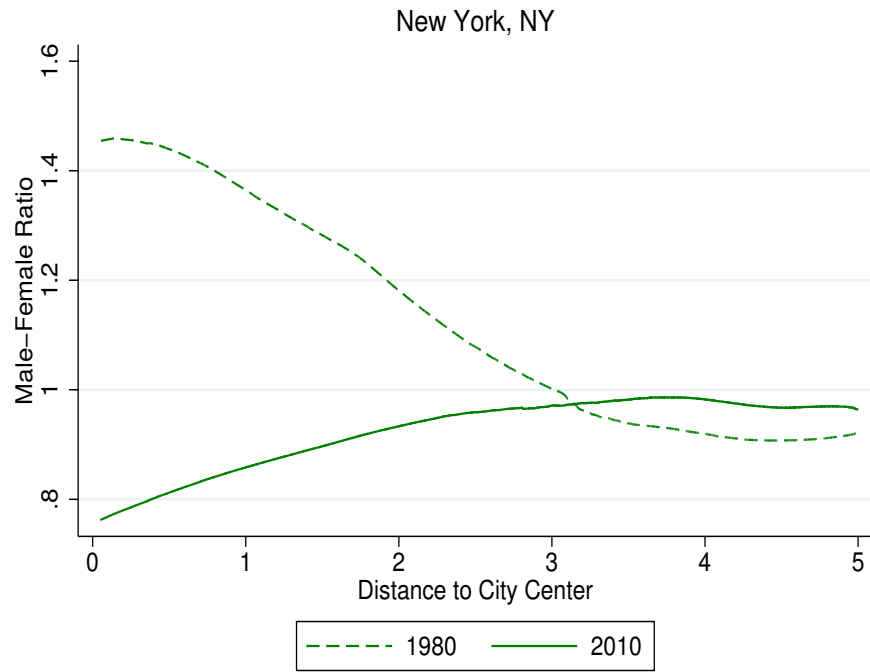
Male−female ratio for population 16−34.

Figure 1.11: Male to Female Ratio (DC), Ages 16-34

Washington, DC
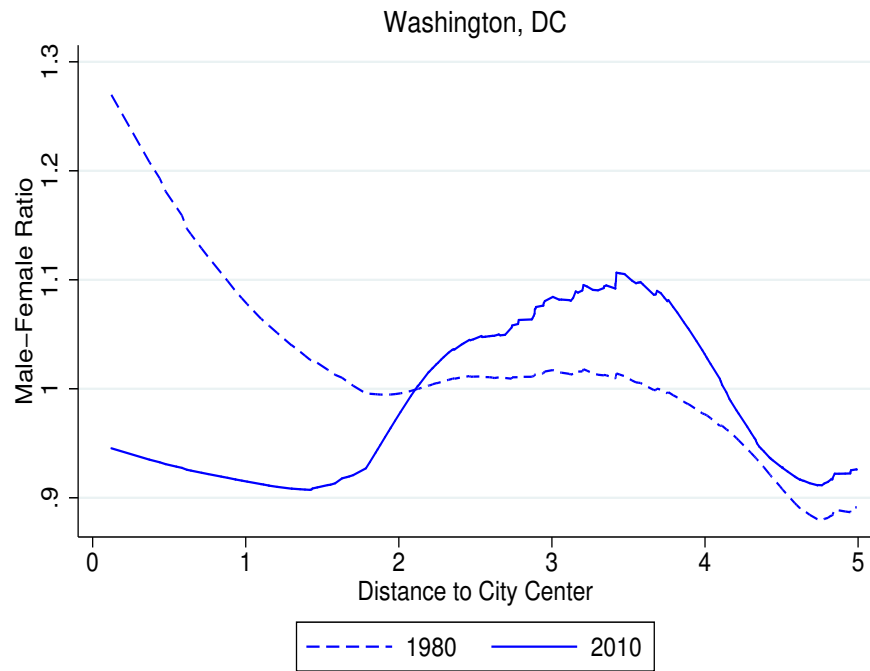
Male−female ratio for population 16−34.

Table 1.3: Gender Wage Gap Effect on Gentrificaton

	(1) Income Percentile	(2) Income Percentile	(3) Income Percentile	(4) Income Percentile
Single Gender Wage Gap × Distance	1.396***	1.624***	1.575***	2.331***
	(0.0659)	(0.369)	(0.269)	(0.663)
Single Gender Wage Gap	-9.954***	-13.47***	-12.19***	-23.20***
	(0.795)	(3.035)	(2.302)	(5.622)
Observations	146480	146480	143975	108999
R^2	0.692	0.694	0.740	0.786
Year FE	No	Yes	Yes	Yes
Metro FE	No	Yes	Yes	Yes
Alt Gentrification Controls	No	No	Yes	Yes
IV for Gender Wage Gap	No	No	No	Yes

Standard errors in parentheses

Gender wage gap while single estimated by LASSO procedure for each metro area in each year.

Cluster-robust standard errors by metro area.

* $p < .1$, ** $p < 0.05$, *** $p < 0.01$

Figure 1.12: Decline in Gender Wage Gap Effect on Gentrification (across specfications)
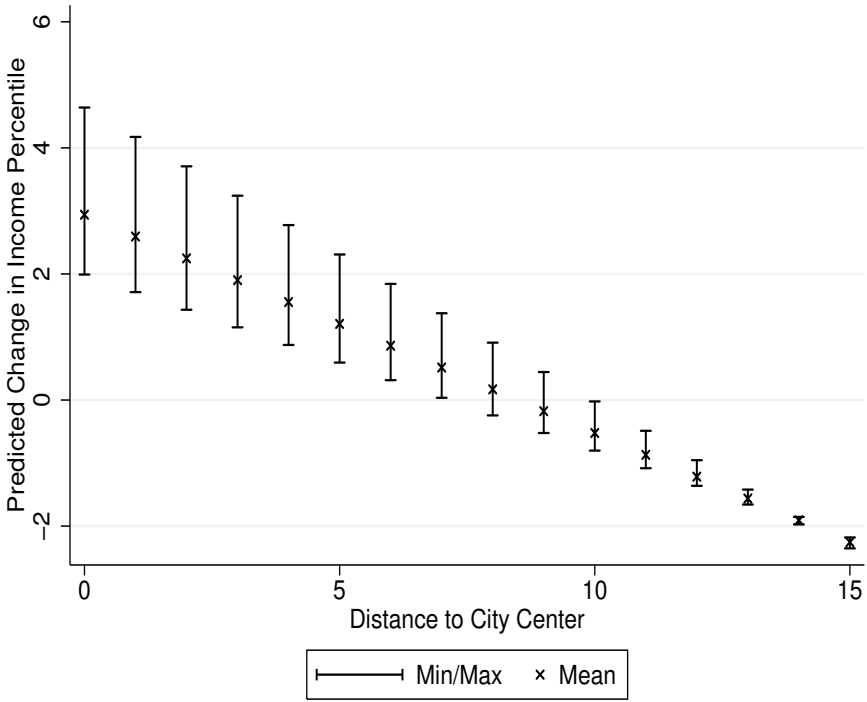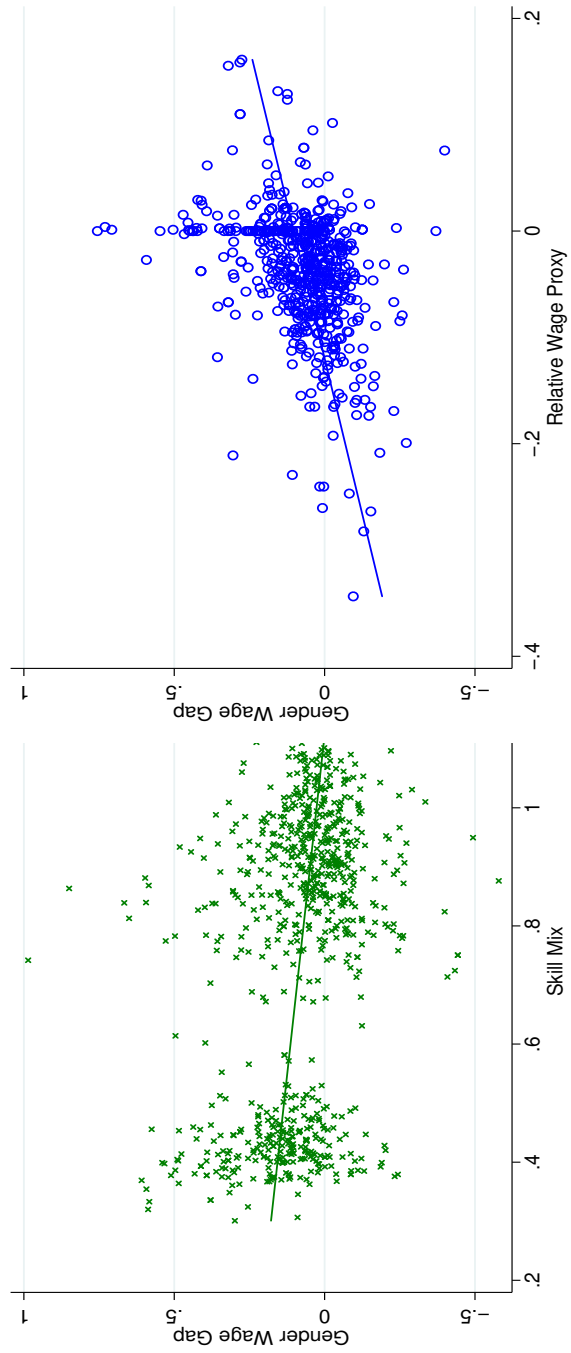
(a) Skill Mix

(b) Relative Wage Proxy and Skill Mix

Figure 1.13: First-Stage Relationships: Wage Proxy and Skill Mix

45

Table 1.4: Effect of Gender Wage Gap Robust to Alternative Gentrification Explanations

	(1) Income Percentile	(2) Income Percentile	(3) Income Percentile	(4) Income Percentile	(5) Income Percentile
Single Gender Wage Gap × Distance	1.624***	1.365***	1.347***	1.558***	1.575***
	(0.369)	(0.288)	(0.278)	(0.270)	(0.269)
Single Gender Wage Gap	-13.47***	-12.35***	-11.05***	-12.10***	-12.17***
	(3.035)	(2.362)	(2.308)	(2.314)	(2.280)
Observations	146480	144151	144151	143975	143975
R^2	0.694	0.722	0.727	0.739	0.740
Year FE	Yes	Yes	Yes	Yes	Yes
Metro FE	Yes	Yes	Yes	Yes	Yes
Transportation	No	Yes	Yes	Yes	Yes
Racial Composition	No	No	Yes	Yes	Yes
Housing Stock	No	No	No	Yes	Yes
Endogenous Gentrification	No	No	No	Yes	Yes

Standard errors in parentheses

Gender wage gap while single estimated by LASSO procedure for each metro area in each year.

Cluster-robust standard errors by metro area.

* $p < .1$, ** $p < 0.05$, *** $p < 0.01$

Table 1.5: Instrumental Variable Estimates for Gender Wage Gap Effect on Gentrificaton

	(1) Income Percentile	(2) Income Percentile	(3) Income Percentile	(4) Income Percentile
Single Gender Wage Gap × Distance	2.047***	2.596***	3.975**	2.331***
	(0.420)	(0.667)	(1.566)	(0.663)
Single Gender Wage Gap	-11.76***	-35.29*	-62.68*	-23.20***
	(3.735)	(21.03)	(33.58)	(5.622)
Observations	143975	108999	143975	108999
R^2	0.570	0.785	0.732	0.786
Year FE	Yes	Yes	Yes	Yes
Metro FE	Yes	Yes	Yes	Yes
Alt Gentrification Controls	Yes	Yes	Yes	Yes
Skill-Mix IV	No	Yes	No	Yes
Wage Proxy IV	No	No	Yes	Yes

Standard errors in parentheses

Gender wage gap while single estimated by LASSO procedure for each metro area in each year.

Cluster-robust standard errors by metro area.

* $p < .1$, ** $p < 0.05$, *** $p < 0.01$

Table 1.6: Gentrification Outcomes and Gender Wage Gap Components

	(1) Income Percentile	(2) Rent Percentile	(3) Perc Bach Deg	(4) Minority Share
Gender Wage Gap	-11.72***	-10.84***	-3.988	0.00871
	(-5.05)	(-3.89)	(-0.75)	(0.00)
Gender Wage Gap × Distance	1.536***	1.348***	1.177*	-0.714**
	(5.69)	(4.25)	(2.04)	(-2.86)
Observations	148311	112804	148118	148311
R^2	0.734	0.593	0.819	0.878

t statistics in parentheses

All specifications include metropolitan area and year fixed effects.

Cluster-robust standard errors by metro area.

Additional controls include specifics on minority share, mode of commuting, number of bedrooms,

units per structure, age of housing stock, and a lag in the dependent variable.

Estimates account for neighborhood-level spillovers in income.

$^+$ $p < .1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Chapter 2

Urban Gentrification: The Role of Condominiums

2.1 Introduction

For several decades after World War Two the population density gradients of American metropolitan areas "flattened" as a relatively larger share of urban residents chose to live in the suburban ring rather than the central city. Initially, postwar suburbanization was not neutral with respect to socioeconomic characteristics; relative to pre-WWII differences, postwar suburban residents were better educated, had higher incomes and wealth, and were more likely to be white than central city residents (Margo (1992)). From 1980 to the present, however, the population shares of central cities have been increasing, and socioeconomic disparities between the urban core and periphery have begun to narrow. That is, over the past three decades or so, higher SES households have increasingly chosen to live in central cities relative to the suburban ring – a process commonly referred to as "gentrification".

Various explanations have been offered to account for gentrification. On the demand side, for example, labor economists have argued that a rising opportunity of time spent commuting – a byproduct of rising wage inequality and increasing returns to schooling – has increased the demand for housing by high SES households in the urban core, particularly in cities with efficient public transportation systems (Kahn (2007); Guerrieri, Hartley, and Hurst (2013); Edlund, Machado, and Sviatschi (2015)). Complimentary to this explanation is that urban cores increasingly offer cultural and consumption amenities – museums, the arts, fashionable restaurants and shopping – that have high income elasticities (Couture and Handbury (2015)). On the supply side, gentrification could also occur in response to a reduction in the cost of supplying housing desired by higher income, better educated consumers in the urban core relative to costs in the suburbs. The goal of this paper is to isolate the causal effect of one possible supply-side factor, the legalization post-1970 of what was for the United States, a novel housing structure-type – the condominium.

By the 1950s provisions of the US tax code favored owner-occupancy for higher income house-

holds in most cases regardless of the type of structure, unless offset by other costs that would make ownership more costly than renting (Hansmann (1991)). However, prior to 1960 owner-occupying a single unit within a multi-family dwelling without simultaneously owning the entire building was essentially impossible except in a very few metropolitan areas – in which the cooperative apartment was a legally permissible structure type. This situation changed, however, in the early 1960s, when the condominium first became a legally recognized form of ownership and conveyance in the American housing market. By the end of the decade all states had passed condo enabling legislation. Thus, from ca. 1970 to the present, a high income household who wished to owner occupy had a third housing structure type – the condominium – to consider, along with traditional detached single family dwellings and coops.

Cities develop from the core outward. Historically, because of high commuting costs, housing in the urban core was much denser – that is, multifamily rather than detached single family dwellings. Depending on the age of the city, therefore, some urban cores were well-stocked with housing whose characteristics made it ripe for conversion to condominiums – a high "condo propensity" – while in other cores were less well endowed – they had a low "condo propensity". Variation in propensity exists across cities, as just described, but also across neighborhoods (census tracts) within cities.

To identify a causal effect of the condominium on gentrification, we seek exogenous variation across locations in condo supply, and then examine if this exogenous variation affects the SES characteristics of the residents. We identify such variation by considering how a particular factor – a municipal condo ordinance – affected the costs of condo conversion.

By a "municipal condo ordinance" we are referring to legislation passed at the city level whose intent was to protect tenants from condo conversions and, therefore, raise the cost and slow development – for example, by imposing protections against evictions, temporary moratoria, and advance notification requirements, among other provisions. In the absence of such an ordinance, we argue

that, post-1970, the percent condo will be an increasing function of condo propensity, but the elasticity will be less than one, because some multi-family housing will still be rental. In a city with a municipal ordinance, however, developers will concentrate their effects on neighborhoods where conversion is most profitable, if the ordinance is binding. In such cities, we expect to see a negative main effect of the ordinance on the percent condo, but only in neighborhoods where there is a high condo propensity in the pre-condo periods. We focus on the percent of housing structures with 5+ units in 1970 as our main metric for condo propensity. To bolster identification, we focus our analysis on central city and suburban neighborhoods near municipal boundaries, in cities with and without condo ordinances. By the usual type of "border" argument (Platt Boustan (2012), Black (1999)), we expect these neighborhoods to be very similar except for variation in the percent condo induced by cost variation, as just described. This gives us a first stage set-up similar in spirit to a DDD approach where we estimate an effect of the ordinances on condo rates. We find that the ordinances negatively impacted condo rates relative in neighborhoods within municipalities that passed restrictive regulations. We use a similar approach within city centers and find comparable significant, negative effects of the ordinances on condo rates.

We use the predicted value of percent condo from the first stage in a 2SLS regression of resident household income, education and racial composition. We find that an exogenous increase in condo availability is significantly associated with a higher average income of residents and a higher proportion of college graduates. Greater condo availability also seems to increase the share of black residents in a neighborhood, contrary to common concerns about displacement, but this relationship is not robust to our different specifications. These findings are robust to the exclusion or inclusion of a long list of control variables, and to the precise manner that we control for condo propensity in the first stage.

The remainder of the paper is organized as follows. Section 2.2 provides a capsule economic

history of condos in the United States. We show that, absent the availability of cooperative apartments, owning and occupying a slice of a multi-family dwelling was next to impossible in the United States prior to the 1960s, when the first condo enabling legislation was enacted. We also review the previous literature on gentrification, and a make a prima facie empirical case that the availability of a new housing structure type could be an important factor in urban gentrification. In Section 2.3 we present our data and identification strategy. Our empirical set-up and description of our estimation approach is in Section 2.4. The main results are presented in Section 2.5, and Section 2.6 concludes.

2.2 Condos: A Brief Economic History

Owner-occupancy rates of single family dwellings have always been quite high in the United States, even when agriculture was the primary economic activity. Well into the twentieth century farm owner-operators were home owners practically by definition because the farm homestead was part of the farm itself and the vast majority of owner-operators lived, as it were, where they worked. As labor shifted out of agriculture and the population became more urbanized the share of households living in multifamily dwellings increased. This shift cause aggregate owner-occupancy rates to decline from the mid-nineteenth into the early twentieth century (Collins and Margo (2011)).

There are many reasons why owner-occupancy rates were low in multi-family dwellings in urban America in the early twentieth century but one key reason was that it was very costly for households to owner-occupy a "slice" of a building. In theory, a person could always purchase an entire multi-family dwelling and live in one of the units, renting out the others; alternatively, a group of households could pool resources and live in the separate units. But these arrangements were evidently very uncommon; the vast majority of households living in apartment building in, say, Boston, in 1910 were renters.

There was one exception to the general rule and this was the cooperative apartment. The first cooperative apartment was built in Manhattan in the 1870s. From that point cooperative buildings were constructed elsewhere in the New York metropolitan area and occasionally in other cities, such as Chicago but nationwide were very uncommon. In a cooperative setting, a household does not have title to a specific unit bur rather owns shares in the cooperative, with a right to occupy a particular dwelling. For both entry and exit, coops were "gated communities" – a household could not simply buy shares at will but had to be admitted to the coop. Nor could the unit be sold per se; the buyer had to be approved by the cooperative. Compared with a single family dwelling the transactions costs of "owning" and occupying a cooperative apartment were relatively high.

Although it was difficult and therefore uncommon to owner-occupy a slice of a multi-family dwelling in the early 20th century US, such arrangements were actually known and somewhat frequent elsewhere in the world, particularly in European cities with a history of civil (as opposed to common) law. In early twentieth century Paris or Madrid, for example, one could, in fact, own a particular unit and also have the right to use common facilities and also to sell the unit. This institutional form in the real estate market was the condominium.

2.2.1   Urban Renaissance in the Era of the Condominium

Prior to the advent of condos as a legal form, there was a well-documented exodus of households, especially white, middle and high income, from urban centers following World War II. Mieszkowski and Mills (1993) articulate several explanations for suburbanization: federal mortgage insurance post-World War II, highway construction (and thus lower commuting costs), high urban crime rates, and poor city school systems. Boustan, Bunten, and Hearey (2013) surveys recent work estimating the causes of suburbanization, the most prominent being the rise of real income and a desire for land and housing (Margo (1992)), the construction of interstate highways, which reduced

the commuting times for suburban households (Baum-Snow (2007)), rising crime rates (Cullen and Levitt (1999)), and racial tensions and "white flight" (Boustan (2010)).

There is mounting evidence that since 1990 the income gradient with respect to distance from the city center has tilted away from suburbanization in favor of urban living. Glaeser and Gottlieb (2006) document that some of the larger, more dense cities in the US began to redevelop in the 1990s. Their explanation is that by facilitating social interaction certain cities have developed an array of consumer amenities that attract residents to the core. Their analyses demonstrate a higher propensity for urban dwellers to go to museums, out to a bar or dinner, a movie, and various types of musical events. The ability of urban consumers to enjoy these activities is obviously reduced in the presence of high crime rates, and their work mentions the reduction of crime in cities like New York and Chicago as potential reasons for the increases in urban population. Their analysis, while documenting many interesting observations, does not estimate these urban trends at a general level, nor does it attempt to disentangle the puzzle of which came first: lower crime, or better amenities. Glaeser and Shapiro (2003) focus on central cities and population growth generally across cities. They find that weather (e.g. high average annual temperatures in Sun Belt cities) and high human capital are highly correlated with central city growth.

Urban revival is often described in the context introduced above; namely that density facilitates social interaction, which leads to better consumer amenities, walkability, and accessibility (as portrayed by the terms "New Urbanism" and the "Consumer City", Glaeser, Kolko, and Saiz (2001)). These benefits attract residents relative to suburban offerings. However, the enjoyment of these amenities is lower in the presence of high crime rates. Drastic drops in crime rates in places like New York and Chicago are postulated to have led to increases in their respective populations (Glaeser and Gottlieb (2006)). Additionally, to the extent it is high income individuals who pursue higher quality amenities and accessibility, this clustering of income may itself induce gentrifica-

tion. Work by Guerrieri, Hartley, and Hurst (2013) considers models of endogenous gentrification wherein households have an explicit preference for locating closer to high income populations, hence gentrification emanates outward from high-income neighborhood clusters. Much of this literature neglects to consider the relevance of downtown areas and neighborhood heterogeneity in urban renewal, focusing instead on the central city as a whole. Glaeser and Shapiro (2003) explain that, "the return to city living is really about downtowns, not big metropolitan regions, so [central] cities make more sense as a unit of analysis." Even within a central city, never mind a MSA, there is extreme variation in population density, income, racial composition and housing prices, in addition to crime, pollution, and the availability of amenities. Miller (2017) demonstrates that analyses at the MSA and central city level mute much of the changes occurring in downtown areas, documenting that while in aggregate central cities do not appear to be gaining income relative to suburban areas, when a downtown is carved out within a central city, highlighting neighborhoods nearest major employment centers, these "city center" neighborhoods are improving dramatically in income relative to outlying suburban areas in numerous US cities. This is motivation to explore gentrification patterns at the neighborhood level. Couture and Handbury (2015) explore urban revival in the US between 2000 and 2010, focusing on the role of amenities. Edlund, Machado, and Sviatschi (2015) consider the role of more intense working hours among the highly educated as increasing the penalty for commuting and increasing incentives for living close to employment, and Miller (2017) explores the role of increasing female labor force participation and delayed marriage in shifting the income gradient in favor of dense, urban neighborhoods. This body of work largely relies on Bartik-style, shift-share instruments as the main sources of exogenous variation. Our exploration of the causal role of condominiums is also at the neighborhood level, with a notable contribution being that we introduce a novel source of exogenous, supply-side variation specific to condominium development.

2.3 Data and Identification

The neighborhood data analyzed econometrically in this paper come from two sources. The first source is the Census Bureau's Neighborhood Change Data Base (NCDB). Beginning in 1970 and continuing to the present, the NCDB provides census tract data using normalized tract boundaries as of a given date – in our case, 2010. In our base sample, we restrict our analysis to census tracts in the 100 largest cities (as of 2010); in most regressions, we further restrict the analysis to census tracts that are within a certain physical distance (for example, 2000 meters) of the municipal boundary of the parent city (e.g. Baltimore, Houston).

The NCDB provides tract level data on a set of housing characteristics variables – in particular, housing structure type. With some exceptions, these variables are all in percent form – for example, the percent of the housing stock that consists of detached, single family dwellings. However, while the information is collected by the Census, no waves of the public-use version of the NCDB provide information on the percent of housing units that are condominiums, which is our independent variable of interest. At our request, the Census Bureau made special tabulations of the 2010 data, and provided us with counts of condos for 2010 in each census tract in the NCDB, for which we can compute the percent condo at the neighborhood level. This data is sourced for 2010 from the American Community Survey. In Figure 2.1 below is an example of variation in the 2010 percent condo for two cities: Washington, DC and Baltimore, MD. The dark boundary is the municipal boundary for each city. Percent condo is highest at the city center for both cities, with notable upticks in select suburban neighborhoods on the opposite side of the municipal boundary.

No such data exist for 1970; however, aggregate data indicate that there were very few condos in the United States anywhere in 1970, and thus there is little harm in assuming that the percent condo in 1970 was zero in all of our census tracts. While condos were rare prior to 1970, many neighborhoods contained housing stock well-suited for future condo conversion. Certain housing

stock characteristics are highly predictive of future percent condo, including the age of the housing stock and the number of units in a housing structure. Figures 2.2 and 2.3 document this relationship. Figure 2.2 contains a similar map of Washington, DC and Baltimore, MD, with tracts colored by the percentage of housing units with five or more units (percent 5+). Like percent condo, percent 5+ tends to be highest in the city center, remaining high as neighborhoods emanate outward from the center to form a portion of a suburban ring around the city.

Beyond these example cities, the correlation between 1970 percent 5+ and 2010 percent condo persists. Figure 2.3 presents a nonparametric plot of 2010 percent condo on 1970 percent 5+ unit, including 95% confidence intervals. The two are significantly, positively correlated, in near-linear fashion, on average, for neighborhoods with more than 20% 5+ unit housing in 1970. Throughout the analysis that follows we limit our sample to tracts above a threshold of percent 5+ to ensure baseline comparability in neighborhoods.

Starting from the 1970 baseline of zero, condos became a more prevalent form of housing ownership, and one that increasingly is associated with gentrification. Table 2.1 presents results from OLS regressions of a set of neighborhood-level outcomes on percent condo. These outcomes include the log of mean household income in a tract, the population share of a tract with at least a bachelor's degree, and the population share of a tract that is black. Each specification includes a full set of neighborhood-specific housing controls and commuting controls, as well as tract fixed effects. In each case the result is consistent with standard gentrification trends, namely that income and share educated are increasing with percent condo and that share black is decreasing. The standard deviation in percent condo is approximately 10 percentage points, which suggests that a standard deviation increase in percent condo is associated with a 3.85% increase in income and a 2.38 percentage point increase in the share with at least a bachelor's degree. Conversely, percent condo is negatively associated with the share black, with a standard deviation increase in percent

condo correlated with a -0.3-percentage point change in the population share that is black. These estimates are not causal, however. To assess causality of condo conversion on these outcomes we explore a series of condo-related ordinances passed heterogeneously across US cities.

2.3.1 Identification Strategy: Condominium Ordinances

To isolate a causal effect of condo conversion on our outcomes of interest we need a source of exogenous variation that differentially affects the cost for neighborhoods to convert similar 1970 housing stock into condos. To do this we collected data on a series of municipal condo ordinances passed between 1970 and 2010. The broad intent in passing these ordinances was to protect tenants from displacement, in the process increasing the cost for developers of converting existing housing stock into condo units. Conversions typically occurred to rental units and were frequently met in the 1970s and 1980s with public opposition. Policymakers were primarily concerned with reductions to the rental housing stock and the displacement of current tenants. To address these issues, the regulations frequently include rights of first refusal for existing tenants to purchase the unit, various periods of tenant notice of the upcoming conversion, and relocation assistance (Fine (1980)). We construct a new data set of both municipal and state condo ordinance strictness by reviewing the individual state and city ordinances and assigning them a ranking. These ordinances all apply to the conversion of existing housing stock to condos, not to the construction of new condominium buildings. A city is then assigned a zero or one based on the strictness, and therefore costliness, of the policies detailed in the ordinance. The first municipal ordinances followed state and federal housing policy enacted in the early 1970s, accelerating in earnest following 1980. By 1973 only one city, Newark, NJ, had successfully passed a condo ordinance, followed in 1975 by Indianapolis, IN and Washington, DC in 1976. By 2009, 34 cities had passed some sort of municipal ordinance, 28 of which we categorized as imposing notable marginal cost on condo conversion.

59

To identify these condo ordinances, we reviewed the municipal code of each of the 100 largest cities, reviewing each occurrence of the word "condominium" in the municipal code and identifying any local laws restricting condominium construction or conversion. We also searched the Proquest newspaper archive for the terms "condominium" and "ordinance" together with each city name to identify any earlier ordinances that could have been enacted and subsequently repealed. We developed a three-point ranking for the local ordinances. Ordinances that recognize or facilitate condominium conversions but do not provide any protections for tenants or impose costs on developers, we ranked as a "0". Ordinances that set out time frames and requirements for tenant notification of condominium conversions or offer tenants the right of first refusal we ranked as a "1". Ordinances that went further to require tenant relocation assistance or tenant relocation payments we ranked as a "2". Ordinances that went still further to impose a cap on the number of permissible annual condominium conversions, established a minimum city-wide rental vacancy rate before conversions were permitted, granted some categories of tenants' lifetime leases, required the replacement of low-income rentals elsewhere, or required tenant approval for condominium conversion, we ranked "3". For our analyses, we categorize any municipality with an ordinance score of one or higher as having passed an ordinance. Any city with a "0", or no local ordinance, is categorized as not having passed an ordinance.

2.4 Empirical Strategy and Estimation

We start with an overall "structural" equation as estimated using OLS in Table 1, with outcomes expressed as a function of percent condo, a time-invariant neighborhood-specific effect, and neighborhood controls:

$$y_{ijt} = \alpha_i + \beta c_{ijt} + X_{ijt}\gamma + \varepsilon_{ijt} \tag{2.1}$$

Where i denotes a tract, j denotes an MSA, and t denotes year. Our endogeneity concern is that: $E(c_{ijt}|\varepsilon_{ijt}) = 0$. This concern arises primarily from simultaneity concerns, specifically in case there is an extent to which the outcomes themselves influence percent condo. In the case of income this is a plausible concern as a higher mean income in a neighborhood may lead developers to convert existing rental stock into condos. To overcome this limitation, we exploit the differential passage of municipal condo ordinances. Many neighborhoods that are comparable along housing stock characteristics in the period prior to mass condoization (i.e. prior to 1980) face different regulatory environments as only some of these areas ultimately pass ordinances.

We focus on municipal ordinances rather than those enacted by states. Municipal ordinances are most often passed in the "central city," or city containing the largest employment center, of a metropolitan statistical area (MSA). This offers us two margins of comparison. First, within an MSA that passes an ordinance there are neighborhoods within the central city that are subject to the ordinance while all other neighborhoods, including those adjacent to the municipal boundary, are not limited by the regulations. We refer to this margin of comparison as the central city boundary comparison, illustrated in the map in Figure 2.4. This figure again includes maps of Washington, DC and Baltimore, MD, with the central cities distinguished by darker coloring (blue for Washington, DC, red for Baltimore, MD). The intuition for this empirical approach is that honing in on the municipal boundary mitigates concerns about differing trends between central city and suburban tracts that may arise in the broader sample. As a real world example of this comparison, this approach isolates how a neighborhood in Washington, DC changes relative to a neighborhood in Arlington, VA (separated only by the Potomac River).

As a second margin of comparison, housing characteristics are often comparable within the downtown, or "city center", areas of a city. The city center in the set of neighborhoods containing the employment epicenter of a metropolitan area. Two comparable city center neighborhoods, such

61

as DuPont Circle in Washington, DC and Fells Point in Baltimore, MD, face different regulatory pressure for conversion of their housing stock to condos. Figure 2.5 illustrates this city center empirical approach. As evidenced by the percent condo maps in 2010 shown in Figure 2.1, condo percentages are high in city center neighborhoods. Interestingly, they are also high in central city boundary neighborhoods, which similarly provides additional motivation for our boundary approach. Note that not all boundary or city center tracts contain 1970 housing stock that is necessarily conducive to condo conversion, so for most analyses we limit the sample above a threshold of percent 5+ units. Finally, these maps demonstrate that while both approaches attempt to standardize the housing stock between treatment and control groups, their respective samples – especially in larger cities — have minimal overlap.

2.4.1 Estimation and First-Stage

To implement these two approaches, we implement a first-stage for percent condo that takes the form of a differences-in-differences-in-differences (DDD), augmented by a geographic and base period housing characteristic restriction. The full specification has three margins of comparison. For simplicity, we start by describing the first two, which are: (i) central city versus suburban, and (ii) pre-and post-ordinance passage. Utilizing only these two dimensions is possible in MSAs where a central city passes an ordinance. Given that ordinances apply only to city neighborhoods, this approach would be a standard differences-in-differences, with central city neighborhoods serving as the treatment group and suburban neighborhoods serving as the control group. Denote a tract's distance from the nearest municipal boundary as b_i , distance to city center d_i , and a tract's base period percent 5+ units as h_i . Applying restrictions of the form $b_i < \bar{b}$ and $h_i > \bar{h}$ limit the sample to tracts near the boundary above a certain percentage of 5+ units. Even with these sample restrictions, an additional concern that violates necessary assumptions for DD is that the trends

differ for central city relative to suburban neighborhoods. To account for this, we make use of the set of MSAs that do not contain a constituent city that ever passes a condo ordinance, but for whom we can still compare central city to suburban neighborhoods. Below is the full specification. For all $b_i < \bar{b}$ and $h_i > \bar{h}$, we estimate a $\hat{\delta}(\bar{b}, \bar{h})$ from:

$$c_{ijt} = \mu_i + \tau P_t + \pi_1 P_t M_j + \pi_2 P_t C_i + \delta C_i M_j P_t + X_{ijt}\theta + \nu_{ijt} \tag{2.2}$$

Where C_i is one if tract i is in the central city (i.e. zero if suburban), P_t indicates the post-period (2010), and M_j indicates that MSA j has a central city that at some point in time passed a condo ordinance. μ_i is a tract fixed effect. It follows that an ordinance impacts condo conversion only if $P_t M_j C_i = 1$.

Figures 2.6 and 2.7 present estimates of $\hat{\delta}(\bar{b}, \bar{h})$, varying $\bar{b}$ and $\bar{h}$, respectively. The intuition is neighborhoods facing a condo ordinance bear a higher conversion cost, which likely limits condo conversion relative to what it would have been absent the regulation. The empirical results support this hypothesis in Figure 2.6, where for each estimate within approximately two miles of the central city boundary we find a negative and significant effect of the condo ordinance on percent condo. This effect is stable, ranging from -10 to -15-percentage points. Figure 2.3 contains similar results. This figure changes the restriction for minimum percent 5+ units in the housing stock, fixing the boundary distance to within two miles. Once the percent 5+ threshold exceeds approximately 40% the estimate of the condo ordinance effect on percent condo is significantly negative, initially around -5, dropping to -10-percentage points as the threshold rises. Interestingly, there is also a small but positive effect when the percent 5+ threshold is very low (<20%). This may suggest that neighborhoods with less dense housing stock may receive higher condo housing stock in the form of new development with developers opting to build new condos (less costly in less congested neighborhoods), rather than convert existing high density stock under the ordinance.

The city center approach is similar in set-up, and nests within Equation (2), albeit with a limitation of the form $d_i < \bar{d}$ and $h_i > \bar{h}$, meaning all tracts in the sample are within a certain distance of the city center. However, because $C_i = 1$ for all city center neighborhoods, the city center approach is simply a DD estimate (i.e. there is no need to account for differences in suburban versus central city trends).

Table 2.2 includes estimates from this approach, including the sample of all city center neighborhoods, partitioned by whether their housing stock has as least 60% 5+ units (1), or not (2). For city center neighborhoods with housing stock conducive to condo conversion the ordinances limit condo conversion by five percentage points, whereas in neighborhoods with less dense housing stock the effect of the ordinances is that condo rates are two percentage points higher. Both estimates are significant at the 5 percent level. This pattern mimics that found in Figure 8 in the boundary analysis where the percent 5+ restriction was varied. In both cases, the results suggest that in neighborhoods with a high percentage of housing structures with multiple units, condominium ordinances had a significant, negative effect on percent condo. Moreover, this effect is heterogeneous and depends on a neighborhood's housing stock propensity to convert to condos.

2.5   Effect of Condos on Outcomes

The first stage results are consistent with the type of variation that is anticipated ex ante regarding how conversion regulations impact percent condo. The increased costs that accrue to areas limit the condo development in areas subject to regulation relative to what they would be absent the ordinances. We now utilize this variation to isolate a causal effect of condos on neighborhood gentrification outcomes. Table 2.3 includes two-stage least square estimates of our structural equation (1), using a first-stage for percent condo as in (2) and the central city boundary approach. All estimates are limited to a sample within two miles of the municipal boundary and with at least half

64

of the housing stock consisting of structures with five or more units. The table also includes OLS estimates that correspond to the central city boundary samples used in the estimation strategy, hence they differ in magnitude from those presented in Table 1.

For income, there is evidence of strong causal relationship between percent condo and income at the neighborhood level. A 10-percentage point increase in percent condo improves neighborhood income by approximately 7.8% based on the 2SLS estimates. The effect is very similar in magnitude to the OLS estimate for the same sample of 6.1%. The same holds for the percentage of neighborhood residents who hold a bachelor's degree or higher. We find that a 10-percentage point increase in percent condo improves the education levels in the neighborhood by 5.7 percentage points, again slightly higher than the analogous OLS estimate of 3.7 percentage points. These results are consistent in sign and significance to the original OLS results on the full sample for income and education, with slightly higher magnitudes.

In the boundary sample, there is no discernible correlation between percent condo and share black based on OLS estimates. 2SLS estimates, however, demonstrate a positive effect of percent condo on share black, with a 10-percentage point increase in percent condo leading to a 2.7 percentage point increase in the share of black neighborhood residents.

The causal evidence using the boundary approach is largely corroborated within the city center sample. Table 2.4 includes these estimates based on city center neighborhoods. Like what is observed at the boundary, OLS and 2SLS estimates are similar in magnitude and significance. A 10-percentage point increase in percent condo results in an 9.5% increase in median income (8.1 pp based on OLS), and a 5.7 percentage point increase in share with at least a bachelor's degree (4.1 pp based on OLS).

However, the results differ for the neighborhood share of residents who are black. Interestingly, at the city boundary the 2SLS estimates indicate a positive and significant effect of percent on

condo on share black that is large: a 10-percentage point increase in percent condo leads to a 2.7 percentage point increase in black share of the population. In contrast, in the city center approach there is no discernible effect, with a negative coefficient in the OLS estimate failing to sustain through 2SLS estimates exploiting the conversion ordinance variables. These results illustrate how these boundary and city center results, while similar for income and education, do not necessarily converge for other outcomes. With respect to condominium development as a force of gentrification, this evidence does not suggest widespread displacement in the form of racial composition within our sample of dense urban neighborhoods.

2.5.1 Ownership and Condominiums

A condominium is the most efficient method of residential property ownership within a high-density neighborhood. Next we explore the evidence on how ownership patterns are determined by percent condo. If true that the main mechanism underlying the condominium's influence on gentrification is access to ownership, there should be a clear relationship between ownership and percent condo in the data. Percent condo and percent ownership can differ for a variety of reasons. A given condo unit, despite its ownership structure, may be occupied by the owner or rented out. One gentrification mechanism related to condominiums is that higher income individuals wish to own rather than rent, so that even with a desire to live in a city neighborhood a preference for ownership among high earners may push them to suburbanize if the only city ownership option is to the purchase of an entire building. The condo potentially afforded these high earners an option within the city. If this is the case, we should observe a high correlation between condominiums and ownership.

Alternatively, it is possible that early condominium purchases were done as investments, with owners refurbishing apartments and subsequently leasing them. In this case a condo can be thought of as a heterogeneous product. Relative to rentals available under homogenous, macro-managed

and operated buildings, condos possibly represent a distinct and possibly higher quality option, thus attracting higher income renters. With this story in mind, ownership rates and condo rates do not necessarily move in tandem. Finally, in a hybrid of the two previous stories, early condo adopters may have occupied the units they owned and then eventually rented out their units, in which case the owner-occupancy rate and percent condo relationship varies over time. This relationship is difficult to test given our limited time range of condo data.

We present estimates that evaluate this relationship, but with some significant caveats. First, in our data it is not possible to cross-tabulate, at the unit level, if a residence is owner-occupied and if it is a condominium; we can only observe the overall levels of percent condo, and separately percent of owner-occupancy in the neighborhood. The most obvious issue with this set-up is that neighborhoods with a large stock of detached, single family housing may have high rates of ownership and low rates of condos. Our sample restrictions based on percent 5+ units somewhat limit this issue, but it is still a possibility that after this restriction is imposed the remaining housing stock could vary widely in its composition of other types of structures. Other issues aligning ownership and condo rates potentially arise as well. Vacant units can be categorized as condominiums, but not as owner-occupied. One possibility is that a neighborhood with a high vacancy rate, even if the housing stock is dense, may have a high percent condo but low ownership-occupancy.

Tables 2.5 presents OLS and 2SLS estimates of an effect of percent condo on ownership rates, using both the boundary and city center approaches. For both empirical approaches OLS estimates suggest a positive correlation between percent condo and owner-occupancy rates. However, 2SLS estimates under both the boundary and city center approaches are insignificant. Given some of the issues with jointly observing the condo and ownership rates, in Table 6 we pose a different question. Table 2.6 uses the full sample of years, from 1970 to 2010, for which we observe ownership rates but not condo rates, and presents reduced form estimates of the boundary DDD estimate and city

center DD estimate of the condo ordinances on ownership rates (rather than condo rates). We find evidence that in the boundary tracts the ordinances negatively impact ownership rates. This is consistent in sign with what we documented for the ordinance effect at the boundary for condo rates, although smaller in magnitude. The effect at the city center is positive, but insignificant.

2.6   Discussion and Future Work

The advent of the condominium allowed households an option to own a "slice" of a larger building, offering a new legal form of ownership in cities. We document that individuals who live in condos are more likely to be higher income and have a college degree. Moreover, neighborhoods with a high percentage of condominiums also tend to have higher mean household income and higher shares of educated residents. The causal direction of this relationship is unclear; while the ownership option afforded by the condominium legal form may attract gentrifiers, it is also true that external demand shocks that attract high earning households to cities may incentivize developers to build more condos. To determine if condo development contributed to the gentrification of cities from 1970 to 2010 we exploit condo ordinances passed over the period in question at the municipal level, compiling a new data set of these ordinances for this exercise. In certain cities, these regulations were strict enough to significantly limit condo development. We introduce two empirical approaches to estimate this effect: first, focusing at the central city boundary comparing suburban to urban neighborhoods; and second, honing in on dense, city center neighborhoods. In both cases, we find that ordinances significantly limit percent condo in neighborhoods relative to what they would have been absent the regulation.

We use this variation to re-evaluate the causal channel of condominiums affecting gentrification outcomes. Using an instrumental variable strategy based on the ordinances, we find positive and significant effects of condominiums on neighborhood income and education levels, but evidence

on other outcomes, such as the share of black residents and ownership in a neighborhood, is less clear. Still, the evidence of the condominium effect on income and education at the neighborhood level is very strong, suggesting that condominium ordinances, used in tandem with neighborhood characteristics as in our empirical approach, offers a novel, neighborhood level supply-shifter that predicts inflows of high-income residents. This may prove fruitful in future work exploring the consequences of gentrification on other outcomes of interest, including crime and displacement, education, and local economic activity.

Variation in Condominium Rates (2010)

Baltimore, MD

Washington, DC

Text

Central City

Condo %

0.00 - 0.11

0.11 - 0.31

0.31 - 0.62

0.62 - 1.00
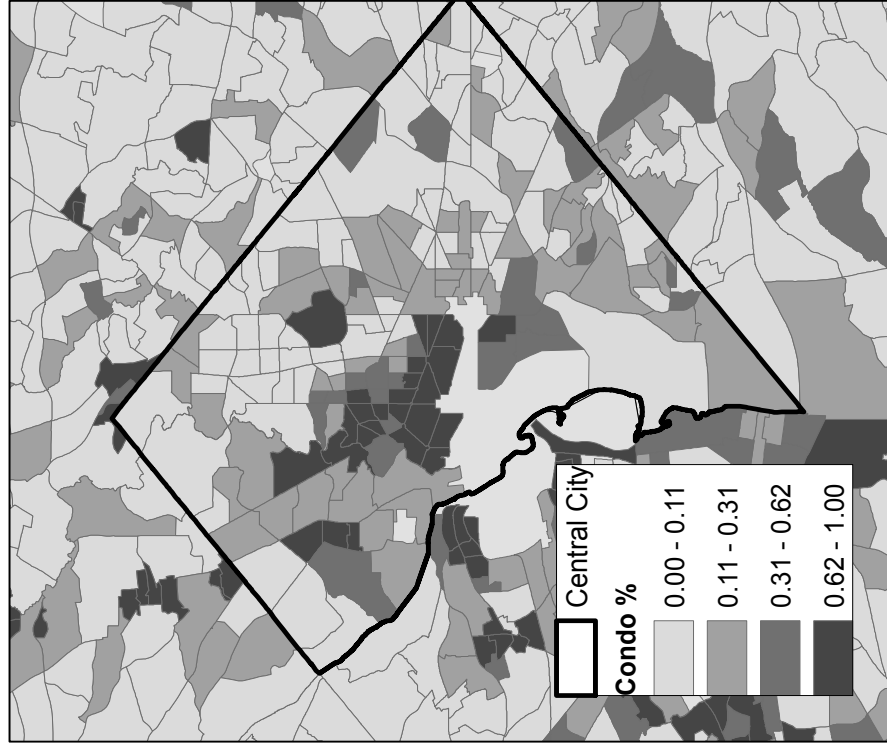
Source: American Community Survey 2010, Custom Table on Condominium Housing; Census Tiger Shapefiles 2010.

Figure 2.1: Condo Rates in 2010

Variation in % 5+ Unit Structures (1970)

Washington, DC
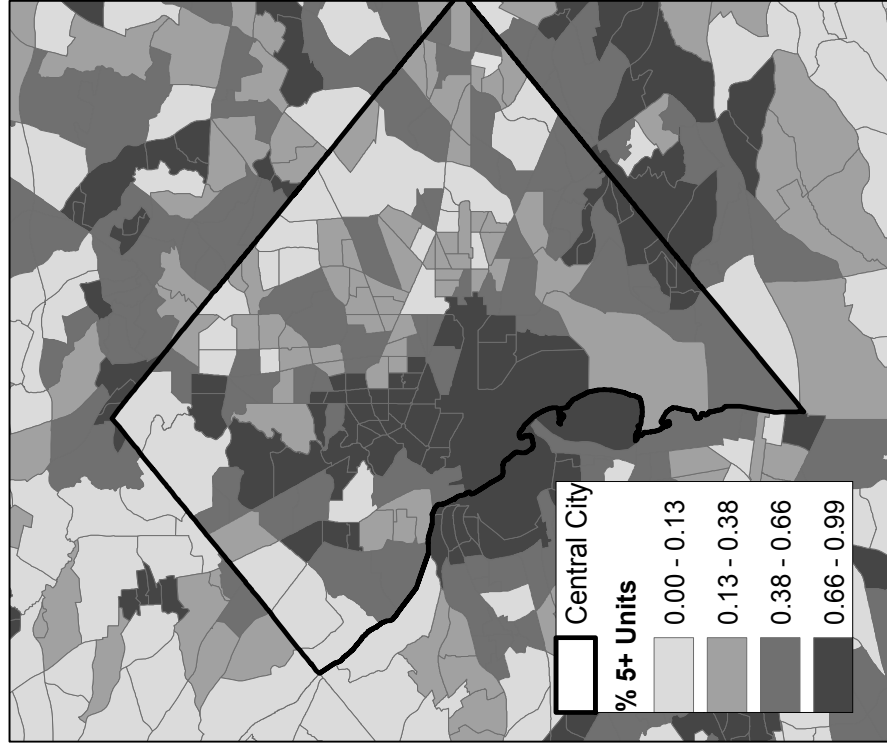
Baltimore, MD

Central City

% 5+ Units

0.00 - 0.13

0.13 - 0.38

0.38 - 0.66

0.66 - 0.99

Source: Neighborhood Change Database 2010; Census Tiger Shapefiles 2010.

Figure 2.2: Variation in Percent 5+ Units in 1970
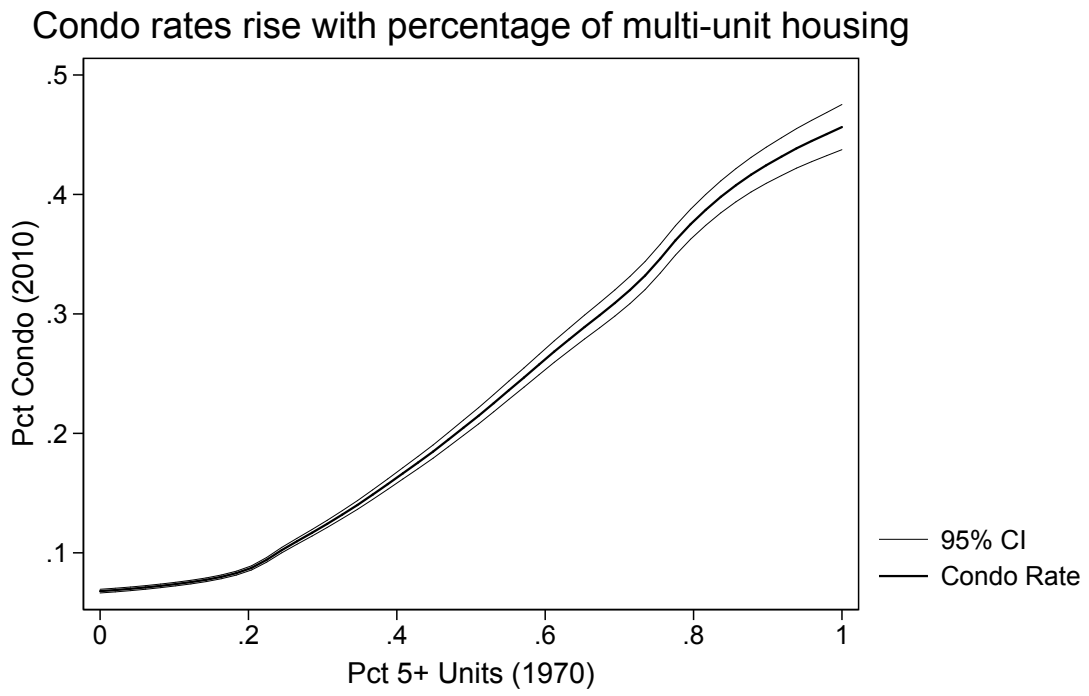
Figure 2.3: Condo Propensity: Multi-unit Structures

Table 2.1: Condo Rates Correlate with Neighborhood Outcomes

	(1)	(2)	(3)	(4)
	Log Mean Income	Share Bachelor Degree +	Share Black	Ownership
Condo Pct	0.415***	0.256***	-0.0605***	0.193***
	[0.0163]	[0.00539]	[0.00528]	[0.00614]
Observations	98577	99518	99533	99533
R^2	0.143	0.646	0.191	0.450
Housing Controls	Yes	Yes	Yes	Yes
Commuting Controls	Yes	Yes	Yes	Yes
Tract Fixed Effects	Yes	Yes	Yes	Yes

Standard errors in brackets

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Boundary of Central City Empirical Approach

Washington, DC (ordinance)

Baltimore, MD (no ordinance)

Legend:
- Boundary Tracts
- Central City (C=1)
- Suburbs (C=0)
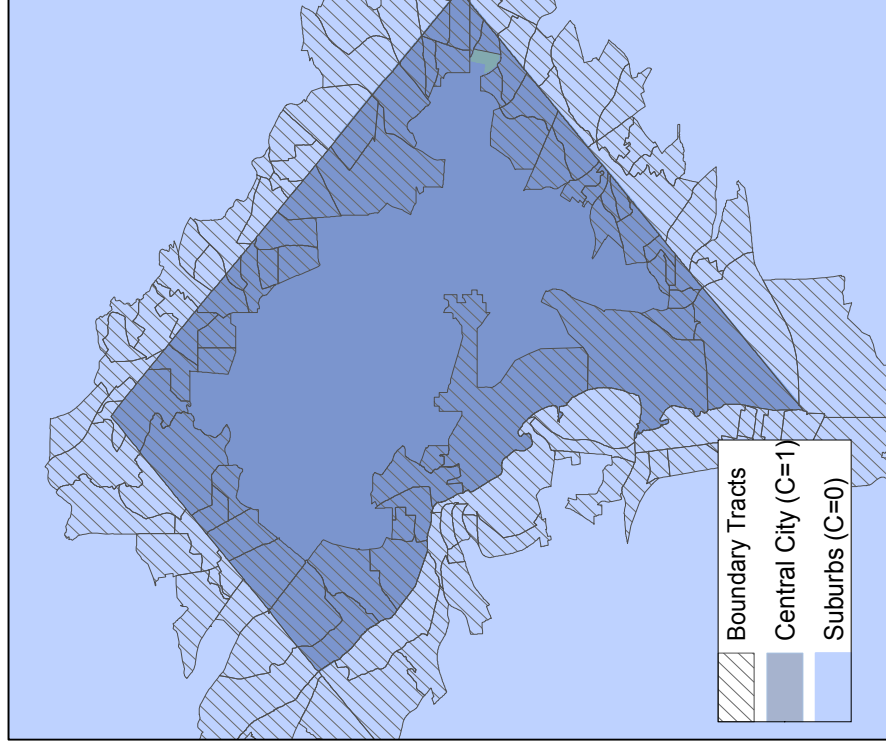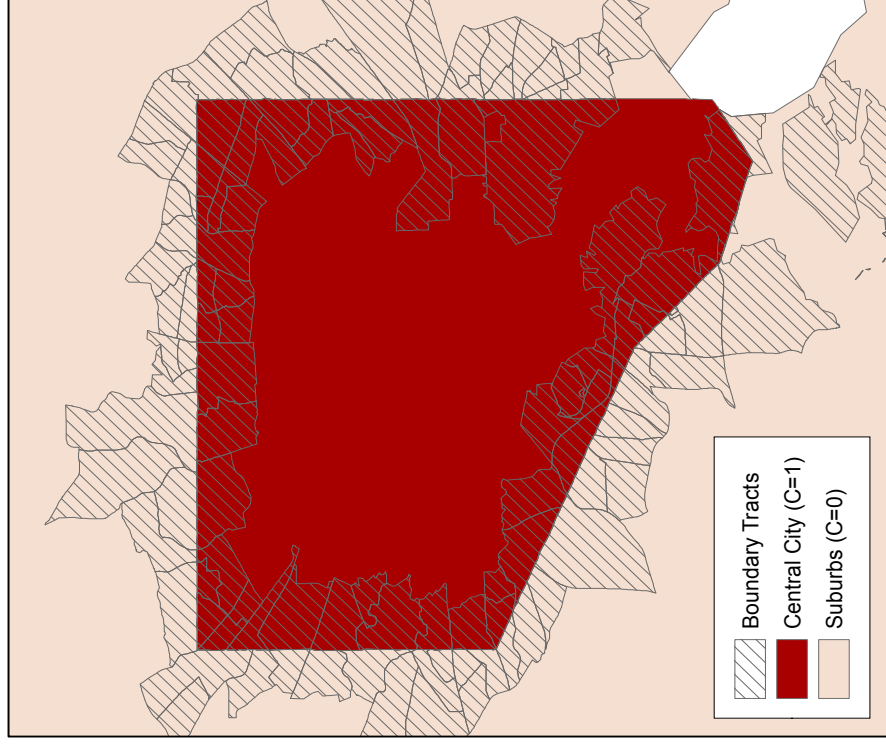
"Boundary" tracts lie within 1/2 miles of central city border.
Source: Census Tiger Shapefiles 2010

Figure 2.4: Central City Boundary Empirical Approach

City Center Empirical Approach
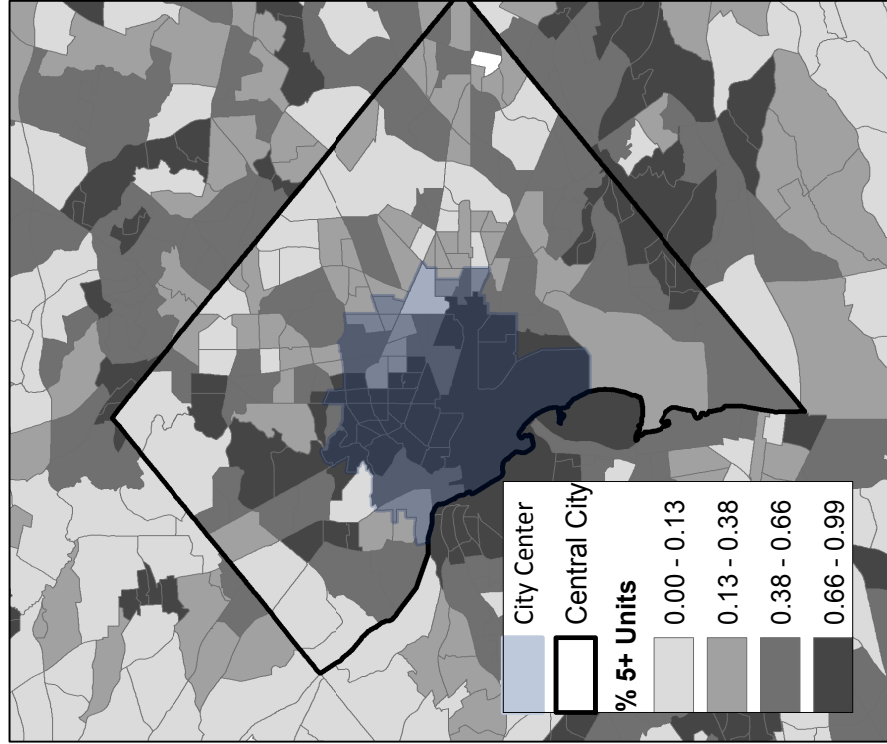
Washington, DC (ordinance)   **Baltimore, MD (no-ordinance)**
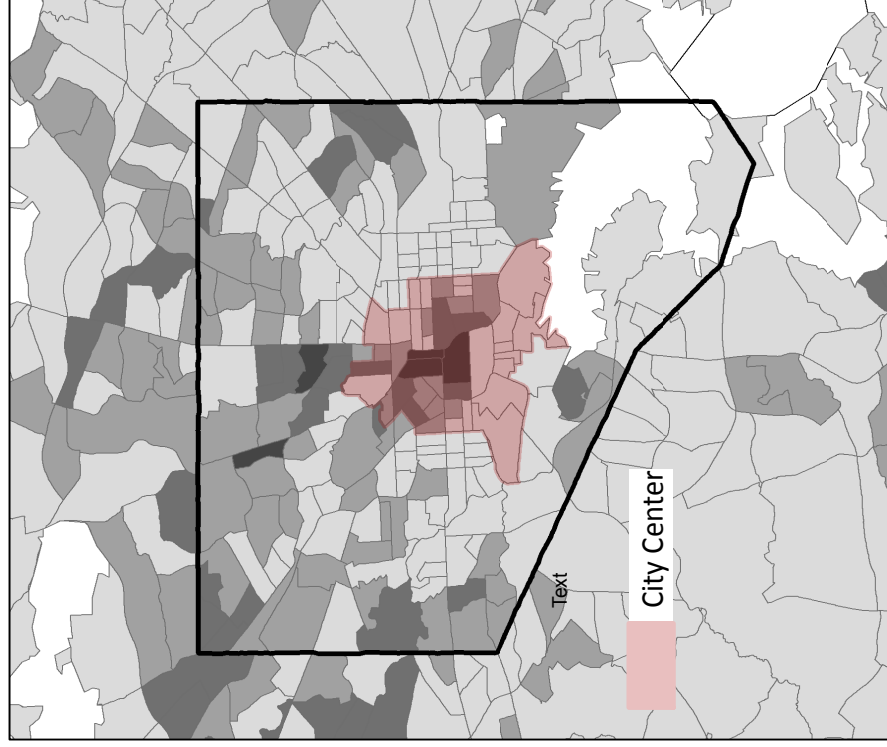
Comparison is blue region (subject to ordinance) versus red region (no ordinance).
Source: Neighborhood Change Database 2010; Census Tiger Shapefiles 2010.

Figure 2.5: City Center Empirical Approach

Figure 2.6: Heterogeneous Effects by Distance to Boundary
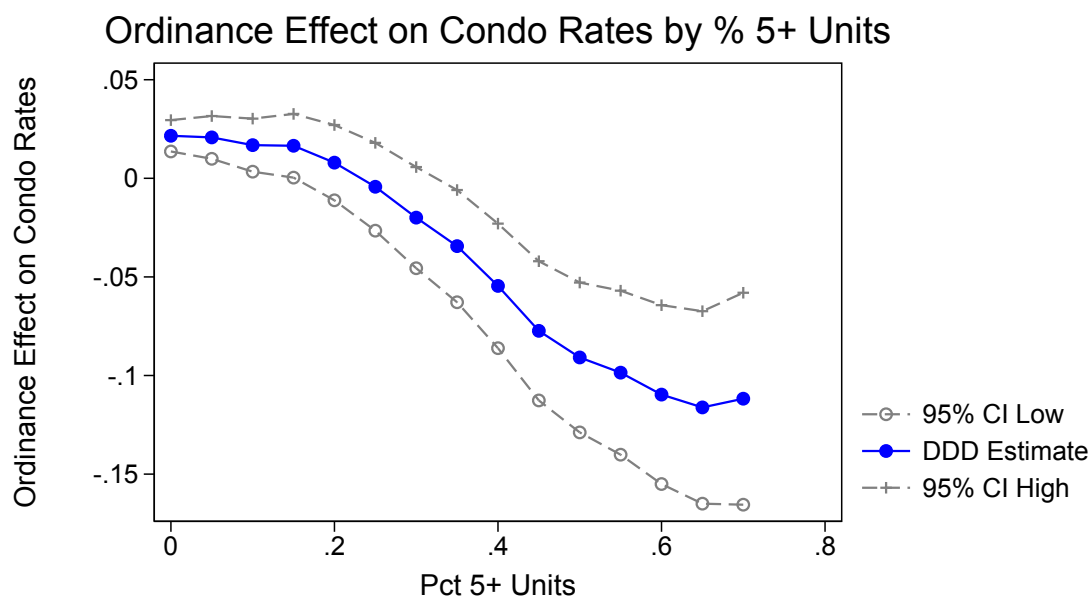
Figure 2.7: Heterogeneous Effects by Percent 5+ Units

Table 2.2: Ordinance Effect on Condo Rates in City Center

	(1)	(2)
	>60% 5+ Units	< 60% 5+ Units
Post-Condo Ord × MSA with Condo Ordinance	-0.0541*	0.0280*
	[0.0267]	[0.0109]
Observations	1886	3395
R^2	0.469	0.236

Standard errors in brackets

Limited to tracts in the city center. Uses zero assumption in pre-period.

Includes tract area fixed effects.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 2.3: Two Stage Least Squares Estimates of Condo Effects at City Boundary

	(1)	(2)	(3)	(4)	(5)	(6)
	Income (OLS)	Income (IV)	BD+ (OLS)	BD+ (IV)	Share Black (OLS)	Share Black (IV)
Condo Pct	0.613***	0.779***	0.371***	0.570***	-0.00247	0.274***
	[0.0255]	[0.0591]	[0.00715]	[0.0178]	[0.00917]	[0.0231]
Observations	6711	6469	6753	6534	6761	6538

Standard errors in brackets

Limited to tracts within two miles of central city boundary and with at least 50% 5+ unit structures.

Uses zero assumption in pre-period.

Includes fulll controls and tract fixed effects.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 2.4: Two Stage Least Squares Estimates of Condo Effects in City Center

	(1) Income (OLS)	(2) Income (IV)	(3) BD+ (OLS)	(4) BD+ (IV)	(5) Share Black (OLS)	(6) Share Black (IV)
Condo Pct	0.805***	0.948***	0.413***	0.572***	-0.0437***	0.0366
	[0.0405]	[0.0960]	[0.0118]	[0.0287]	[0.0117]	[0.0280]
Observations	1982	1982	1982	1982	1983	1983

Standard errors in brackets

Limited to tracts within city center and with at least 50% 5+ unit structures.

Uses zero assumption in pre-period.

Includes tract fixed effects.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 2.5: Two Stage Least Squares Estimates of Condo Effects on Ownership

	(1) Boundary (OLS)	(2) Boundary (IV)	(3) City Center (OLS)	(4) City Center (IV)
Condo Pct	0.0635***	-0.0551	0.0663**	0.0915
	[0.0125]	[0.0328]	[0.0226]	[0.0533]
Observations	6761	6761	1983	1983

Standard errors in brackets

Boundary limited to tracts within two miles of central city boundary and with at least 50% 5+ unit structures.

City Center limited to tracts within city center and with at least 50% 5+ unit structures.

Uses zero assumption in pre-period.

Includes tract fixed effects.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 2.6: Reduced Form Estimates of Condo Ordinance Effects on Ownership

	(1)	(2)
	Own (Boundary)	Own (City Center)
Post-Condo Ord × CC Tract × MSA with Condo Ord	-0.0628***	0.0887
	[0.0102]	[0.0681]
Observations	22202	3433

Standard errors in brackets

Boundary limited to tracts within two miles of central city boundary and with at least 50% 5+ unit structures.

City Center limited to tracts within city center and with at least 50% 5+ unit structures.

Uses zero assumption in pre-period.

Includes tract fixed effects.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Chapter 3

Party On: The Labor Market Returns to Socializing

3.1 Introduction

This paper introduces a framework for understanding the propensity for social individuals to experience high returns in the labor market. The channels through which social skills and social networks affect earnings are complementary to those pertinent to the seminal question of whether or not there are returns to schooling and cognitive ability (Card (1999), Card (2001)). The role that socializing plays in determining labor market outcomes is a related mechanism in that often decisions to socialize are made in tandem with studying decisions. High school and college are formative years for cognitive education, yet also periods of intense social interaction and friendship formation. We explore whether or not manifestations of sociability, such as having friends, improve labor market outcomes. We estimate a labor market return to having friends that is causal and suggests making five to six friends has an impact on wages of approximately 10%, comparable to a broad set of estimates of the return to an additional year of schooling.

After establishing this effect, we explore the mechanisms through which socializing influences labor market outcomes. Social skills can take a variety of forms. For example, how well an individual can communicate and present ideas, how persuasive they are, their perceptiveness of social cues or their ability to negotiate, instruct, or coordinate activity among different people (Kerr and Tindale (2004), Woolley et al. (2010), Bandiera, Barankay, and Rasul (2010), Mas and Moretti (2009)). These abilities stand in contrast to technical skills acquired in the classroom or through on-the-job training related to cognitive abilities. These include types of reasoning, memorization, and written and oral comprehension.[1] These same social skills that contribute productively in the workplace also foster social interactions elsewhere and contribute to the formation of social networks. The size and connectedness of an individual's social network can help immensely in finding a job, receiving promotions, bonuses, and preventing job loss (Calvo-Armengol and Jackson (2004),Montgomery

---

[1]A useful categorization of these types of abilities is O*NET https://www.onetonline.org/.

(1991), Elliott (2001), Granovetter (1973), Granovetter (1985), Corcoran, Datcher, and Duncan (1980)). We emphasize the importance of social skill development and the formation of one's social network, in particular as a trade-off with certain forms of cognitive development. Much of the existing literature on this topic groups social skills into a group of non-cognitive or "soft" skills largely derived from work by James Heckman (Heckman (2006), Heckman, Stixrud, and Urzua (2006), Heckman, Pinto, and Savelyev (2013)) that highlights that both cognitive and non-cognitive skills play a crucial role in determining labor market outcomes and social behaviors. These papers draw attention to the formation of both skill types at an early age, with a focus on parental involvement, child care, and preschool programs. We use this literature as a point of departure and focus on the decision-making of individuals when, having an endowment of social and cognitive ability from these earlier phases of life, they must make decisions that require a trade-off between social and cognitive investments.

We construct a model to isolate the mechanisms relevant when individuals in either high school, college, or in the workplace face these decisions. For intuition, we model a decision faced by the typical high school student any given day of the week: whether to socialize with friends or spend time studying. We refer to this colloquially as the *party or study* decision. Students enjoy socializing, dislike studying, and split their time between the two activities. Both socializing and studying contribute to labor market productivity and the wages individuals earn. Studying improves education and socializing improves one's social network. As an important distinction, there is strategic interaction in socializing but not in studying. In other words, socializing when others socialize leads to more friends, but the impact of studying on your educational output is independent of the study behavior of others. It is this strategic complementarity that distinguishes our framework from others in the non-cognitive and cognitive skill investment literature.[2] Students in the model

---

[2]For examples of recent work on social networks and the labor market, see Deming (2015), Conley et al. (2015), Lindenlaub and Prummer (2014), Weinberger (2014), Conti et al. (2013).

make socializing and studying decisions across utility in two periods, facing a constraint on their time, in a non-cooperative game. Both studying and socializing decisions, along with endowments, determine a level of education and number of friends, both of which serve as type of capital and earn returns in the labor market.

Two questions follow from the model set up, which we investigate empirically. First, is there a trade-off between studying and socializing? We find that individuals with higher social endowments invest less in studying and more in socializing, along with related activities such as drinking alcohol. In turn, high social types make more friends while earning lower grade point averages and achieving fewer years of schooling. Conversely, those who choose to study earn higher grade point averages and achieve higher levels of schooling, but make fewer friends than those who socialize. Second, is there evidence of strategic decision-making in socializing decisions? We find evidence that socializing decisions, especially those related to the consumption of alcohol, are dependent on peer behavior.

One of the prominent empirical challenges in this environment is finding a measurement of social skill. We condense the social skills that individuals bring to the labor market into observable characteristics of their social network. For the primary portion of our empirical work, we use data from Waves 1 through 4 of the restricted-use National Longitudinal Study of Adolescent to Adult Health (henceforth "Add Health", http://www.cpc.unc.edu/projects/addhealth), which contains detailed data on the friendship networks of individuals in middle and high school and how those networks persist into their professional lives. Our primary measure of social skill in this setting is a person's *in-degree*. For a given student i, i's in-degree is the number of people who nominate i as one of their friends. This social network data is matched to a rich set of information on behaviors, institutional characteristics, personal characteristics, and a variety of outcomes later in life, including education, employment, health, and marriage. Within Add Health we can identify students in the same school and grade, which we refer to as a *cohort*. We use this cohort data to construct

homophily instruments for both education and network characteristics. Using these instruments we find positive, significant 2SLS estimates of the log wage returns to years of schooling and in-degree, as well as support of our hypotheses on how socializing and studying decisions influence education and network formation. We supplement the Add Health wage results with empirical evidence that socialization not only impacts wages, but also marriage and health outcomes. Additionally, we include evidence from other panel data sets, including the British Cohort Study, and National Longitudinal Study of Youth, that are broadly consistent with our findings from Add Health.

This work contributes to a puzzle in the literature summarized by Weiss (1995), who explains that years of schooling and work experience are highly correlated with labor market outcomes, but "estimates suggest that courses, test scores, and measurable learning in secondary school can explain ... less than one-quarter of the higher earnings of high school graduates ... what accounts for the other three quarters of earnings?" The prevailing answer to this question is that a degree (e.g. high shool diploma, bachelor's degree, etc.) signals to employers the presence of unobservable, desireable traints in potential employees. Recently Johnson (2017) commented on this signal/credential explanation, stating that economists, while cognizant of the fact that college substantially impacts future earnings, "treat higher education as a black box." The focal point in our model, the party or study decision, offers insight. As students trade off studying for socializing, this has short-term implications on observable education output such as test scores and grade point averages, but still contributes to the student's productive capital for the labor market. It merely operates through the social channel, rather than cognitive. The importance of socialization in schooling has further applications to issues in higher education. Many colleges and universities are facing criticism for investing in infrastructure for non-academic, recreational facilities believed to contribute to increasing tuition levels (Jacob, McCall, and Stange (2013)). Similarly firm expenditure on entertainment and travel in many cases is also a high proportion of costs (Cai, Fang, and Xu (2011)). These

investments are more reasonably justified in the presence of high returns to socialization. Moreover, the proliferation of massively open online courses (MOOCs), controversy over greek life and alcohol consumption, among other topics, all affect the extent and nature of social interactions on campuses. While we leave the full details of this analysis to future work, we do present suggestive evidence on returns to drinking alcohol, maintaining the assumption that drinking is largely a socializing decision, especially on college campuses. Throughout our analyses we find evidence of a significant positive but concave impact of drinking alcohol on wages. Additionally, using salary data from the College Scorecard (https://collegescorecard.ed.gov/data/) and academic and party school rankings from Niche.com (https://www.niche.com/), we find associations between high wages after graduation and high rankings in both party and academic categories.

The remainder of this paper is organized as follows. Section 3.2 contains details on the variables we use from the Add Health data and some insight into how we use the timing of different waves of surveys, as well as details on supplementary data sets. Section 3.3 contains motivating empirical evidence of a return to socializing on wages, and how students choose trade-off social and education outcomes. Section 3.4 presents our model of the party and study decision, including comparative statics and the resulting Bayesian Nash equilibrium. Section 3.5 details how our model motivates instrumental variable estimation and details how we construct homophily variables at the cohort level as instruments. 3.6 presents our main results. Section 3.7 includes supplementary empirical evidence and applications. Section 3.8 concludes.

3.2 Data

The largest data challenge we face is finding accurate and intuitive measures of socialization. Ultimately, we need three social measures: a social endowment, a social investment (or effort), and a social output. Labor economists have generally converged to adopt certain analogous measures, such

as years of schooling as educational output, as standard and acceptable measures in the study of schooling on wages. IQ and a range of test score data are frequently used as measures of educational endowment or intelligence "type". There is less of a consensus for socialization.

Some studies focus on self-reported measures about extroversion, many of which are collected at an early age. Extroversion, and conversely introversion or shyness, is a member of the "Big Five" psychological traits, with evidence this trait is both stable over time (Cobb-Clark and Schurer (2012)) and associated with degree attainment and labor market outcomes (Lenton (2014)). Extroversion and shyness measures are widely available in data sets available to labor economists. Therefore, given its availability we use shyness as our main measure for social endowment.

Reliable metrics for social investment prove the most elusive. A popular choice is participation in school clubs or sports teams.[3] Time-use surveys collect fairly detailed information on time spent studying, reading, participating in physical activity, and recreational activities often engaged in alone or over the Internet, such as playing video games. Surveys are less comprehensive in their documentation of the intensity with which individuals socialize. We take two approaches. First, we assume throughout that a minute spent socializing comes at the cost of a minute studying. As we detail below in our model, with finite time students must choose a time allocation between these two activities. As such, measures of time spent studying contained in surveys we discuss below represent an inverse measure of time spent socializing. In the data sets where we observe more explicit measures of social effort, evidence suggests our assumption of a trade-off between the two seems valid.

Second, we take the stance that the decision to consume alcohol during the high school and college developmental periods we cover is primarily a decision to socialize. Alcohol is a notorious "disinhibitor" or "social lubricant" that lowers the cost of interacting with others. Moreover, alcohol

---

[3]For example, Weinberger (2014) focuses on extracurricular participation including: sports, clubs, performing arts, student publication, in addition to leadership roles within these extracurricular areas.

consumption among adolescents is believed to be largely motivated by its capacity to facilitate interactions and boost the energy of social events (Feldman et al. (1999), Kuntsche et al. (2005)). Drinking also improves the imbiber's mood, as well as the consumer's ability to perceive the moods of others (Stevens, Gerlach, and Rist (2008)). Indeed, we find that much like extroverts and those who spend time hanging out with friends, those who drink have more friends and earn higher wages. It is worth emphasizing that this is true for those who drink *in moderation*, and that there is sharp concavity to both social and labor market returns to drinking, aligning with the range of consumption wherein alcoholism becomes a concern.

Lastly, we leverage the structure and detail of social network data to compute metrics of social output. Again our focus is on the in-degree. The best data for this purpose is the restricted-use version of the Add Health data, whose structure is outlined in Table 3.1. In addition to having all of the social, education, and labor market variables of interest to us, the restricted-use data allows for observation of the precise social networks as they exist in the schools. Beyond observing each student's in-degree and other network characteristics we also observe who those exact friends are. In addition, we can see course overlap, examine the cohort characteristics of a particular student, and construct homophily scores.[4] This is perhaps the greatest benefit of the Add Health data: it allows us to observe who connects to whom in a number of different ways, including through friendship nominations, course overlap, cohort sharing, and homophily. Our identification strategy, detailed below, depends on constructing valid instruments for social and educational outputs and investments using these multiple viewpoints of social closeness.

There are limitations to the restricted-use Add Health data, however. While the Add Health data contain excellent measures of social and educational output, and excellent detail about the interactions among respondents, there are two main drawbacks. First, the investment measures

---

[4]We use the term "cohort" to indicate a cohort of students in the same grade of the same school.

aside from those related to alcohol consumption are poor and often partitioned into broad categorical variables. Second, the labor market outcomes are collected at a fairly early stage in the labor market, around the age of 30. As a result, we additional include OLS results from two other panel data sources: the British Cohort Study (henceforth "BCS"), and the National Longitudinal Study of Youth 1979 (henceforth "NLSY79"). The BCS is a panel of around 19,000 students born in the United Kingdom in a single week, observed at a higher frequency than the Add Health sample, between birth and age 42. It contains measures an array of measures on both investments, endowments, and output, including some, such as number of friends, at a number of different ages. However, relative to Add Health it lacks schooling variables exactly comparable to the years of schooling measure common to the labor literature and the sample of respondents for whom full information exists across the survey dates is small. The NLSY79 contains a panel from cohorts born between 1960-65 first interviewed in 1979. Like Add Health and the BCS it contains endowment measures of IQ and shyness, investment measures such as time spent studying, time spent socializing, and drinking alcohol, and a set of outcomes from ages 29 - 49. However, it contains no "degree"-type social measure such as number of friends, and the social measures it does contain suffer from poor timing and are generally of lower quality. Importantly, neither alternative data set includes information on the exact linkage between respondents

3.3   Motivation: Friends and Wages

The economic study of socialization has benefited greatly from recent advances in network analysis and social networks. Innovations have come in the form of theory, which originated primarily in social science among sociologists. As the study of groups blossomed into the more nuanced study of networks, data collection progressed to include efforts to capture information on how group members

related to one another.[5]

Two major contributions arose from the economics literature relevant to our analysis. First, while data in a network setting more explicitly violates classical econometric assumptions of observational independence, the *structure* of networks has proven an important tool in identification (see Bramoullé, Djebbari, and Fortin (2009), as well as more broad discussions in Graham (2015), Blume et al. (2010), and Jackson (2005)). Second, social interactions are well modeled using tools from game theory. Recent work has demonstrated that various types of Nash equilibria in games in a social setting can be expressed as functions of network characteristics. These include Boncacich centrality (Calvó-Armengol, Patacchini, and Zenou (2009)) and the lowest eigenvalue (Bramoullé, Kranton, and D'amours (2014)). The dependence of behavior decisions on social network structure is an important result, since non-network based measures of social skill are often self-reported, retrospective, and contain minimal or irrelevant empirical content.

We assume something fundamental about social skill, namely that high social skill means a higher conversion rate of social interactions into friendships. As such, characteristics of a person's ego network, including their in-degree and out-degrees, among other centrality measures, are functions of their innate social skill and socialization investments. While much of the earlier literature explores whether or not there are peer effects in different behavior, we go one step further to explore the impact of social network structure on labor market outcomes. There is substantial empirical evidence that suggests positive and notable large associations between network characteristics and wages. See Figure 3.1, which contains plots of nonparametric estimates of the correlation between log earnings and a measure of education (years of schooling) and a measure of socialization (in-degree). As

---

[5]The Add Health data that we use is perhaps to most well-known example of longitudinal, survey data enhanced by the collection of social relationships among respondents. We describe in detail the friend nomination data in Add Health below. Other examples, employed by some of the other papers we cite, include the Wisconsin Longitudinal Study ("WLS", http://www.ssc.wisc.edu/wlsresearch/) and the Berea Panel Study ("BPS", http://economics.uwo.ca/chcp/the_berea_panel_study.html).

expected, in the first panel years of schooling is positively correlated with earnings for both men and women, a stylized fact well-documented in the labor literature. Of note here is the positive correlation between log earnings and in-degree. As we will illustrate in detail below, this relationship withstands significantly more scrutiny and is the focus of the remainder of the paper. As a preview, while holding the earnings-years of schooling relationship constant, we find across a range of data sets and specifications a positive, significant relationship between number of friends and earnings later in life.

To understand the mechanisms that underlie this relationship, see Table 3.3. In addition to the in-degree and years of schooling variables, this table also includes summary statistics on a set of social and education variables that plausibly influence the outcomes correlated with earnings. All variables include overall means, in addition to conditional means based on binary versions of social and education endowment variables shyness and IQ, respectively. Those who are shy, on average, earn less money, have fewer friends, are less likely to spend time with friends and drink alcohol, and spend more time studying.[6] Higher IQ is associated with higher earnings, more friends, more drinking in later periods, less time spent with friends, shyness, more years of schooling, higher GPA and, most interestingly, less study effort.

The first goal of this analysis is to estimate a return to the social output - in-degree - jointly with an estimate of the return to education output. Figure 3.1 shows nonparametric plots of log earnings on years of schooling in panel (a) and in-degree in panel (b). Both series include 95% confidence intervals, separate out male and female earnings, and include a sample of only white individuals. The pattern with respect to years of schooling is, as expected, increasing steadily. Interestingly, this near-linear, positive trend bears out in the log wage relationship with in-degree as well. Individuals who are nominated as friends by more people have higher earnings.

---

[6]Differences along other dimensions, including years of schooling, GPA, and IQ/Picture Vocabulary Test, are less pronounced.

There is an existing stigma that having a large number of friends, or being popular, has a tendency to be inversely related to some marks of intelligence or education-forward behavior. Experimental evidence, including Bursztyn and Jensen (2015), supports this idea in that students may suppress education investments, such as studying or test preparation, due to social peer pressure. While this response to prevailing norms may very well exist, at broader levels social and education markers are in fact fairly positively correlated. Figures 3.2 and 3.3 present correlations between years of schooling and in-degree, and IQ and shyness, respectively. For both men and women there is a strong positive correlation between years of schooling and in-degree. Additionally, despite priors to the contrary there is no evidence of an inverse relationship between shyness and IQ. In fact, the distribution for non-shy individuals is slightly further right than for shy individuals.

3.3.1 Preliminary Estimates of the Return to Socialization

The descriptive evidence indicates that the ability to form and sustain meaningful friendships has a notable, positive association with wages later in life. Next we present evidence that this relationship holds up to more careful empirical scrutiny. This association remains after controlling for a host of potentially confounding factors, including years of schooling, a host of background controls, including race, gender, nativity, and parental background, as well as fixed effects at the school level. Moreover, the relationship persists in a number of different data sources. See Table 3.4, which presents results from OLS regressions of wages during Wave 4 on years of schooling (Wave 4) and in-degree (Wave 1) using the restricted-use Add Health data. Note these are correlations only, and do not account for endogeneity concerns. The two specifications differ in the inclusion of endowments IQ and shyness. For both specifications, we find a positive and significant coefficients. For specification (2), the coefficients on in-degree is 0.0172, suggesting that an additional friend nomination in high school is correlated with a 1.72% increase in wages in adulthood. In the same

specification we estimate a coefficient on years of schooling that indicates an additional year of schooling is associated with a 11.8% increase in wages in adulthood. The ratio of the return to a year of schooling to a return to a friend is approximately seven; so the wage benefit from seven additional friends equates to the wage benefit of an extra year of schooling.[7]

The basic result holds in other data sets as well. In Table 3.5 we present comparable estimates from the British Cohort Study in which an additional friend at age 34 is associated with approximately 3% higher wages. Again, this holds while controlling for levels of education.[8] Additional results documented by Conti et al. (2013) provides additional, descriptive evidence of comparable magnitude and significance. For a return to sociability they find in the BPS that an additional friend nomination (in-degree) is associated with 2% higher wages.

3.3.2 The Production of Education and Sociability

The key component of our mechanism is that individuals make a choice between two activities. The first is studying, which improves educational outcomes, and therefore wages, while negatively impacting utility. The second is partying or socializing, which is enjoyable and thus positively impacts utility and – potentially – wages by improving social outcomes like friendships. Before delving fully into our model, in this section we present motivational evidence of a trade-off pattern.

Table 3.6 includes a set of regressions that offer insight into the trade-off between education and social skill by investment. The outcome variables in this table are in-degree, our main social outcome, along with years of schooling and GPA as education outcomes. The regressors include our set of social and education endowments and investments. Starting with the endowments, shyness is

---

[7]While our focus is on wages, other life outcomes are significantly, positively correlated with both years of school and in-degree. See 3.12 for examples of marriage and health outcomes.

[8]The British data uses a slightly different measure for schooling, the highest age the respondent left the education system, rather than standard years of schooling. Still, this variable is positive, significant and of the same order of magnitude as standard schooling return estimates.

significantly and negatively correlated with in-degree (see specifications (1) and (2)). Self-reported shy is associated with 0.6 fewer friend nominations in high school. Shyness has no discernible correlation with either education outcome. The education endowment, an Add Health administered picture vocabulary test, is positively correlated with both social and education outcomes. These results are largely consistent with the comparative static predictions of our model in Section 3.4.

One of our key results is contained in the results for the endowment variables. Individuals who exert study effort face 0.2 fewer friend nominations, but achieve 0.3 more years of schooling and earn GPA that is 0.25 points higher than those who report not exerting study effort. Conversely, those who report spending time with friends (which we redefine to be binary), face 0.5 more friend nominations while achieving 0.2 fewer years of schooling and GPAs that are approximately 0.06 points lower than those who do not report spend less time with friends. Moreover, when we include alcohol consumption, another form of socializing, we find similar results. Those who drink frequently do so to the benefit of their friendships and to the detriment of schooling outcomes. Individuals who drink alcohol frequently (more than once a month) earn 0.3 more friend nominations, while achieving 0.2 fewer years of schooling and 0.16 fewer GPA points relative to those who drink less than once a month.

Similar to the wage results, this finding is not limited to the Add Health data. Table 3.7 presents similar results from the NLSY79 data. The social investment variable, average days drinking alcohol in a month, is positively and significantly associated with being outgoing as an adult, while hours study, a measure of education investment, has no significant correlation with adult sociability. Years of schooling is significantly correlated with both variables, positively so with hours spent studying and negatively with days of drinking. These are consistent with the broader results in the Add Health data.[9]

---

[9]The outgoing as an adult variable is different than in-degree, but similar in spirit and sufficient for this exercise. Also, note the endowment correlations, namely that being an outgoing child is positively and significantly affects

Beyond the influence of endowed characteristics, individual decisions on how to invest time and effort between social and educational activity suggest a clear trade-off: social activities, including hanging out with friends and drinking alcohol, have a substantial positive association with number of friends, although those same decisions negatively impact schooling outcomes. The reverse is true of effort and time allocated toward studying and scholastic activity. Studying is strongly and positively correlated with education outcomes, but negatively influences a person's ability to develop and maintain friendships.

3.3.3 Who Invests?

To complete our investigation of this study-party mechanism we examine which types of individuals opt to allocate their time and effort between studying and partying. In Table 3.8 we regress a set of social and study investment variables on our set of endowments, along with a set of controls. A few of the results are worth mentioning. Study effort is positively and significantly correlated with shyness: shy individuals are more likely to report expending effort studying by three percentage points relative to more outgoing individuals. Those with higher IQ, however, are significantly less likely to study.[10]

Shyness is systematically negatively and significantly associated with social effort. Shy individuals spend less time with friends and less time drinking, both at high school and college ages. IQ results are more varied, which we will demonstrate below is consistent with the model. Higher IQ is associated positively with frequent drinking at college ages and negatively with time spent

---

being an outgoing adult but has not effect on schooling. Conversely, AFQT, an IQ measure, has no correlation with being outgoing but is positively associated with schooling outcomes.

[10]This result is consistent with the comparative statics of our model below. There are two reasons in our set-up that higher IQ individuals study less. One is that as if intelligence is an attractive characteristic for a friend, meaning that for each unit of time spent socializing they are able to convert interactions more easily into friendships. Second, in a variant of our set-up we allow the cost of studying to vary with intelligence. In this case, smarter individuals can achieve the same education output with less studying and more partying relative to less smart individuals. These are reasons we may expect to see high IQ people invest more in socializing and less in studying.

with friends, both correlations are statistically significant. High-school age drinking and IQ are insignificantly correlated.

3.4 Model

Next we introduce a model that helps formalize the decisions we explored in the previous sections. Consider a group of individuals indexed by i. Individuals are initially endowed with some level of social skill θ_i^s, as well as a given cognitive ability, for example intelligence or proclivity for analytical or cognitive tasks, θ_i^e. In the first period, best thought of as a schooling period such as high school or college, individual i receives utility from socializing, which can be thought of as a form of leisure:

$$U_{1i}\left(s_i\right) = \left(\beta\theta_i^s + \upsilon_i\right)s_i - \frac{1}{2}s_i^2 \tag{3.1}$$

where s_i is the time spent in socializing, θ_i^s represents the observed social skill endowment (e.g. shyness), and υ_i represents the unobserved preferences. We normalize the total amount of time to 1, so $0 \leq s_i \leq 1$.

Individuals accumulate two types of capital, social capital S_i (for instance the number of friends) and cognitive capital E_i (for example, years of schooling, GPA, or a test score). S_i and E_i depend on initial skills θ_i^s and θ_i^e, and on inputs into their production, specifically a time allocation. In addition to socializing, we assume students spend time studying, which we denote by e_i, so that $s_i + e_i = 1$. This constraint is the essence of what we refer to as the *party or study* decision. For example, a college student must choose how much time to allocate to social activities (e.g. attending a party) versus focusing on studying (e.g. spending an evening in the library). The final levels of S_i and E_i are given by production functions:

- *Education.* We assume that the education E_i follows a distribution F_e supported on $\left\{0, 1, \ldots, \bar{E}\right\}$

for some $\bar{E} < \infty$ with mean

$$\mathbb{E}\left[E_i | s, w\right] = \alpha \theta_i^e e_i \tag{3.2}$$

where $e_i = 1 - s_i$ is the time spent in studying and θ_i^e represents the observed cognitive skills (e.g. IQ). Since education is measured by years of schooling, we assume E_i has a support on non-negative integers with a maximum $\bar{E}$.

- *Social Capital.* The social capital S_i is measured by the in-degree of i, and is assumed to follow a distribution F_s supported on $\{0, 1, \ldots\}$ with mean

$$\mathbb{E}\left[S_i | s, w\right] = \sum_{j \neq i} p_{ij} s_i s_j \tag{3.3}$$

To justify the specification in (3.3), suppose each i in a school has the opportunity to meet any of their classmates, indexed by j. The likelihood that i meets j depends on the time that i and j spend socializing, s_i and s_j. Moreover, given that they meet, they convert the interaction into a friendship with probability:

$$p_{ij} = p\left(x_i, x_j, \theta_i^e, \theta_j^e, \theta_i^s, \theta_j^s\right) \tag{3.4}$$

which could depend on the social distance $|x_i - x_j|$ between i and j (homophily), θ_i^e, θ_j^e (quality of link) and θ_i^s, θ_j^s. These meeting and formation probabilities together determine the network formation process, resulting in an individual's social capital they bring to the labor market.

The key component of this structure is in the interdependence of socialilzing: students only form friendships, and thus a network, if they socialize along with others in their group of potential friends, i.e. those in the same school or cohort.

In the second period individuals take their E_i and S_i as skills on the labor market, where individual i's utility depends on her wage W_i (consumption)

$$U_{2i}\left(W_i\right) = \ln W_i = r_e E_i + r_s S_i + \gamma' x_i + \varepsilon_i \tag{3.5}$$

which is determined by education E_i, social capital S_i, other observed characteristics x_i, and un-observed heterogeneity ε_i. In the wage equation (3.5), r_e and r_s represent the returns to education and social capital, respectively. Let $w = \left(\theta^e, \theta^s, x\right)$ denote the observed variables.

3.4.1 Equilibrium

Suppose that in period 1 when i chooses s_i, v_i is realized, but E_i, S_i, and ε_i are not, so i makes the decision according to her expected utility. Assume $w = \left(\theta^e, \theta^s, x\right)$ and v are publicly observed. Moreover, assume that

$$\mathbb{E}\left[E_i \middle| s, w, v\right] = \mathbb{E}\left[E_i \middle| s, w\right]$$
$$\mathbb{E}\left[S_i \middle| s, w, v\right] = \mathbb{E}\left[S_i \middle| s, w\right]$$

and

$$\mathbb{E}\left[\varepsilon_i \middle| s, w, v\right] = \mathbb{E}\left[\varepsilon_i \middle| w, v\right]$$

In a Nash equilibrium, given s_{-i}, each i chooses s_i to maximize their expected utility

$$U_{1i}\left(s_i\right) + \mathbb{E}\left[U_{2i}\left(W_i\right) \middle| w, v\right]$$
$$= \left(\beta\theta_i^s + v_i\right) s_i - \frac{1}{2}s_i^2 + r_e \alpha\theta_i^e \left(1 - s_i\right) + r_s \sum_{j \neq i} p_{ij} s_i s_j + \gamma' x_i + \mathbb{E}\left[\varepsilon_i \middle| w, v\right]$$

Interior Solution

The optimal strategy profile $s^* = (s_1^*, \ldots, s_n^*)$ is obtained from the first order condition

$$s_i^* = r_s \sum_{j \neq i} p_{ij} s_j^* + \beta \theta_i^s - r_e \alpha \theta_i^e + v_i \tag{3.6}$$

Define the matrix

$$P = \begin{bmatrix} 0 & p_{12} & \cdots & p_{1n} \\ p_{21} & 0 & \cdots & p_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ p_{n1} & p_{2n} & \cdots & 0 \end{bmatrix}$$

We write (3.6) in matrix form

$$s^* = r_s P s^* + \beta \theta^s - r_e \alpha \theta^e + v \tag{3.7}$$

Assume $|r_s| < 1$. The equilibrium s^* is unique and is given by

$$s^* = (I - r_s P)^{-1} \left(\beta \theta^s - r_e \alpha \theta^e + v \right) \tag{3.8}$$

The equilibrium E_i^* and N_i^* have means given by $\alpha \theta_i^e e_i^*$ and $\sum_{j \neq i} p_{ij} s_i^* s_j^*$.[11]

---

[11]

Our setup does not guarantee an interior solution. With the presence of corner solutions, the optimal s_i^* is of the form

$$s_i^* = \begin{cases} 1 & \text{if } y_i^* \geq 1 \\ y_i^* & \text{if } 0 < y_i^* < 1 \\ 0 & \text{if } y_i^* \leq 0 \end{cases}$$

where $y_i^* = r_s \sum_{j \neq i} p_{ij} s_j^* + \beta \theta_i^s - r_e \alpha \theta_i^e + v_i$. We can write s_i^* more concisely as

$$s_i^* = 1 \{y_i^* > 0\} \min \{y_i^*, 1\}$$

This system could have multiple equilibria. Since s_i^* is weakly increasing in y_i^* and strictly increasing in y_i^* if $0 < y_i^* < 1$, any multiple equilibria, once exist, must take the form that there is a subset of the agents who either all choose $s_i^* = 1$ or all choose $s_i^* = 0$.

3.4.2 Comparative Statics

We examine how the equilibrium s^* changes as θ^e or θ^s changes. By the specification of p_{ij} in (3.4),

$$\frac{\partial p_{ij}}{\partial \theta_k^e} = 0, \quad \frac{\partial p_{ij}}{\partial \theta_k^s} = 0, \quad \text{if } k \neq i, j$$

Differentiating both sides of (3.7) with respect to θ_i^e and θ_i^s respectively yields

$$
\begin{pmatrix} \frac{\partial s_1}{\partial \theta_i^e} \\ \vdots \\ \frac{\partial s_i}{\partial \theta_i^e} \\ \vdots \\ \frac{\partial s_n}{\partial \theta_i^e} \end{pmatrix}
= r_s P
\begin{pmatrix} \frac{\partial s_1}{\partial \theta_i^e} \\ \vdots \\ \frac{\partial s_i}{\partial \theta_i^e} \\ \vdots \\ \frac{\partial s_n}{\partial \theta_i^e} \end{pmatrix}
+
\begin{pmatrix} r_s \frac{\partial p_{1i}}{\partial \theta_i^e} s_i^* \\ \vdots \\ r_s \sum_{j \neq i} \frac{\partial p_{ij}}{\partial \theta_i^e} s_j^* - r_e \alpha \\ \vdots \\ r_s \frac{\partial p_{ni}}{\partial \theta_i^e} s_i^* \end{pmatrix}
$$

and

$$
\begin{pmatrix} \frac{\partial s_1}{\partial \theta_i^s} \\ \vdots \\ \frac{\partial s_i}{\partial \theta_i^s} \\ \vdots \\ \frac{\partial s_n}{\partial \theta_i^s} \end{pmatrix}
= r_s P
\begin{pmatrix} \frac{\partial s_1}{\partial \theta_i^s} \\ \vdots \\ \frac{\partial s_i}{\partial \theta_i^s} \\ \vdots \\ \frac{\partial s_n}{\partial \theta_i^s} \end{pmatrix}
+
\begin{pmatrix} r_s \frac{\partial p_{1i}}{\partial \theta_i^s} s_i^* \\ \vdots \\ r_s \sum_{j \neq i} \frac{\partial p_{ij}}{\partial \theta_i^s} s_j^* + \beta \\ \vdots \\ r_s \frac{\partial p_{ni}}{\partial \theta_i^s} s_i^* \end{pmatrix}
$$

After simple algebra we obtain

$$
\begin{pmatrix}
\frac{\partial s_1}{\partial \theta_i^e} \\
\vdots \\
\frac{\partial s_i}{\partial \theta_i^e} \\
\vdots \\
\frac{\partial s_n}{\partial \theta_i^e}
\end{pmatrix}
= (I - r_s P)^{-1}
\begin{pmatrix}
r_s \frac{\partial p_{1i}}{\partial \theta_i^e} s_i^* \\
\vdots \\
r_s \sum_{j \neq i} \frac{\partial p_{ij}}{\partial \theta_i^e} s_j^* - r_e \alpha \\
\vdots \\
r_s \frac{\partial p_{ni}}{\partial \theta_i^e} s_i^*
\end{pmatrix}
\tag{3.9}
$$

and

$$
\begin{pmatrix}
\frac{\partial s_1}{\partial \theta_i^s} \\
\vdots \\
\frac{\partial s_i}{\partial \theta_i^s} \\
\vdots \\
\frac{\partial s_n}{\partial \theta_i^s}
\end{pmatrix}
= (I - r_s P)^{-1}
\begin{pmatrix}
r_s \frac{\partial p_{1i}}{\partial \theta_i^s} s_i^* \\
\vdots \\
r_s \sum_{j \neq i} \frac{\partial p_{ij}}{\partial \theta_i^s} s_j^* + \beta \\
\vdots \\
r_s \frac{\partial p_{ni}}{\partial \theta_i^s} s_i^*
\end{pmatrix}
\tag{3.10}
$$

By construction the returns r_e and r_s and the parameters α and β should be nonnegative, so the components of matrix $(I - r_s P)^{-1} = \sum_{k=0}^{\infty} r_s^k P^k$ must all be nonnegative. Moreover, it is reasonable to assume that social skills increase the probability of forming a link, i.e.,

$$
\frac{\partial p_{ij}}{\partial \theta_i^s} \geq 0, \quad \frac{\partial p_{ji}}{\partial \theta_i^s} \geq 0
$$

Hence,

$$
\frac{\partial s_i^*}{\partial \theta_i^s} \geq 0, \quad \frac{\partial s_j^*}{\partial \theta_i^s} \geq 0
$$

The signs of $\frac{\partial s_i^*}{\partial \theta_i^e}$ and $\frac{\partial s_j^*}{\partial \theta_i^e}$ are less clear. For example, supppose $\frac{\partial p_{ij}}{\partial \theta_j^e} = 0$. From (3.10) we get

$$
\frac{\partial s_i}{\partial \theta_i^e} = (I - r_s P)_{ii}^{-1} \left(r_s \sum_{j \neq i} \frac{\partial p_{ij}}{\partial \theta_i^e} s_j^* - r_e \alpha \right)
$$

103

where $(I - r_s P)_{ii}^{-1} > 0$ is the (i, i) element of matrix $(I - r_s P)^{-1}$. If we assume $\frac{\partial p_{ij}}{\partial \theta_i^e} > 0$, the sign of $\frac{\partial s_i}{\partial \theta_i^s}$ could be either positive or negative. Intuitively, smarter students tend to study more because their study effort yields higher return. On the other hand, smarter students are more attractive among peers, so a unit of time they spend on socializing will bring them more connections, providing incentive to socialize more. The net effect depends on the relative magnitude of these two opposite effects.

3.4.3 Estimation

Our primary parameters of interest are r_e and r_s. We seek to estimate them from the wage equation

$$\ln W_i = r_e E_i^* + r_s S_i^* + \gamma' x_i + \varepsilon_i \tag{3.11}$$

Recall that

$$\mathbb{E}\left[E_i^* \mid s, w\right] = \alpha \theta_i^e e_i^*$$

$$\mathbb{E}\left[S_i^* \mid s, w\right] = \sum_{j \neq i} p_{ij} s_i^* s_j^*$$

Since s_i^* is a function of v_i, if we allow v_i to be correlated with ε_i,

$$\mathbb{E}\left[\varepsilon_i \mid w, v_i\right] \neq 0$$

then both E_i^* and S_i^* are correlated with ε_i, through their means. Such correlation will yield biases in the OLS estimates of r_e and r_s.

3.5 Homophily and Instrument Construction

Our model of the party-study decision, and the subsequent production functions and earnings, leaves us with one primary endogeneity concern for our wage equation. Both E and S are functions of optimal socializing decision s^*, which in part depends on unobservable utility derived from social interaction v_i. To the extent this unobserved component of utility during socialization is correlated with ε_i, the unobservable in our wage equation, our OLS estimates of both r_e and r_s will be biased.

As a point of emphasis, recall that our model predicts specific signs for these biases with regard to estimates of r_e and r_s under reasonable assumptions. To reiterate, we predict a positive bias in OLS for the estimate of r_e, and a negative bias in the estimate for r_s. This means that generally we expect the estimate of eight friends per year of schooling (from Table 3.4) to be an approximate lower bound on the ratio of r_e/r_s.[12]

We adopt a few approaches to mitigate this issue. First, consider the case in which the unobservable social factor is due to the schooling environment. It is possible that some schools are better at fostering a dynamic social environment that is conducive to friendship formation. This may be due to active after-school programs, including athletics and arts programs, or may be due to location specifics. For example, in many suburban high schools students are picked up from school by parents immediately after class lets out, whereas in urban schools some students walk home, which they may do in groups. Still in other circumstances students may take a bus. In each scenario the after-school environment tied to a particular school differs in the intensity of social interaction it fosters. We present results with school-level fixed effects below to account for this concern.

Once we include school fixed effects, our remaining concern is on variation in socialization that

---

[12]The sign of the bias on r_s is determined by $b_s = \underbrace{Cov(v,\varepsilon)}_{\leq 0} \underbrace{\frac{\partial S}{\partial v}}_{\geq 0} \leq 0$, and similarly $b_e = \underbrace{Cov(v,\varepsilon)}_{\leq 0} \underbrace{\frac{\partial E}{\partial v}}_{\leq 0} \geq 0$. We assume $Cov(v,\varepsilon) \leq 0$, so that social preference v, which mimics a preference for leisure, is weakly negative correlated with unobservables that benefit wages later in life. The partial derivatives on E and S follow from the fact that v improves s^*.

occurs within a school at the individual level. We observe a large set of other variables about how individuals perceive themselves in a social setting, both during high school and in later periods. These include variables about responses to whether or not someone is popular, the life of the party, active in social situations, or has feelings of isolation. This set of variables, which we refer to as the set of "social controls", are included in select specifications below.[13]

Still, there remains the possibility of unobservable social factors that may influence our results. Fortunately, our model establishes a framework for focusing on the type of variation we need in candidate instruments. In particular, recall equations (3.7) and (3.8), our best response function and Bayesian Nash equilibrium, respectively. Our instruments should affect s^* through channels distinct from v. Any characteristics that affect p_{ij} (the elements of P), including individual and peer demographic, education and social characteristics, are candidate instruments in this framework.[14]

We introduce a set of candidate instruments based on homophily to address these issues. In our full specification we need at least two instruments z_i: one for E and one for S, such that:

$$\mathbb{E}\left[\varepsilon_i \,|\, w, z_i\right] = 0 \tag{3.12}$$

and z is correlated with E_i^* and S_i^*. We can then estimate our wage equation by 2SLS and the estimators will be consistent.

Moreover, we need there to be sufficient, independent variation in these instruments to isolate r_e and r_s separately. In the instrument sets discussed below, we aim to satisfy the following conditions.

1. *Return to education.* The resulting estimate of r_e should fall within the existing range of

---

[13]This is a similar approach to that taken by Conti et al. (2013) who include an additional set of regressors in their estimates of a return to friend nominations, including social participation, number of outings with friends, relative, marital status, and whether or not one's job was found through a network.

[14]This is the most natural, but not the only, channel for a valid instrument to operate. It is also true that other factors that influences E and S separate s^*. We discuss some of these, including class size and teacher homopily, below.

estimates.[15]

2. *OLS Biases.* We expect the predicted biases from our model that persist in OLS to be remedied, such that the estimated r_e/r_s declines.

3. *Instrument partition.* Our set of instruments should partition into a set strongly correlated with in-degree, and another strongly correlated with years of schooling, with some that potentially overlap and are significantly correlated with both.[16]

3.5.1 Student Homophily Instruments

Even with our expanded set of social controls, there persists the concern that something remains omitted that is both related to social utility and wages and thus still may bias our estimates. We include a set of estimates that adopt an instrumental variable approach following the social network literature. One of the distinguishing characteristics of social networks relative to other network types is their exhibition of homophily, or the property "that people are more prone to maintain relationships with people who are similar to themselves" (see Jackson (2005)). Other recent work on social networks highlights the role of homophily in friendship formation (Jackson, Rogers, and Zenou (2016) and Graham (2015) include discussions of social network homophily).

We construct a set of *student homophily distance* variables, starting with a set of individual student characteristics including gender, age, and years a student has been at their school. We also compute homophily distance based on our endowment variables: IQ and shyness. We aggregate these characteristics at the cohort (school-grade) level, and compute for each student the distance between their own characteristics and the aggregate characteristics of their cohort. Intuitively,

---

[15]Surveys in Card (1999) and Card (2001) set this range from 4.6-16.4% per year of schooling.

[16]This is meant in the spirit of other empirical work with two endogenous regressors in that we will test for weakness of the instruments, and there is sufficient orthogonality between $\hat{E}$ and $\hat{S}$. We will make a point to discuss this carefully when reviewing first-stage results.

the closer students are to their cohort peers, the more likely they are to find themselves in social interactions and thus the more likely they are to form friendships.

As an example, consider a student whose has been at their current school for two years. New students to a school may find it easier to make friends with other new students rather than students who have been in the school system longer, themselves having formed friendships with peers with whom they've already interacted. Define the vector of *student* homophily distances, $h_i^s = |x_{ik}^s - \bar{x}_k^s|$, where k denote a cohort and x^s denotes the set of characteristics used for student homophily.[17].

While the source of the endogeneity is identical for both E and S, their separate appearance in the wage equation necessitates using at least two instruments. Luckily, the benefits of social closeness are not limited to friendship formation. Some of these student homophily distance variables influence studying and education outcomes as well. For example, Black, Devereux, and Salvanes (2011) explore the role of school starting age on education outcomes, finding a range of small and large effects on a range of outcomes, including educational attainment and teenage pregnancy, in a sample of individuals from Norway. Starting school early or late is captured in our age homophily distance variable, and we will demonstrate it has predictive power for both in-degree and years of schooling.

Figures 3.4 and 3.5 include nonparametric plots of in-degree and education outcomes and functions of student homphily distances. In each plot, all variables include residuals from fixed effects estimates at the school level that "partial-out" variation from our basic set of individual controls. Table 3.9 includes related estimates from regressions that equate to the first-stage using these homophily variables. First, notice that certain student homophily variables are more strongly correlated with in-degree and others are more strongly correlated with education outcomes, with some

---

[17]Note that we do not use the full set of characteristics x, but rather the set available from the AH in-school sample. This allows us to compute actual means from the full school populations rather than weight-adjusted means from the smaller in-home samples. Also, note that $\bar{x}_k^s = \frac{1}{n_k} \sum_{i \in k} x_i^s$.

variables have significant associations with both.

- **In-degree influencers:** these variables are significantly associated with in-degree but not years of schooling.

  - *Shyness distance* is negatively correlated with in-degree, which is partially consistent with intuition: shy people in an outgoing cohort are less likely to form friendships. Interestingly, the preference for similar friends also plays a role here in that outgoing people in a relatively shy cohort, who also have large shyness distance values, convert fewer relationships into friendships. A change from being the only outgoing person in a shy cohort, with a shyness distance near one, to a cohort of all outgoing people, thus shyness distance equal to zero, equates to 1.6 more friend nominations.

  - *Female distance* is positively correlated with in-degree. Larger female student homophily distance indicates the individual is among the gender minority in their school, which may be conducive to making friends of the opposite gender. The magnitude of the coefficient suggests that being the difference in in-degree from being female in an all female school to female in an all male school is two friend nominations.

- **Years of schooling influencers:** these variables are significantly associated with years of schooling but not in-degree.

  - *IQ distance* is positively associated with years of schooling, suggesting that being both more or less intelligent within a cohort is better for education outcomes relative to being at the mean. One explanation is that smarter students can stand out more in a less intelligent cohort, and less intelligent students can benefit from the peer effects of being in a smarter cohort, therefore distance plausibly improves outcomes in both instances.

- **Influencers of both:**

– *Age distance* is negatively associated with both in-degree and years of schooling, so that students who are older or younger within a cohort on average have fewer friends and lower years of schooling.

These results suggest that the homophily instruments satisfy standard relevance assumptions: the Cragg-Donald F-statistics in Table 3.9 are 13.09 and 23.77 for in-degree and years of schooling, respectively, which pass standard tests for weak instruments.[18]

Alternative Candidates for Education Instruments

In addition to the student homophily instruments there are other variables that fit the set of initial criteria we used for instrument selection. Many of these operate through channels of E or S and not s^*, and are motivated by examples in the existing literature. However, we describe below the shortcomings of these variables, and prefer the student homophily variables.[19]

- *Teacher homophily.* Gershenson, Holt, and Papageorge (2016) demonstrate that there are notable differences in achievement between students who have teachers with similar demographic characteristics. Specifically, they find among 3rd - 5th grade low income black boys, having a black teacher increases interest in college by 29% and reduces dropout rates by 39%. Therefore teacher homophily is a logical candidate instrument in that it may influence student studying patterns, affecting wages only through its subsequent impact on education outcomes like years of schooling. We constructed a set of teacher homophily variables based on gender and race, and while their inclusion returned comparable estimates to those using only student homophily instruments the resulting first-stage estimates were weak. Additionally, the use of

---

[18]The relevant critical values we can reject from Stock and Yogo (2002) are 10.83 (10% maximal IV relative bias) for in-degree and 18.37 (5% maximal IV relative bias) for years of schooling.

[19]Supplementary estimates using these sets of instruments can be found in Table 3.11. The results are broadly consistent with those using only the set of student homophily instruments, however these sets of instruments are generally weaker.

race-based homophily variables raised concerns about selection into schools with certain characteristics, potentially violating the exclusion restriction. And finally, while baseline estimates corroborated some findings in the literature, such as a benefit to black students having black teachers, these relationships did not remain significant in the presense of other controls.[20]

- *Class size.* Class size could plausibly influence both in-degree and years of schooling. In-degree may vary with class size in a number of different ways; larger classes may benefit friendship formation given that the pool of potential friends is larger, or they may inhibit friendship formation, as matching with the right people may be more challenging in larger cohorts and classes. A well-known study of the influence of class size on academic achievement is Angrist and Lavy (1999), which uses Maimonides' Rule to identify the effect of class size on test scores for elementary school students, finding positive effects of smaller class sizes for certain groups of students. Our first-stage estimates of class size found similar patterns, with smaller classes positively and significantly correlated with years of schooling, although these effects were highly non-linear. We also found in most specifications that in-degree was correlated with larger class size. Estimates that include class size instruments are largely consistent with the results we present below. However, most of class size estimates failed to reject that these instruments were weak.

- *Academic network.* The academic network data in AH contains information on the course overlap among students in schools. We can construct an alternative network formulation, distinct from the nomination network, where links exist if students take courses with one another. Moreover, these links are weighted by the intensity of the course overlap (i.e. taking

---

[20]One other issue specific to the teacher homophily variables in AH is that teacher demographic composition is available only at the school level, rather than cohort. This means that the remaining variation after including full controls and school fixed effects is minimal, and thus this approach may be valid with more detailed teacher information.

all of the same courses has a higher weight than if two students only share one course). However, this exercise dramatically limits our sample, cutting it down from around 10,000 students to around 3,000.[21] This set up offers us a new set of candidate instruments for education: i's average of endowments among classmates (the academic network analog of friends). We computed the average IQ and average shyness among an individual's classmates, to potentially serve as additional instruments for education and in-degree, respectively. Again, the first stages for these instruments was very weak.

3.6 Results for Wages

In this section we present results that test whether or not the preliminary estimates of returns to socializing, r_s, remain in approaches that attempt to ameliorate specific endogeneity concerns. Whenever we estimate r_s we also estimate r_e, a return to a year of schooling. Jointly estimating both returns allows us to compute a ratio, equating a year of schooling to a number of nominations by friends. Additionally, given the wealth of estimates in the literature on the returns to schooling, one check of the validity of our estimates is to observe if the return to schooling estimate is a reasonable magnitude relative to existing results. Also, recall that the main source of bias, outlined in our model, suggests that the OLS estimate r_s is negatively biased, while the OLS return for r_e has a positive bias, meaning we expect the ratio r_e/r_s to decrease as the bias is minimized. Table 3.10 includes these estimates using the homophily instruments:

- *(1)* is the same specification from the preliminary estimates, indicating a r_e/r_s ratio of approximately eight.

- *(2)* includes school fixed effects, highlighting within-school variation to identify the returns.

---

[21]This is due to attrition during follow up in Wave 3.

Net of these school-level effects, the coefficients of interest move in the direction anticipated by the model, with the ratio r_e/r_s dropping from eight to six friends per year of schooling. While the school fixed effects address time-invariant, unobservable social factors at the institution level, individual social characteristics that are unobserved may still influence our estimates.

- *(3)* includes a set of additional self-reported social controls. The coefficients change minimally after these social variables are included, the r_s estimate decreases slightly from the fixed effect estimate, but is still higher than the OLS estimate. The estimate of r_e remains essentially unchanged with these social controls.

- *(4)* uses a homophily approach - rather than social controls - to address endogeneity of in-degree. In this specification we only instrument for in-degree using student homophily variables. As predicted by our model, the estimate of r_s increases and r_e decreases results in a $r_e/r_s < 1$.

- *(5)* uses teacher homophily to instrument for only years of schooling, with no additional controls for in-degree.

- *(6)* uses both sets of homophily instruments. Both return estimates are substantially higher than OLS estimates, yet the overall ratio drops so that $r_e/r_s \approx 3$.

One issue that persists with these estimates is that the return to schooling estimate when education is treated as endogenous are large. The full IV estimate in (6) of 0.177 is within range at the higher end of existing return to schooling IV estimates, and specification (5) where we only instrument for schooling returns a very high estimate for years of schooling of over 0.2.

The results from Table 3.10 support the notion that there are significant, non-zero returns to socializing. The coefficient on in-degree found in OLS is the lowest estimate among the specifications tested. However, there are important caveats worth mentioning. First, it is only when we instrument

just for in-degree that the estimates change following the bias predictions from the model. When education is treated as endogenous estimates of r_e increase to levels outside of the usual range in the literature. This raises concerns, especially given our prediction, that there is lingering a positive bias in the return to education estimate.

3.7 Application: Does it Pay to Attend a Party School?

Partying and socializing factor prominently in a number of ongoing debates about higher education. Many colleges and universities are under scrutiny for both the nature and extent of social activity on their campuses. As tuition universally seems to rise, many schools are facing a backlash, with some parents believing these increases are a function, in part, of unnecessary investments in non-academic spaces not fundamental to the perceived core purpose of attending college. The expense of American universities, in particular, has spawned a growing industry of massively open online courses (MOOCs), including sites like Coursera and Udemy, where university level courses are available at zero, or at most a fraction of, the cost charged to physically attend courses that, in content, are virtually identical to those offered on campus. The precedent for this type of education provision, albeit not quite as accessible, was set by many for-profit universities such as the University of Phoenix, who grew due to an early and prolific online presence. The appeal of these services is obvious in this setting: remove the expensive extracurriculars for a focus on the cognitive content and more students can access university level teaching and credentials.

Even within traditional universities, forms of social interaction that revolve around drinking and the Greek system have come under intense scrutiny, including efforts by many elite universities to eliminate these institutions entirely. In this section we want to explore a simple exercise related to our party-study decision above, and offer the results as preliminary motivation for further study. We have documented a sizable return to socializing, of comparable magnitude to a year of schooling,

114

at the individual level. In our model, there are conditions under which people students optimally choose non-zero levels of both studying and socializing, with the possibility still persisting for multiple equilibria that include all socializing or all studying. Using data at the university level, below we provide evidence that students who attend schools with thriving academic *and* social scenes face higher expected earnings upon graduation relative to both schools with either poor academics or poor social environments.

We start with rankings of universities' academic and social standing available from Niche (www.niche.com). Niche provides a range of rankings and letter grades of neighborhoods, schools, and universities in an array of categories. We use their academic grades, as well as their grades of party schools and Greek life from 2015. To provide a sense of the types of schools and their rankings in each variable, Table 3.13 is list of a randomly selected set of schools, including one from each letter grade pair of academic and party rank. We convert these letter grades into numerical GPA units to use in our analysis below.

We match this university level data with earnings averages collected by the US Department of Education and available publicly through the College Scorecard Data (https://collegescorecard.ed.gov/data/). There are two earnings variables that we focus on. First, we see the median earnings of students who matriculated at a given school ten years after entering the school, as well as the fraction of matriculating students who earn more than $25,000 per year six years after entering the school. We run regressions of the form:

$$\ln w_u = \alpha + \beta a_u + \gamma p_u + \delta(a_u \times p_u) + X_u\theta + \nu_u$$

Where u indexes universities, w is one of the aforementioned wage variables, a is an academic grade, p is a party grade, and X is a set of university-level characteristics, including the highest degree awarded by the institution, whether it is public versus private, and whether it is profit or

115

non-profit.

These results of the estimation using OLS are presented in Table 3.14. In all specifications the interaction term between party and academic atmospheres is positive and significant, indicating that students who attend schools with a strong academic program and a vibrant social scene earn more on average than those at institutions who are weaker in either category. Also, while it is true that even in a school limited or no Greek life there is a positive correlation between earnings and academics, the same is not true of the overall party atmosphere. Schools with non-existent party scenes do not face average earnings that increase with academic quality.

To better illustrate these interaction effects, see Figures 3.6 and 3.7. These figures plot the heterogeneous effects of academic quality over the distribution of party grades. For a handful of schools that have the least desirable party atmospheres, their academic grades have little to no association with earnings. Most schools have a positive effect from academics, falling sufficiently far right on the distribution of party grades. Partying is negatively associated with earnings for a small number of schools with low academic grades, but similarly is positive for most schools as they have sufficiently high academic grades.

The role of alcohol consumption in the US college social scene is widely publicized and controversial. As discussed above, alcohol consumption can be characterized as a social investment, and our estimates suggest that drinking effort benefits friendship formation at the expense of education outcomes. In previous analyses we presented linear estimates of this effect, which mask important variation with regard to alcohol consumption. In Figure 3.8 we plot nonparametric plots versus in-degree and years of schooling against days drunk in a month in Wave 3, the wave during which Add Health respondents are most likely in college. Both demonstrate a notable concavity with respect to drunkenness.[22] From these plots we can observe the range, of approximately two to five days

---

[22]Binge drinking is an especially great concern in college social settings, and this motivates our focus on drunkenness, rather than simply instances drinking, so that we can capture some information about intensity.

drunk per month, wherein social effort through drinking is correlated with in-degree but negatively correlated with education. How students make these decisions in college, and how they decide to opt into certain academic or socially ranked schools based on their endowments, is any interesting area of future study.

3.8   Conclusion

A year of schooling has long been a coveted commodity in the labor market. While the cognitive component of schooling remains important for productivity and subsequent labor market success, we have demonstrated that social skills, in particular those measured by the number of friends a person has, are another important and distinct portion of the skills accumulated during schooling.

We find that the wage return to one friend ranges from 1.4-6.7% percent, with the corresponding ratio of the return to education over the return to socializing ranging from three to eight. This means that our lowest estimates, which come from OLS, suggest a year of schooling has the same return as eight friends in high school. Two stage least squares estimates that control for the biases present in OLS indicate that this number may be as low as three.

Socializing, despite benefits that follow it in the labor market, does not come without cost. In deciding to spend time investing in friendships our results suggest that individuals forgo measurable levels of academic achievement when they decide to party. Despite common belief, this is not necessarily to the detriment of future outcomes. We find those friendships contribute dramatically to labor market prospects, and suggestive evidence indicates that other life outcomes including marriage and general health, may also improve with more friends.

The precise nature of the party-study decision, and how it aggregates above the individual and induces selection, is a natural next step in this line of work. Our initial estimates about college academics and social atmospheres suggest that there are benefits to vibrancy in both, and

it would be interesting to explore how high school students opt into certain institutions. Other significant decisions, including migration and occupation selection, may also depend critically on the mechanisms discussed in this paper.

Table 3.1: Add Health Survey Structure

Wave Number	Year	Ages (Grades)	Sample Size
1	1994 - 95	11 - 18 (7th - 12th)	In-School: 90,000; In-Home: 20,745
2	1996	12 - 19 (7th - 13th)	In-Home: 14,738
3	2001-02	18 - 27	In-Home: 15,197
4	2008	25 - 34	In-Home: 15,701

Table 3.2: Important Variables and Data Sources

Variable Name	Add Health Variable	Instrument	Alternative Data Source
Social endowment: θ^s	Shy (Wave 1)	–	BCS: Extrovert (Age 10), NLSY: Outgoing (Age 6)
Education endowment: θ^e	Picture Vocab Test (Wave 1)	–	BCS: Reading score (Age 5), NLSY: AFQT
Social investment: s_i	a. Drinking Intensity (Wave 3), b. Time Spent with Friends	Average Shyness of Friends	BCS: time on homework, goes out, drinking, NSLY: Hours studying, days drink
Social investment of others: s_j	–	–	BCS: –, NLSY: –
Education output: E_i	a. Years of schooling (Wave 4), b. GPA (Wave 1)	Average PVT of classmates	BCS: Age left school, NLSY: Years of schooling
Social output: S_i	In-degree (Wave 1 & 2)	Average Homophily Distance to Cohort (Wave 1)	BCS: Number of Friends (Age 34), NLSY: Outgoing as Adult
Wages: W_i	Annual Earnings	–	BCS: Annual Earnings, NLSY: Annual Earnings

Table 3.3: Add Health Summary Statistics

	(1) Total	(2) Shyness (θ^s):		(3) IQ (θ^e):	
	_	0	1	0	1
Annual Earnings (W4)	34151.2	34865.9	32813.7	29569.3	38169.9
	(41572.7)	(42474.0)	(39801.9)	(37687.0)	(44318.3)
In-Degree (W1)	4.493	4.770	3.983	4.111	4.832
	(3.711)	(3.871)	(3.336)	(3.481)	(3.872)
Frequent Drinking (W1)	0.157	0.171	0.133	0.151	0.164
	(0.364)	(0.376)	(0.339)	(0.358)	(0.370)
Frequent Drinking (W2)	0.185	0.203	0.151	0.172	0.197
	(0.388)	(0.403)	(0.358)	(0.378)	(0.398)
Frequent Drinking (W3)	0.341	0.355	0.316	0.264	0.417
	(0.474)	(0.479)	(0.465)	(0.441)	(0.493)
Time w Friends (W1)	0.400	0.422	0.360	0.415	0.386
	(0.490)	(0.494)	(0.480)	(0.493)	(0.487)
Shyness (W1)	0.350	0	1	0.372	0.329
	(0.477)	(0)	(0)	(0.483)	(0.470)
Years of Schooling (W4)	14.53	14.58	14.44	13.80	15.20
	(2.224)	(2.220)	(2.229)	(2.215)	(2.014)
GPA (W2)	2.792	2.805	2.770	2.618	2.953
	(0.757)	(0.755)	(0.760)	(0.723)	(0.751)
Study Effort (W1)	0.398	0.386	0.420	0.434	0.367
	(0.490)	(0.487)	(0.494)	(0.496)	(0.482)
Pic Vocab Test (W1)	99.81	100.5	98.48	88.33	112.1
	(14.97)	(14.40)	(15.88)	(10.30)	(7.574)
Observations	14713	14713		14713	

mean coefficients; sd in parentheses

Means by group reported; standard deviations in parentheses.

Wave number is roman numeral in parethesis.

denotes drinking more than 2-3 times per month.

Log Earnings on Schooling

Schooling by Wave IV.
Limited to white individuals.
Both series include 95% confidence intervals.
Source: Add Health restricted-use data.

(a)

Log Earnings on In-Degree

In-degree is equal to the number of friends who nominate ego.
Limited to white individuals.
Both series include 95% confidence intervals.
Source: Add Health restricted-use data.

(b)

Figure 3.1: Log Earnings on Education and Social Outcomes

Figure 3.2: In-Degree and Years of Schooling

IQ by Shyness

Limited to white individuals.
Both series include 95% confidence intervals.
Source: Add Health restricted-use data.
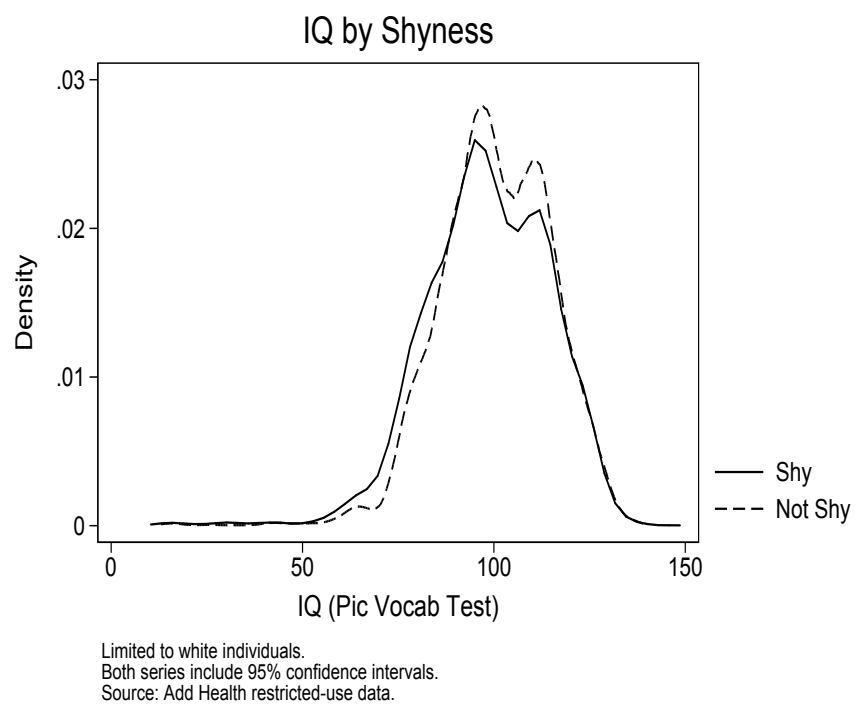
Figure 3.3: Shyness and IQ

Table 3.4: Log Wages on Friends and Schooling - Add Health

	(1)	(2)
	Log Wages (W4)	Log Wages (W4)
In-Degree (W1)	0.0143***	0.0172***
	[0.00270]	[0.00291]
Years of Schooling (W4)	0.113***	0.118***
	[0.00533]	[0.00649]
Observations	9738	6945
R^2	0.121	0.135
Full Controls	Yes	Yes
Endowments	No	Yes

Source: Add Health restricted-use data.

* $p < .1$, ** $p < 0.05$, *** $p < 0.01$

Table 3.5: Log Wages on Friends and Schooling - BCS

	(1)	(2)
	Log Wages (Age 42)	Log Wages (Age 42)
Number of Friends (Age 34)	0.0314***	0.0296***
	[0.00526]	[0.00735]
Highest Age Left Ed (Age 42)	0.0484***	0.0454***
	[0.00335]	[0.00524]
Observations	5103	2823
R^2	0.252	0.248
Full Controls	Yes	Yes
Endowments	No	Yes

Source: BCS.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 3.6: Social and Education Outputs on Investments and Endowments - Add Health

	(1) In-Degree (W1)	(2) In-Degree (W1)	(3) Years of Schooling (W4)	(4) Years of Schooling (W4)	(5) GPA (W2)	(6) GPA (W2)
Shyness (W1)	-0.603*** [0.0780]	-0.597*** [0.0781]	0.00992 [0.0450]	0.00317 [0.0451]	0.00742 [0.0158]	0.000830 [0.0158]
Pic Vocab Test (W1)	0.0157*** [0.00290]	0.0159*** [0.00290]	0.0468*** [0.00184]	0.0468*** [0.00184]	0.0111*** [0.000619]	0.0111*** [0.000619]
Study Effort (W1)	-0.221*** [0.0809]	-0.194** [0.0809]	0.287*** [0.0449]	0.274*** [0.0450]	0.258*** [0.0154]	0.249*** [0.0154]
Time w Friends (W1)	0.487*** [0.0809]	0.460*** [0.0809]	-0.226*** [0.0444]	-0.211*** [0.0446]	-0.0670*** [0.0154]	-0.0553*** [0.0155]
Frequent Drinking (W2)		0.315*** [0.105]		-0.191*** [0.0576]		-0.164*** [0.0206]
Observations	8791	8775	7653	7640	8556	8542
R^2	0.058	0.059	0.250	0.251	0.161	0.168
Full Controls	Yes	Yes	Yes	Yes	Yes	Yes
Drinking	No	Yes	No	Yes	No	Yes

Source: Add Health restricted-use data.

* $p < .1$, ** $p < 0.05$, *** $p < 0.01$

Table 3.7: Social and Education Outputs on Investments and Endowments - NLSY

	(1) Outgoing as Adult	(2) Outgoing as Adult	(3) Years of Schooling	(4) Years of Schooling
Days Drinking (Month)	0.00321** [0.00113]	0.00264* [0.00110]	-0.0145** [0.00480]	-0.0155*** [0.00443]
Hours Study	0.000136 [0.00110]	0.000173 [0.00108]	0.0309*** [0.00495]	0.0269*** [0.00424]
Outgoing in Youth		0.294*** [0.0145]		0.0698 [0.0663]
AFQT		-0.000109 [0.000368]		0.0419*** [0.00158]
Observations	2881	2881	2845	2845
R^2	0.017	0.122	0.335	0.476
Endowments	No	Yes	No	Yes

Source: NLSY79 (restricted to observations in 1994).

$* p < 0.05$, $** p < 0.01$, $*** p < 0.001$

127

Table 3.8: Social and Education Investments on Endowments

|	(1) Study Effort (W1)	(2) Time w Friends (W1)	(3) Frequent Drinking (W2)	(4) Frequent Drinking (W3)
Shyness (W1)	0.0317***	-0.0600***	-0.0479***	-0.0201**
	[0.0102]	[0.00861]	[0.00652]	[0.00806]
Pic Vocab Test (W1)	-0.00201***	-0.00160***	0.000127	0.00406***
	[0.000381]	[0.000312]	[0.000228]	[0.000281]
Observations	9783	14049	14015	14050
R^2	0.049	0.010	0.061	0.077
Full Controls	Yes	Yes	Yes	Yes

Source: Add Health restricted-use data.
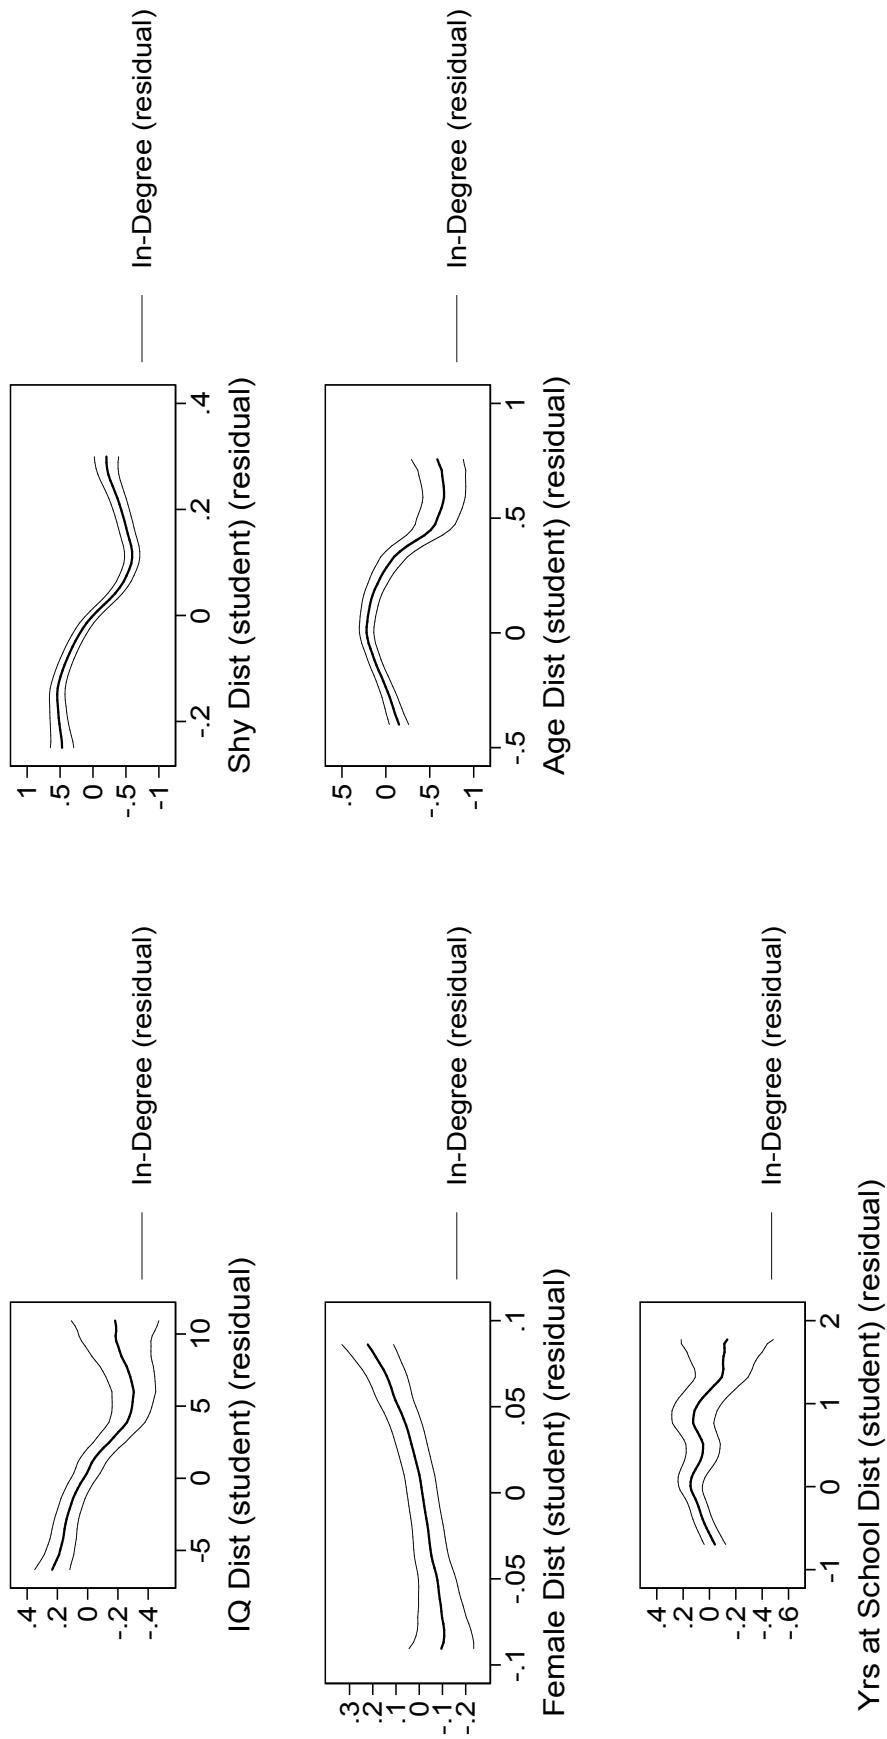
* $p < .1$, ** $p < 0.05$, *** $p < 0.01$

Residuals indicate variation from full controls and school fixed effect has been partialled out.
5th-95th percentile of homophily variables plotted.
Both series include 95% confidence intervals.
Source: Add Health restricted-use data.

Figure 3.4: In-Degree and Student Homophily
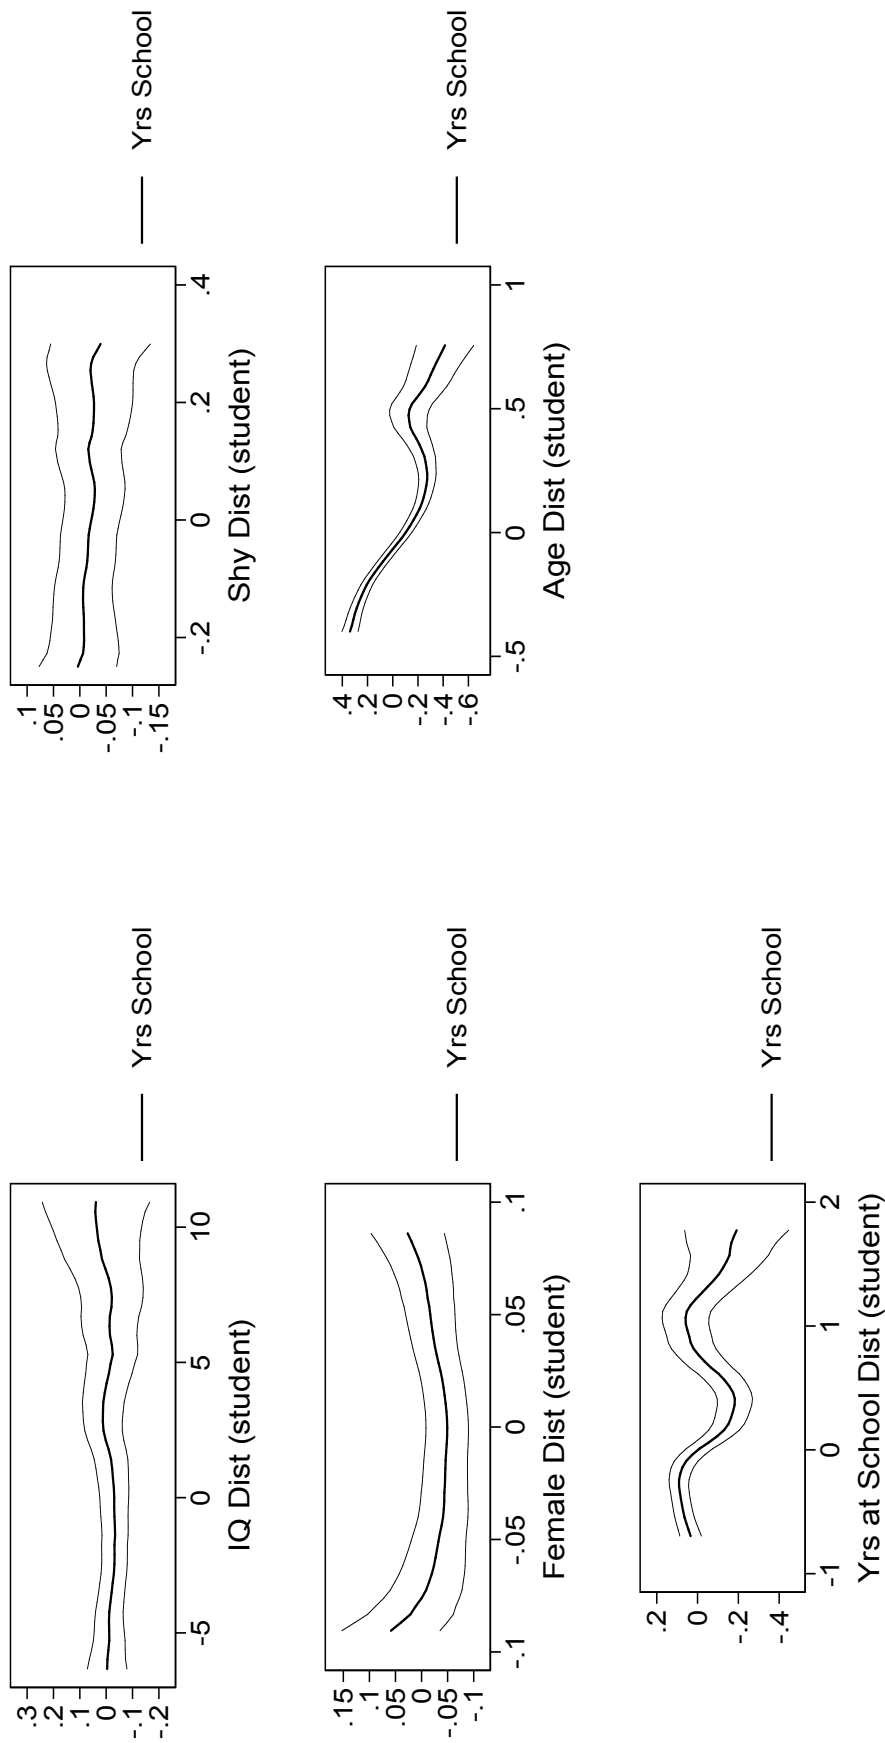
129

Figure 3.5: Education Outcomes and Student Homophily

Residuals indicate variation from full controls and school fixed effect has been partialled out.
5th-95th percentile of homophily variables plotted.
Both series include 95% confidence intervals.
Source: Add Health restricted-use data.

Table 3.9: First Stage Estimates - Student Homophily

	(1) In-Degree (W1)	(2) Years of Schooling (W4)
Female Dist (student)	2.129***	0.401
	[0.682]	[0.376]
Age Dist (student)	-0.569***	-0.708***
	[0.131]	[0.115]
Yrs at School Dist (student)	-0.287***	-0.0314
	[0.0672]	[0.0366]
IQ Dist (student)	-0.0125	0.0106***
	[0.00796]	[0.00395]
Shy Dist (student)	-1.645***	-0.0732
	[0.204]	[0.123]
Observations	9444	8230
Cragg-Donald Wald F	13.09	23.77

Standard errors in brackets

Source: Add Health restricted-use data.

All include school-level fixed effects.

* $p < .1$, ** $p < 0.05$, *** $p < 0.01$

Table 3.10: Two Stage Least Square Estimates - Student Homophily

	(1) Log Wages (W4)	(2) Log Wages (W4)	(3) Log Wages (W4)	(4) Log Wages (W4)	(5) Log Wages (W4)	(6) Log Wages (W4)
In-Degree (W1)	0.0143***	0.0173***	0.0154***	0.0843**	0.0151***	0.0662*
	[0.00282]	[0.00292]	[0.00349]	[0.0341]	[0.00439]	[0.0360]
Years of Schooling (W4)	0.113***	0.105***	0.103***	0.0975***	0.206***	0.177***
	[0.00628]	[0.00590]	[0.00669]	[0.00983]	[0.0480]	[0.0524]
Observations	9738	9738	7949	6873	6873	6873
R^2	0.121	0.093	0.101	0.054	0.074	0.059
Cragg-Donald Wald F	.	.	.	13.09	23.77	10.67
Full Controls	Yes	Yes	Yes	Yes	Yes	Yes
School Fixed Effects	No	Yes	Yes	Yes	Yes	Yes
Social Controls	No	No	Yes	No	No	No
Social Endogenous	No	No	No	Yes	No	Yes
Education Endogenous	No	No	No	No	Yes	Yes

Standard errors in brackets

Source: Add Health restricted-use data.

Social controls include life of the party, popularity, social at parties, and spending time with friends.

Student homophily IVs used for any endogenous regressor.

* $p < .1$, ** $p < 0.05$, *** $p < 0.01$

Table 3.11: Two Stage Least Square Estimates - Adding Class Size and Teacher Homophily

	(1) Log Wages (W4)	(2) Log Wages (W4)	(3) Log Wages (W4)	(4) Log Wages (W4)
In-Degree (W1)	0.0662*	0.0575*	0.0897**	0.0787**
	[0.0360]	[0.0348]	[0.0366]	[0.0352]
Years of Schooling (W4)	0.177***	0.202***	0.158***	0.185***
	[0.0524]	[0.0446]	[0.0520]	[0.0445]
Observations	6873	6873	6841	6841
R^2	0.059	0.054	0.032	0.036
Cragg-Donald Wald F	10.67	7.36	7.56	5.86
Full Controls	Yes	Yes	Yes	Yes
Fixed Effects	Yes	Yes	Yes	Yes
Student Homophily	Yes	Yes	Yes	Yes
Class Size	No	No	No	Yes
Teacher Homophily	No	No	Yes	Yes

Standard errors in brackets

Source: Add Health restricted-use data.

* $p < .1$, ** $p < 0.05$, *** $p < 0.01$

Table 3.12: Other Outcomes on Friends and Schooling

	(1) Married at Least Once (W4)	(2) General Health Status (W4)	(3) BMI (W4)
In-Degree (W1)	0.00368***	0.0131***	-0.295***
	[0.00136]	[0.00230]	[0.0454]
Years of Schooling (W4)	0.00636**	0.0726***	-0.336***
	[0.00244]	[0.00528]	[0.0837]
Observations	10917	10924	10924
R^2	0.042	0.051	0.016

Source: Add Health restricted-use data. Includes school fixed effects and full controls.

* $p < .1$, ** $p < 0.05$, *** $p < 0.01$

Table 3.13: Examples of Universities by Academics and Party School Grades

Name	City	State	Academics	Party School
Wesleyan University	Middletown	CT	A	A
University of San Diego	San Diego	CA	A	B
University of Texas - Dallas	Richardson	TX	A	C
University of North Texas	Denton	TX	B	A
Belmont University	Nashville	TN	B	B
Columbia International University	Columbia	SC	B	C
Warner University	Lake Wales	FL	B	D
University of Wisconsin - Oshkosh	Oshkosh	WI	C	A
Presentation College	Aberdeen	SD	C	B
Arkansas Tech University	Russellville	AR	C	C
Penn State Brandywine	Media	PA	C	D
Paul Quinn College	Dallas	TX	D	B
Livingstone College	Salisbury	NC	D	C

A-, A, A+ all categorized to A, and so on. Schools chosen at random.

Table 3.14: Earnings, Academics and Partying - by Grade/Score

	(1) Ln Median Earnings (10 YSE)	(2) Ln Median Earnings (10 YSE)	(3) Earnings > 25K (6 YSE)	(4) Earnings > 25K (6 YSE)
Academics	-0.00586 [0.0256]	0.0699*** [0.0122]	-0.00894 [0.0129]	0.0132* [0.00597]
Party School	-0.104*** [0.0225]		-0.0418*** [0.0116]	
Academics × Party	0.0487*** [0.00815]		0.0194*** [0.00414]	
Greek Life		-0.0980*** [0.0139]		-0.0554*** [0.00709]
Academics × Greek Life		0.0394*** [0.00479]		0.0210*** [0.00234]
Observations	1422	1422	1429	1429
R^2	0.309	0.325	0.203	0.233

Standard errors in brackets

YSE denotes years since entering school.

Controls include highest degree awarded, and public vs private, profit vs. non-profit.

Academic, party, and Greek life grades are on a gpa scale 0 - 4.33.

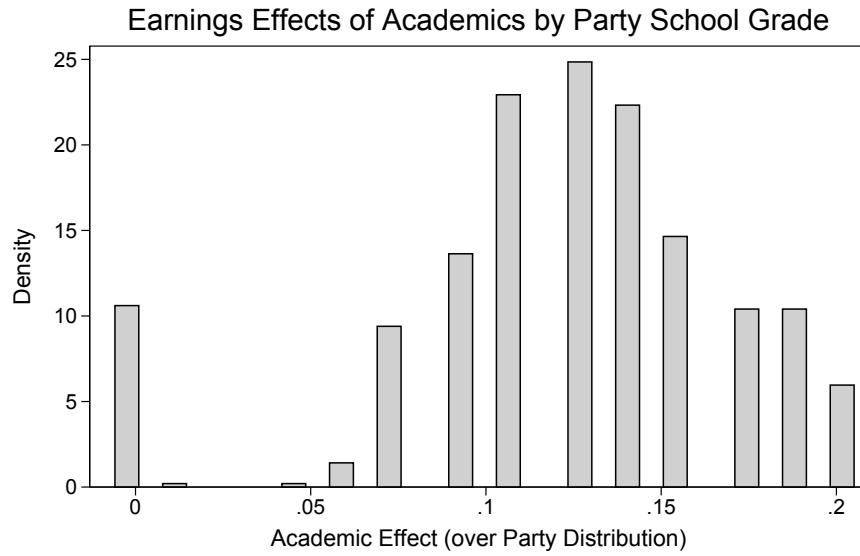* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

136

Figure 3.6: Earnings by Academic Grade

Figure 3.7: Earnings by Party Grade

In-Degree

Years of Schooling
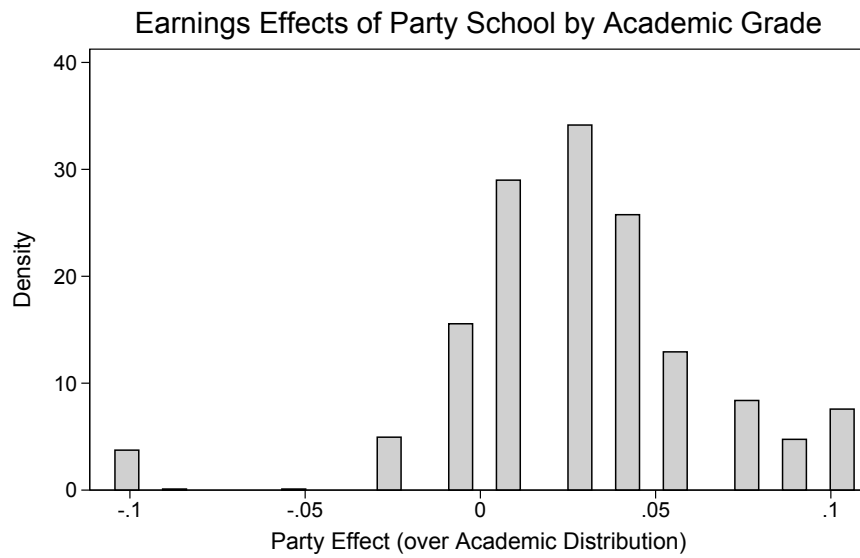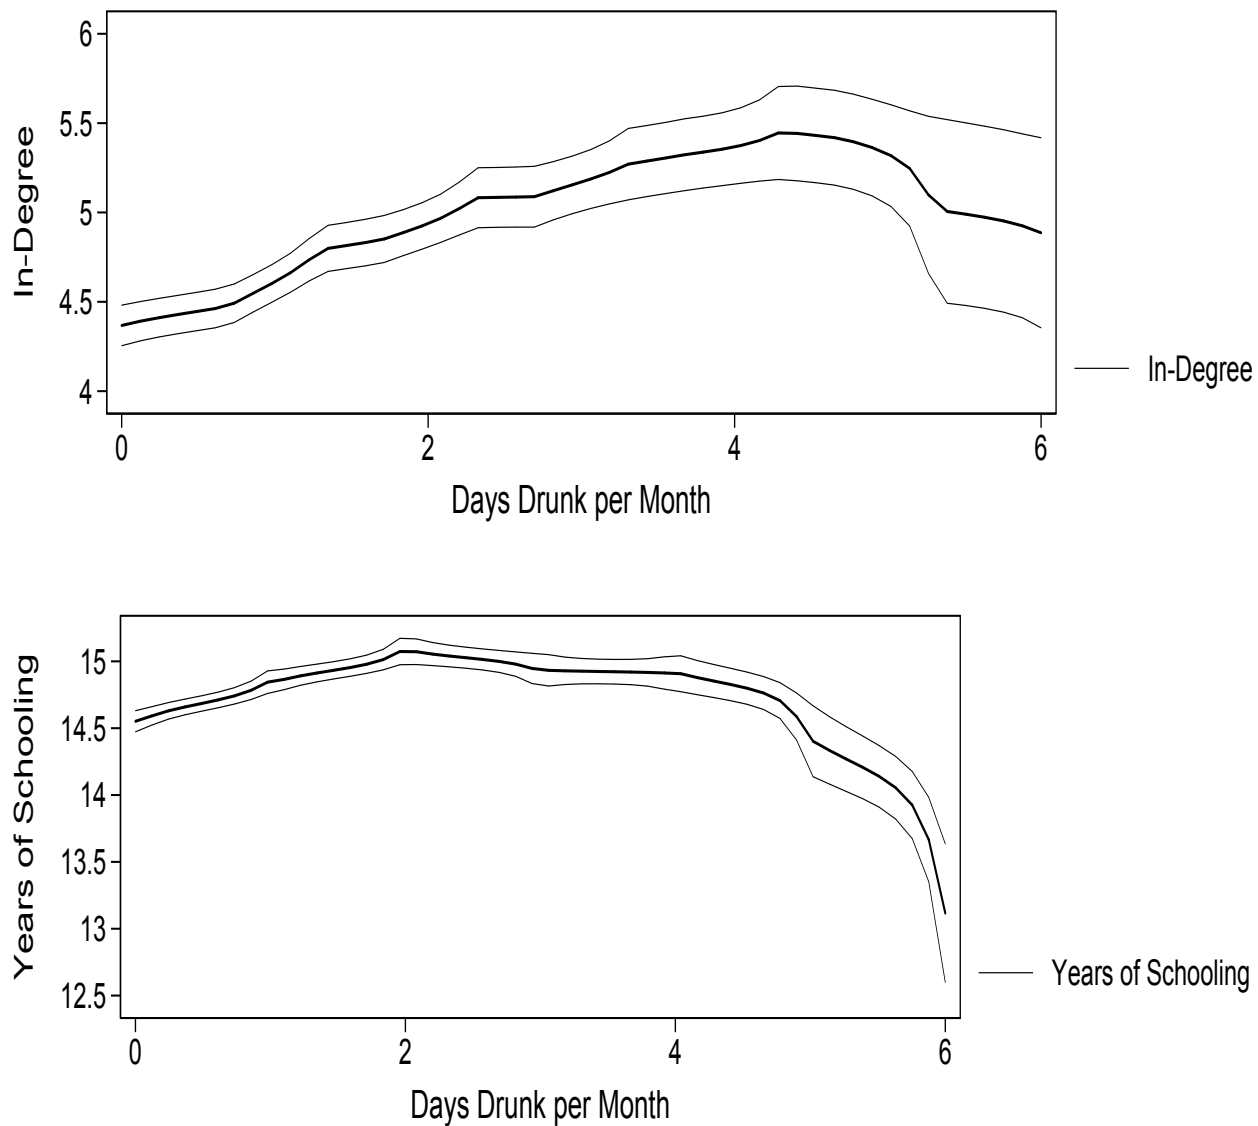
Both series include 95% confidence intervals.
Source: Add Health restricted-use data.

Figure 3.8: Education, Social Outcomes and Drunkenness

Bibliography

Angelov, Nikolay, Per Johansson, and Erica Lindahl. 2016. "Parenthood and the Gender Gap in Pay." *Journal of Labor Economics* 34 (3):000–000. 5

Angrist, Joshua D and Victor Lavy. 1999. "Using Maimonides' rule to estimate the effect of class size on scholastic achievement." *The Quarterly Journal of Economics* 114 (2):533–575. 111

Autor, David H, Frank Levy, and Richard J Murnane. 2003. "THE SKILL CONTENT OF RE-CENT TECHNOLOGICAL CHANGE: AN EMPIRICAL EXPLORATION." *Quarterly Journal of Economics* 118 (4). 27

Bandiera, Oriana, Iwan Barankay, and Imran Rasul. 2010. "Social incentives in the workplace." *The Review of Economic Studies* 77 (2):417–458. 84

Baum-Snow, Nathaniel. 2007. "Did highways cause suburbanization?" *The Quarterly Journal of Economics* 122 (2):775–805. 6, 23, 55

Beaudry, Paul and Ethan Lewis. 2012. "Do Male-Female Wage Differentials Reflect Differences in the Return to Skill? Cross-city Evidence from 1980-2000." Tech. rep., National Bureau of Economic Research. 7, 27

Becker, Gary S. 1973. "A theory of marriage: Part I." *The Journal of Political Economy* :813–846. 4, 11, 12

Beckmann, Martin J. 1973. "Equilibrium models of residential land use." *Regional and Urban Economics* 3 (4):361–368. 4, 12

Belloni, Alexandre, Victor Chernozhukov, and Christian Hansen. 2014. "High-dimensional methods and inference on structural and treatment effects." *The Journal of Economic Perspectives* 28 (2):29–50. 6, 21

Bertrand, Marianne, Claudia Goldin, and Lawrence F Katz. 2010. "Dynamics of the gender gap for young professionals in the financial and corporate sectors." *American Economic Journal: Applied Economics* 2 (3):228–255. 5

Black, Dan, Gary Gates, Seth Sanders, and Lowell Taylor. 2002. "Why do gay men live in San Francisco?" *Journal of Urban Economics* 51 (1):54–76. 8

Black, Sandra E. 1999. "Do better schools matter? Parental valuation of elementary education." *The Quarterly Journal of Economics* 114 (2):577–599. 52

Black, Sandra E, Paul J Devereux, and Kjell G Salvanes. 2011. "Too young to leave the nest? The effects of school starting age." *The Review of Economics and Statistics* 93 (2):455–467. 108

Blume, Lawrence E, William A Brock, Steven N Durlauf, and Yannis M Ioannides. 2010. "Identification of social interactions." . 92

Boustan, Leah, Robert Margo, Matthew Miller, and Justin Steil. 2016. "Urban Gentrification: The Role of Condominiums." *Working Paper* . 24

Boustan, Leah and Allison Shertzer. 2013. "Population Trends as a Counterweight to Central City Decline, 1950–2000." *Demography* 50 (1):125–147. 8

Boustan, Leah Platt. 2010. "Was postwar suburbanization "white flight"? Evidence from the black migration." *The Quarterly Journal of Economics* 125 (1):417–443. 6, 24, 55

Boustan, Leah Platt, Devin Bunten, and Owen Hearey. 2013. "Urbanization in the United States, 1800-2000." Tech. rep., National Bureau of Economic Research. 54

Bramoullé, Yann, Habiba Djebbari, and Bernard Fortin. 2009. "Identification of peer effects through social networks." *Journal of econometrics* 150 (1):41–55. 92

Bramoullé, Yann, Rachel Kranton, and Martin D'amours. 2014. "Strategic interaction and networks." *The American Economic Review* 104 (3):898–930. 92

Brueckner, Jan K and Stuart S Rosenthal. 2009. "Gentrification and neighborhood housing cycles: will America's future downtowns be rich?" *The Review of Economics and Statistics* 91 (4):725–743. 6, 23

Bursztyn, Leonardo and Robert Jensen. 2015. "How Does Peer Pressure Affect Educational Investments?" *The quarterly journal of economics* :qjv021. 94

Cai, Hongbin, Hanming Fang, and Lixin Colin Xu. 2011. "Eat, drink, firms, government: An investigation of corruption from the entertainment and travel costs of Chinese firms." *The Journal of Law and Economics* 54 (1):55–78. 87

Calvo-Armengol, Antoni and Matthew O Jackson. 2004. "The effects of social networks on employment and inequality." *The American Economic Review* 94 (3):426–454. 84

Calvó-Armengol, Antoni, Eleonora Patacchini, and Yves Zenou. 2009. "Peer effects and social networks in education." *The Review of Economic Studies* 76 (4):1239–1267. 92

Card, David. 1999. "The causal effect of education on earnings." *Handbook of labor economics* 3:1801–1863. 84, 107

———. 2001. "Estimating the return to schooling: Progress on some persistent econometric problems." *Econometrica* 69 (5):1127–1160. 84, 107

Card, David, Alexandre Mas, and Jesse Rothstein. 2008. "Tipping and the Dynamics of Segregation." *The Quarterly Journal of Economics* :177–218. 6, 24

Cobb-Clark, Deborah A and Stefanie Schurer. 2012. "The stability of big-five personality traits." *Economics Letters* 115 (1):11–15. 89

Collins, William J and Robert A Margo. 2011. "Race and Home Ownership from the End of the Civil War to the Present." *The American Economic Review* 101 (3):355–359. 53

Conley, Timothy, Nirav Mehta, Ralph Stinebrickner, and Todd Stinebrickner. 2015. "Social Interactions, Mechanisms, and Equilibrium: Evidence from a Model of Study Time and Academic Achievement." Tech. rep., National Bureau of Economic Research. 85

Conti, Gabriella, Andrea Galeotti, Gerrit Mueller, and Stephen Pudney. 2013. "Popularity." *Journal of Human Resources* 48 (4):1072–1094. 85, 95, 106

Corcoran, Mary, Linda Datcher, and Greg Duncan. 1980. "Information and influence networks in labor markets." *Five thousand American families: Patterns of economic progress* 8 (S 1):37. 85

Costa, Dora L and Matthew E Kahn. 2000. "Power Couples: Changes in the Locational Choice of the College Educated." *The Quarterly Journal of Economics* 115 (4):1287–1315. 7

Couture, Victor and Jessie Handbury. 2015. "Urban revival in America, 2000 to 2010." In *American Economic Association annual meeting. https://www. dropbox. com/s/hlmnbjsx32eaafb/UrbanizationPatterns_vCurrentPublic. pdf*. 3, 24, 27, 50, 56

Cullen, Julie Berry and Steven D Levitt. 1999. "Crime, urban flight, and the consequences for cities." 24, 55

Deming, David J. 2015. "The growing importance of social skills in the labor market." Tech. rep., National Bureau of Economic Research. 27, 85

Diamond, Rebecca. 2016. "The determinants and welfare implications of US workers' diverging location choices by skill: 1980–2000." *The American Economic Review* 106 (3):479–524. 7

Edlund, Lena, Cecilia Machado, and Maria Micaela Sviatschi. 2015. "Bright minds, big rent: gentrification and the rising returns to skill." Tech. rep., National Bureau of Economic Research. 3, 24, 27, 50, 56

Elliott, James R. 2001. "Referral hiring and ethnically homogeneous jobs: How prevalent is the connection and for whom?" *Social Science Research* 30 (3):401–425. 85

Feldman, Linda, Bart Harvey, Philippa Holowaty, and Linda Shortt. 1999. "Alcohol use beliefs and behaviors among high school students." *Journal of Adolescent Health* 24 (1):48–58. 90

Fine, David A. 1980. "The Condominium Conversion Problem: Causes and Solutions." *Duke Law Journal* 1980 (2):306–335. 59

Fujita, Masahisa. 1989. *Urban economic theory: land use and city size*. Cambridge university press. 4, 12, 15, 17

Gershenson, Seth, Stephen B Holt, and Nicholas W Papageorge. 2016. "Who believes in me? The effect of student–teacher demographic match on teacher expectations." *Economics of Education Review* 52:209–224. 110

Glaeser, Edward L and Joshua D Gottlieb. 2006. "Urban resurgence and the consumer city." *Urban studies* 43 (8):1275–1299. 55

Glaeser, Edward L and Matthew E Kahn. 2001. "Decentralized Employment and the Transformation of the American City." *Brookings-Wharton Papers on Urban Affairs* 2001 (1):1–63. 8

Glaeser, Edward L, Jed Kolko, and Albert Saiz. 2001. "Consumer city." *Journal of economic geography* 1 (1):27–50. 55

Glaeser, Edward L and Jesse M Shapiro. 2003. "Urban growth in the 1990s: Is city living back?" *Journal of regional science* 43 (1):139–165. 55, 56

Goldin, Claudia. 2006. "The Quiet Revolution That Transformed Women's Employment, Education, and Family." *The American Economic Review* 96 (2):1–21. 3

Goldin, Claudia and Lawrence F Katz. 2002. "The Power of the Pill: Oral Contraceptives and Womenís Career and Marriage Decisions." *Journal of Political Economy* 110 (4):730–770. 3

———. 2008. "Transitions: Career and family life cycles of the educational elite." *The American Economic Review* 98 (2):363–369. 5

Graham, Bryan S. 2015. "Methods of identification in social networks." *Annu. Rev. Econ.* 7 (1):465–485. 92, 107

Granovetter, Mark. 1985. "Economic action and social structure: The problem of embeddedness." *American journal of sociology* 91 (3):481–510. 85

Granovetter, Mark S. 1973. "The strength of weak ties." *American journal of sociology* 78 (6):1360–1380. 85

Guerrieri, Veronica, Daniel Hartley, and Erik Hurst. 2013. "Endogenous gentrification and housing price dynamics." *Journal of Public Economics* 100:45–60. 6, 25, 50, 56

Hansmann, Henry. 1991. "Condominium and cooperative housing: transactional efficiency, tax subsidies, and tenure choice." *The Journal of Legal Studies* 20 (1):25–71. 51

Heckman, James, Rodrigo Pinto, and Peter Savelyev. 2013. "Understanding the mechanisms through which an influential early childhood program boosted adult outcomes." *The American economic review* 103 (6):2052–2086. 85

Heckman, James J. 2006. "Skill formation and the economics of investing in disadvantaged children." *Science* 312 (5782):1900–1902. 85

Heckman, James J, Jora Stixrud, and Sergio Urzua. 2006. "The effects of cognitive and noncognitive abilities on labor market outcomes and social behavior." *Journal of Labor economics* 24 (3):411–482. 85

Jackson, Matthew O. 2005. "The economics of social networks." . 92, 107

Jackson, Matthew O, Brian W Rogers, and Yves Zenou. 2016. "The economic consequences of social network structure." . 107

Jacob, Brian, Brian McCall, and Kevin M Stange. 2013. "College as Country Club: Do Colleges Cater to Students' Preferences for Consumption?" Tech. rep., National Bureau of Economic Research. 87

Johnson, William R. 2017. "Parties or Problem Sets: Review Article on How College Works and Paying for the Party." *Journal of Economic Literature* . 87

Kahn, Matthew E. 2007. "Gentrification Trends in New Transit-Oriented Communities: Evidence from 14 Cities That Expanded and Built Rail Transit Systems." *Real Estate Economics* 35 (2):155–182. 50

Kerr, Norbert L and R Scott Tindale. 2004. "Group performance and decision making." *Annu. Rev. Psychol.* 55:623–655. 84

Kleven, Henrik J, Camille Landais, and Jacob E Sogaard. 2015. "Children and gender inequality: Evidence from Denmark." *Unpublished manuscript, LSE* . 5, 23

Kuntsche, Emmanuel, Ronald Knibbe, Gerhard Gmel, and Rutger Engels. 2005. "Why do young people drink? A review of drinking motives." *Clinical psychology review* 25 (7):841–861. 90

Lenton, Pamela. 2014. "Personality characteristics, educational attainment and wages: an economic analysis using the British cohort study." *The Sheffield Economic Research Paper Series* 201401. 89

LeRoy, Stephen F and Jon Sonstelie. 1983. "Paradise lost and regained: Transportation innovation, income, and residential location." *Journal of Urban Economics* 13 (1):67–89. 6, 23

Lindenlaub, Ilse and Anja Prummer. 2014. "Gender, Social Networks And Performance." . 85

Lleras-Muney, Adriana, Matthew M. Miller, and Shuyang Sheng. 2017. "Party On: Estimating the Labor Market Returns to Social Capital." . 28

Low, Corinne. 2014. "Pricing the Biological Clock: Reproductive Capital on the US Marriage Market." . 12

Madden, Janice Fanning. 1980. "Urban land use and the growth in two-earner households." *The American Economic Review* :191–197. 8

Margo, Robert A. 1992. "Explaining the postwar suburbanization of population in the United States: The role of income." *Journal of urban economics* 31 (3):301–310. 50, 54

Mas, Alexandre and Enrico Moretti. 2009. "Peers at work." *The American Economic Review* 99 (1):112–145. 84

Mieszkowski, Peter and Edwin S Mills. 1993. "The causes of metropolitan suburbanization." *The Journal of Economic Perspectives* 7 (3):135–147. 54

Miller, Matthew M. 2015. "The Resurgence of Cities through Gentrification of their Downtown Areas." *Working Paper* . 2, 3

Miller, Matthew M. 2017. *Equal Pay Saves the City: How Women Initiated an Urban Renaissance.* Ph.D. thesis, UCLA. 56

Montgomery, James D. 1991. *Social networks and persistent inequality in the labor market.* Center for Urban Affairs and Policy Research. 84

Platt Boustan, Leah. 2012. "School desegregation and urban change: Evidence from city boundaries." *American Economic Journal: Applied Economics* 4 (1):85–108. 52

Polachek, Solomon W. 2004. "How the human capital model explains why the gender wage gap narrowed." . 3

Shenhav, Na'ama. 2015. "What Women Want: Family Formation and Labor Market Responses to Marriage Incentives." . 4, 5, 7, 11, 13, 27, 28

Stevens, Stephan, Alexander L Gerlach, and Fred Rist. 2008. "Effects of alcohol on ratings of emotional facial expressions in social phobics." *Journal of anxiety disorders* 22 (6):940–948. 90

Stock, James H and Motohiro Yogo. 2002. "Testing for weak instruments in linear IV regression." 110

Weichselbaumer, Doris and Rudolf Winter-Ebmer. 2005. "A meta-analysis of the international gender wage gap." *Journal of Economic Surveys* 19 (3):479–511. 3, 21

Weinberger, Catherine J. 2014. "The increasing complementarity between cognitive and social skills." *Review of Economics and Statistics* 96 (4):849–861. 85, 89

Weiss, Andrew. 1995. "Human capital vs. signalling explanations of wages." *The Journal of Economic Perspectives* 9 (4):133–154. 87

Welch, Finis. 2000. "Growth in women's relative wages and in inequality among men: One phenomenon or two?" *The American Economic Review* 90 (2):444–449. 7

Woolley, Anita Williams, Christopher F Chabris, Alex Pentland, Nada Hashmi, and Thomas W Malone. 2010. "Evidence for a collective intelligence factor in the performance of human groups." *science* 330 (6004):686–688. 84